Optimal Agency Bias and Regulatory Review

Ryan Bubb and Patrick L. Warren

ABSTRACT

Why do bureaucratic principals appoint agents who hold different policy views from themselves? 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. This creates an incentive for the principal to use review mechanisms that mitigate the resulting bias in the agents' decisions. The availability of such review mechanisms encourages principals to employ more extreme agents. 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 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.

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 (75 Fed. Reg. 2938 [January 10, 2010]). But the EPA's ozone proposal was quashed following a review of the policy

RYAN BUBB is Associate Professor of Law, New York University School of Law. PATRICK L. WARREN is Assistant Professor of Economics, Clemson University. We are grateful to Barry Adler, Jennifer Arlen, Oren Bar-Gill, Rachel Barkow, Charles Cameron, Stephen Choi, 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, John Patty, Rick Pildes, Ricky Revesz, Catherine Sharkey, Steve Shavell, Matthew Stephenson, Charles J. Thomas, and participants in workshops at New York University School of Law, Clemson University, Harvard Law School, Yale Law School, the University of Chicago Law School and Harris School of Public Policy, the University of California, Berkeley, School of Law, and the Political Economy and Public Law Conference at Washington University in St. Louis for helpful discussions and comments.

[*Journal of Legal Studies*, vol. 43 (January 2014)]

by the White House Office of Information and Regulatory Affairs (OIRA). In a letter to EPA administrator Lisa Jackson announcing the decision, Cass Sunstein (2011), the administrator of OIRA, explained that President Barack 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." The president, of course, had appointed both Jackson and Sunstein to their respective posts, choosing a self-described environmentalist to lead the EPA and a proponent of cost-benefit analysis to head OIRA. Moreover, President Obama issued executive orders requiring that significant rules issued by the EPA (and other executive agencies) be subject to OIRA review, continuing a practice that dates back to the Nixon administration (Exec. Order No. 13,563, 76 Fed. Reg. 3821 [January 18, 2011]; Exec. Order 13,610, 77 Fed. Reg. 28,469 [May 14, 2012]).

In this paper we provide a rationale for a bureaucratic principal appointing agents who hold different policy views from the principal's and instituting a review process led by a bureaucrat with views more aligned 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 for policy making, and generating information is costly. When responsibility for information production is delegated to an agent, but that agent does not bear the full net benefits of regulation, 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 that they make.

We show how policy bias can be harnessed to mitigate the problem of shirking. Suppose, for example, that the principal wants to find regulatory opportunities in some domain, say, the environment, but has to delegate to a single agent the responsibility for generating information about such opportunities and setting policy. To whom should the principal delegate? An agent with preferences aligned with the principal would choose the principal's preferred rules, conditional on their information, but would 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 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 (from the principal's perspective). That is, conditional on the information the agent has generated, the agent will not choose the rule (for example, stringency) preferred by the principal. We show that this trade-off generally results in the principal preferring a 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 bureaucrat more aligned with the principal. One might worry that such review would nullify the agency bias approach to incentivizing information production. With the ultimate policy decisions made according to preferences close to the principal's, 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 (Gilligan and Krehbiel 1987; Aghion and Tirole 1997). 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 a single 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 (for example, 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 identifying opportunities to reduce the level of 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.

Of course, appointment decisions involve many other political and managerial considerations, and there are alternative potential explanations for appointments of biased agents and intermediaries. For example, Warren (2012) shows that a president may choose a more extreme appointee than he would otherwise prefer in order to better resist capture of the appointee by career staff on the one hand and influence by Congress on the other. McCarty (2004) presents a related story in which a president may bias his appointee toward Congress in order to induce a larger appropriation to the agency. Similarly, a principal might choose a biased agent to counterbalance the expected effect of lobbying by an interest group. Or 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 by using a review process that is not well understood by the constituency. 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 and its interaction with regulatory review.

We make two main contributions to the formal literature on delegation. First, we posit a different, and we think better motivated, structure to agent policy preferences than the canonical spatial model that delivers a quite general preference of principals for biased agents.[1] This contrasts with the standard ally principle produced by the spatial model in which principals prefer agents whose policy preferences are aligned with their own (Bendor, Glazer, and Hammond 2001). The spatial model abstracts from what policies actually do and instead simply represents policy preferences with a utility function of the form $u(s, k) = h(|s - k|)$, where s is the policy outcome, k is the agent's ideal policy, and

1. Dewatripont and Tirole (1999) analyze a related but different phenomenon: when a principal can only use decision-based monetary rewards and cannot directly pay for information, it can be optimal to design such rewards to create "advocates" who compete by producing information favorable to their assigned cause. In their setup, the principal has to make a decision that requires information, and it is efficient to look for information in favor of each of several potential decisions. The reason it is optimal to delegate each of these tasks to a separate agent who is paid on the basis of whether his assigned decision is made is that, under that scheme, each agent's pay is monotonically increasing in his effort (which obviously has attractive incentive properties). In contrast, we focus on intrinsic motivation and consider the important case in which the structure of preferences is such that incentives to exert effort to produce information about a decision are linked to preferences regarding the decision conditional on a given set of information.

$h(\cdot)$ is a strictly decreasing function. Bendor and Meirowitz (2004) provide an illuminating theoretical framework that synthesizes much of the delegation literature based on the spatial model and point out that the ally principal can break down in the spatial model when potential agents differ along multiple dimensions. A principal may prefer a less aligned agent to a more aligned one if the less aligned agent is more competent or willing to work harder. Bendor and Meirowitz (2004, p. 301) conclude that "[e]ven with homogenous costs [of acquiring information] the ally principal need not hold because the gains of acquiring information may vary across agents, making proximate agents less willing than distant ones to [acquire information]." But why would the benefits of collecting information be systematically related to agent bias? In their framework, these differences are abstract and exogenous, and there is no general reason why biased agents would work harder than aligned agents.

We depart from the spatial model and posit an alternative, more microfounded structure to preferences that endogenously links bias and effort. In particular, we think of regulatory policies as producing social benefits, such as cleaner air, as well as social costs, such as compliance costs. We posit a form of social preferences: agents care about policy because they care about these social benefits and costs. Moreover, some agents care more about the benefits in a particular domain. Environmentalists, for example, care more about improvements to the environment, relative to their own consumption, than do nonenvironmentalists. This leads to a policy utility function of the form $u(s, k) = kb(s) - c(s)$, where $b(s)$ and $c(s)$ represent some fraction of the social benefits and costs of policy s, respectively, and the agent's type k determines the weight that the agent puts on benefits. We show that, unlike the spatial model, this approach to modeling policy preferences produces an endogenous link between bias and effort that is robust to many different information production and policy choice problems.[2]

Our results contrast with existing work analyzing delegation to biased agents using the spatial model. Gailmard and Patty (2007) provide a delegation model based on spatial preferences in which the principal—a legislature in their interpretation—prefers to induce policy-motivated "zealots" rather than non-policy-motivated "slackers" to self-select into civil service. The reason is that zealots are willing to invest in expertise, since that allows them to better calibrate policy, whereas slackers are

2. See Section 3 for a more detailed discussion of the differences between our approach to modeling policy preferences and the canonical spatial model of policy preferences.

not since they, by assumption, do not care about policy. Since the principal does better with an expert agent than a nonexpert agent, the principal prefers to get zealots despite their exogenously assumed bias. Unlike in our model, bias here has only costs and no benefits. The principal would prefer an aligned policy-motivated agent to a biased policy-motivated agent, but the principal is assumed to be unable to directly control the policy preferences of bureaucrats. Rather, all policy-motivated agents are assumed to have a uniform exogenous degree of policy bias. The models of Gilligan and Krehbiel (1987) and Aghion and Tirole (1997) similarly show that a principal can motivate a biased agent to work harder by delegating authority but also share the feature that the principal would prefer an agent who is perfectly aligned with the principal. In contrast, we model the appointment and delegation decisions of the principal and show why the principal will actively appoint biased agents but subject their decisions to review by more aligned bureaucrats.

Prendergast (2007) develops a model in which the principal does actively prefer a biased agent to an aligned agent. He analyzes the problem of inducing street-level bureaucrats to identify the proper recipients for some treatment (for example, a driver's license or a prison sentence)—a setting quite different from the policy-making bureaucracy considered here. However, he considers delegation of only the search task, not the ultimate allocation decision, and hence does not analyze the trade-off between policy and effort that is central to our analysis. Accordingly, bias in his model has only benefits and no costs, and therefore maximally biased agents are always optimal, unlike in our model.

Our second contribution is in analyzing how this use of agency bias interacts with the institution of regulatory review and, more generally, with the opportunity to choose an (also potentially biased) intermediary. 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. Lewis (2008, p. 2) defines politicization as "the act of increasing the number and penetration of [political] appointees." Moe and Wilson (1994, p. 18) 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 decision-making authority to further rein in agencies' residual noncompliance with presidential policy objectives. On this standard account, then, politicization and centralization are substitutes. The classic formal model of political control of agencies by Calvert, McCubbins, and Weingast (1989) takes a similar approach, with centralized control useful only because of uncertainty about the preferences of appointees ex ante. Many subsequent formal models of regulatory review take the policy preferences of agencies as exogenous (for example, 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 policy-making agent. In Section 3 we consider two extensions of our baseline model: asymmetric information between the agent and the reviewer and an existing regulation that the principal wants to revise. We also discuss the generalizability of our results to different types of information production and to different assumptions about the structure of policy preferences. In Section 4 we illustrate the application of our analysis to administrative decision-making institutions using two historical examples from the Nixon administration: the revitalization of the Federal Trade Commission and the creation of the EPA 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

Consider a setting with potential regulatory opportunities that 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 that 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. 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.

We focus on two design issues: the type of bureaucrats that the principal will appoint and whether to appoint a separate bureaucrat to review rules proposed by the agency. We begin with an analysis of the principal's choice of a single agent to both generate information about regulatory opportunities and choose policy conditional on that information. We refer to this model as full delegation, and it is meant to illustrate the fundamental trade-off associated with choosing a biased agent. We then consider a more flexible model in which the principal can delegate to one agent the information production task and to a different agent the policy choice. This second model, which we term regulatory review, shows how the use of agency bias interacts with institutional mechanisms designed to check that bias ex post.

2.1. Setup

A bureaucrat at the agency can generate information about regulatory opportunities within its purview by exerting costly effort to search. We 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, $\psi(e) = Ce^2/2$, with values of C 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 level $s \geq 0$. Think of stringency as how tightly the regulation controls the pollutant. A higher stringency level 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 a player of type k has gross benefit kBs and gross cost $c(s)$ from a regulation of stringency level s. 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.[3] The greater is k, the more the player cares about the benefits of the regulation. Think of k as measuring how pro-regulatory or mission oriented (where the mission is defined in terms of regulatory

3. If preferred, that fraction could be explicit in the utility function, for example, $U_A(e, s|k_A) = \gamma(k_A Bs - s^2/2) - \psi(e)$, 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.

benefits, for example, environmental protection) the player is. The benefits that the players care about include, for example, a reduction in respiratory diseases in society, while the costs include the cost of installation of equipment at power plants to control the levels of the pollutant. For simplicity, we assume that $c(s)$ takes a quadratic form so that $c(s) = s^2/2$.

We discuss in Section 3 in some detail the extent to which the results we derive under these assumptions about the information production problem and the structure of preferences are generalizable to other settings. In brief, our information production assumptions are not important; similar results obtain, for example, in a model in which a regulatory opportunity has already been identified and the agent searches for information about the marginal benefit of the opportunity. In contrast, our assumptions about the structure of policy preferences are important; in particular, our results do not generally hold if agents have spatial preferences.

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

$$U_{\mathrm{A}}(e, s) = k_{\mathrm{A}}Bs - \frac{s^2}{2} - C\frac{e^2}{2} \tag{1}$$

and

$$U_{\mathrm{P}}(s) = k_{\mathrm{P}}Bs - \frac{s^2}{2}, \tag{2}$$

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

The principal faces two incentive problems posed by delegation to the agent. First, the agent and the principal may put a different weight k_j on the benefits of regulation and hence may have conflicting policy preferences. We assume that k_{A} is in the interval $[0, k^{\max}]$ and that k_{P} is in the interval $(0, k^{\max})$ so that it is always possible for the agent to be strictly less or strictly more pro-regulatory than the principal.

Second, 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 induce first-best effort levels, perhaps because of difficulty in measuring bureaucratic effort and output.[4]

4. Of course, in general extrinsic motivations such as the desire for promotion produce

Suppose that the principal has authority to appoint the agent and that k_A is observable. What type of agent will the principal appoint?

2.2. Full Delegation to the Agent

We begin with full delegation. The sequence of moves in the model is that the principal appoints the agent by choosing k_A; the agent chooses search effort e; with probability e, the agent finds a regulatory opportunity (if not, the game ends); and the agent chooses stringency level 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 level according to her preferences. We assume that the principal can commit to delegating policy authority and discuss the case without the ability to perfectly commit in our discussion of regulatory review in Section 2.3.

2.2.1. Stringency Level and Search Effort. We find the equilibrium of the model by starting with the agent's choice of stringency level. The agent chooses s to solve

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

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

$$s^*(k_A) = k_A B. \tag{4}$$

Note that the agent's choice of stringency level 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 level $k_P B$. Thus, the agent chooses the principal's preferred stringency level for any regulatory opportunity he finds if and only if $k_A = k_P$. The value to a player of type k_i from a regulatory policy implemented by an agent of type k_j is given by

$$V(k_i, k_j) = k_i B s^*(k_j) - \frac{s^*(k_j)^2}{2} = \left[k_i k_j - \frac{k_j^2}{2} \right] B^2. \tag{5}$$

Turn now to the agent's search effort. The agent chooses effort e to solve

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.

$$\max_{e \in [0,1]} \left\{ eV(k_A, k_A) - C\frac{e^2}{2} \right\}. \tag{6}$$

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

$$e^*(k_A) = \frac{V(k_A, k_A)}{C} = \frac{k_A^2 B^2}{2C}. \tag{7}$$

Note that e^* increases in k_A. The principal thus faces a trade-off between incentives for effort provision and choice of stringency level. An agent who shares the principal's policy preferences, $k_A = k_P$, will choose the principal's preferred stringency level. 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 level. 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—the level of stringency is set too high. The key reason this trade-off exists is that increasing the weight that the agent puts on the gross benefits of regulation increases both the marginal benefit of stringency level (which increases the agent's optimal choice of stringency level) and the level of the agent's payoff from finding information about a regulatory opportunity (which increases his level of effort).

2.2.2. Agent Bias. Consider now the principal's optimal choice of agent bias given this trade-off. The principal solves the problem

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

Denote the solution by k^*.

`Proposition 1: Equilibrium under Full Delegation.` The principal chooses a relatively pro-regulatory agent, $k^* = \min\{\frac{3}{2}k_P, k^{\max}\} > k_P$.

Proof. All proofs are in the online appendix.

Proposition 1 characterizes how a principal will use her 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 toward 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 large increase in the agent's

value of finding information about a regulatory opportunity. Increasing k_A above $k_A = k_p$ thus increases the agent's search effort. In contrast, it causes only a relatively small decrease in the principal's policy payoff from the agent finding a regulatory opportunity since the policy is very close to the principal's ideal point, and the principal's payoff is concave. Formally, by the envelope theorem, the utility loss to the principal from moving stringency level away from the principal's ideal point is only second order, while the increase in the agent's value of a regulatory opportunity (and hence effort) is first order.[5]

Finally, note that the proposition establishes that k^* is increasing in k_P, strictly so if $k^{max} > \frac{3}{2}k_P$. 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 instead of choosing the stringency level, 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; that is, $s = 0$. If the agent does propose a rule, then the reviewer gets to choose the rule's stringency level.[6]

The reviewer's policy preferences have the same basic structure as those of the principal. In particular, the reviewer weights the gross benefits of regulation by k_R, with $k_R \in [0, k^{max}]$, so the reviewer's preferences over the ultimate policy decision can be represented by the utility function

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

The sequence of moves in the model is now that the principal appoints

5. Of course, there are characteristics of the agent other than the weight that the agent puts on the benefits of regulation that would be important to a principal and that we omit from the model, but incorporating them would not change the basic point. For example, ceteris paribus, the principal would prefer an agent who has a lower cost of effort (for example, because he is more intelligent) or is more public spirited (that is, the γ in note 3). But holding constant these other characteristics, say, at a high level, our model shows that the principal would prefer a biased, very smart, public-spirited agent to an unbiased, very smart, public-spirited agent.

6. Our assumption that the reviewer gets to set the level of 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 note 7.

agent and reviewer by choosing k_A and k_R; the agent chooses search effort e; with probability e, the agent finds a regulatory opportunity (if not, the game ends); the agent chooses whether to propose a rule to the reviewer (if not, the game ends); and the reviewer chooses stringency level s. Note that the full-delegation model above is equivalent to the regulatory review model under the restriction that $k_R = k_A$.

2.3.1. Stringency Level and Search Effort. We begin with the subgame in which the reviewer chooses stringency level. The reviewer's optimal choice of stringency level 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 level is chosen according to preferences k_R, is given by $V(k_A, k_R)$, the function defined in equation (5). 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 $e = V(k_A, k_R)/C$. Otherwise, the agent will exert no effort and find no regulatory opportunities. Denote the agent's optimal search effort by $e^*(k_A, k_R)$.[7]

2.3.2. Agent and Reviewer Bias. The principal's problem is to choose the agent and reviewer to optimally trade off policy bias and effort. Formally, she solves

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

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

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 principals who are more pro-regulatory appoint reviewers who are more pro-regulatory, $\partial k_R^*/\partial k_P > 0$.

7. If we instead assume that the reviewer has only the right to veto the agent's proposal, and not the right to set stringency level, 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 level of 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 level of stringency to be implemented by choosing the reviewer. In the baseline model she implements stringency level s^* by choosing $k_R = s^*/B$, while in the alternative model she does it by choosing 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, which 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.

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 searching and policy making are separated. The agent's search effort increases in his bias, and with policy controlled by the reviewer, there is no policy cost to the principal from having a biased agent.

Nevertheless, the trade-off between policy and effort remains at a different level. The agent is willing to work even harder as the reviewer's policy preferences become more similar to the agent's. 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. This is once again because of the envelope theorem: increasing k_R from the principal's ideal point at $k_R = k_P$ causes only a second-order decrease in the principal's policy payoff conditional on finding an opportunity but a first-order increase in the agent's policy payoff and effort. Similarly, $k_R = k^{max}$ is also never optimal, since moving the reviewer slightly toward the principal from that point causes a first-order increase in the principal's policy payoff but only a second-order decrease in the agent's effort.

Proposition 2.3 says that the presence of regulatory review reduces equilibrium policy bias. To see the intuition for why, begin with the observation that under full delegation, at $k_A = k^*$, the marginal benefit of increasing k_A (from increased effort) is equal to the marginal cost (from worse policy). Now consider a principal who chooses a reviewer and an agent with the same preference parameter as the optimal agent under full delegation, $k_A = k_R = k^*$. The principal can obviously achieve greater utility by moving just the agent out to $k_A = k^{max}$, which increases the agent's effort at no cost to policy, which remains set by the reviewer. The key question is, how does this change affect the marginal costs and benefits of increasing k_R relative to the marginal costs and benefits of adjusting k_A at $k_A = k^*$ under full delegation? The marginal cost of increasing k_R, stemming from worse policy choices when an opportunity is found, has now gone up because the agent is finding more regulatory opportunities relative to when $k_A = k^*$ under full delegation. The effect on the marginal benefit, stemming from increased effort, is ambiguous

but always less than the effect on the marginal cost. Because of this, the principal prefers to reduce k_R from k^*.

Finally, the principal always strictly prefers to have regulatory review rather than to delegate fully to the agent. This is a corollary of proposition 2.1 and 2.2. The principal could replicate the full-delegation outcome under regulatory review by choosing $k_R = k_A = k^*$ but chooses not to. Substantively, 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, we see that this move alone increases agent effort, resulting in more regulatory opportunities discovered, and increases the principal's payoff. Second, as explained above, the principal will choose policy under review that is closer to her ideal point, which further increases her equilibrium payoff.

This fact highlights the importance of the principal's ability to commit to delegating authority over stringency level. Full commitment 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.[8] However, in other settings the principal may not be able to commit to not intervening ex post. That state of affairs is just a special case of the regulatory review model considered above with k_R constrained to be equal to k_P, and the principal's inability to commit policy authority to a reviewer with an ideal point different from her own obviously makes the principal worse off.[9] Otherwise, the no-commitment case produces results similar to those of the general regulatory review case, with the principal choosing a maximally biased agent to increase effort at no cost to policy, since policy is controlled by the principal.

8. Aghion 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 rule making under Office of Information and Regulatory Affairs (OIRA) review (Livermore and Revesz 2012).

9. The complete lack of commitment captured in this example could be easily weakened to imperfect commitment with little change to our conclusions. Imagine, for example, that the principal could seize control back after the regulation was proposed at some cost $K > 0$. Essentially, this cost places a limit on the bias of a reviewer whom the principal could credibly select. Whichever k_R^{**} satisfies $[V(k_P, k_P) - V(k_P, k_R^{**})] = K$ with equality is that limit, and the principal would, in equilibrium, choose the minimum of k_R^* and k_R^{**} and never actually seize control. Of course, she would continue to choose $k_A^* = k^{\max}$ throughout.

The size of the utility gain for the principal from regulatory review depends in intuitive ways on the parameters. The principal's gains from regulatory review are proportional to B, since each extra unit of effort or extra degree of policy congruence is more valuable to the principal when B is larger. The gains from regulatory review also increase in $k^{\max}$. Since the key advantage of regulatory review is the ability to appoint a more extreme agent without bearing a policy cost, this advantage increases as $k^{\max}$ increases. These two comparative statics suggest that if using regulatory review imposes an unmodeled cost to the principal, and that cost is not affected by B or $k^{\max}$, we should see review employed more for more important policies and for policies for which it is possible to recruit agents with very extreme preferences. Importantly, this is true not because the principal is afraid of attracting those extreme agents by accident but rather because she wants to recruit them and offset them.

Finally, an implication of the model is that the availability of regulatory review encourages greater use of agency bias as a motivational instrument. The complementarity between centralization and politicization generated by our model is in contrast to the standard account, described in the Introduction, of the two as substitutes.

3. EXTENSIONS

In our baseline model above, we abstract 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 generalizability of our results to other settings.

3.1. Asymmetric Information and Regulatory Review

In our baseline model, we assume 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 that the risk that the agent will strategically withhold information curbs the principal's willingness to appoint biased agents.

3.1.1. Setup. To be concrete, suppose that if the agent finds a regulatory opportunity, it is one of two types: a high-value opportunity with marginal benefit B_H with probability q or a low-value opportunity with marginal benefit B_L with probability $1 - q$, with $0 < B_L < B_H$. However, the agent learns what type the opportunity is only 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 that the agent chooses search effort e; with probability e, the agent finds a regulatory opportunity (if not, the game ends); nature chooses marginal benefit B of the regulatory opportunity; with probability p, the agent learns B; the agent chooses whether to propose a rule to the reviewer (if not, the game ends); if the agent knows B, he chooses whether to disclose B to the reviewer; and the reviewer chooses stringency level s. Our equilibrium concept is perfect Bayesian equilibrium.

3.1.2. Stringency Level, Disclosure, and Search Effort. To find the equilibrium, we start with the reviewer's choice of stringency level. If the agent has conveyed the marginal benefit of the regulatory opportunity to the reviewer, then the reviewer's choice problem is

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

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, expression (11) is also the reviewer's

10. Dessein (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 reviewer 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 choose someone like herself. We allow the principal control over both but with an effort-policy trade-off.

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 $s^*(k_R, B) = k_R B$.

Next consider the agent's disclosure strategy, for now taking the preference parameters k_A and k_R as given. 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 level $k_R B_L$, which is lower than the agent's preferred stringency level $k_A B_L$. 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 $\overline{B} = qB_H + (1 - q)B_L$. Thus the reviewer would choose a higher level of stringency, $k_R \overline{B} > k_R B_L$. For this deviation not to be attractive to the agent, hiding must result in the reviewer overshooting by selecting a level of stringency that is above the agent's ideal stringency level of $k_A B_L$. In fact, this implemented stringency level must be so far above the agent's ideal that the agent prefers the too-low stringency level 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 level, which leads 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 limits 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 $-[k_R(1 - q)/2](B_H - B_L/B_H) \leq k_A - k_R \leq (k_R q/2)(B_H - B_L/B_L)$.

Let $\overline{k}(k_R) \equiv k_R + (k_R q/2)(B_H - B_L/B_L)$ 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

$$
\begin{aligned}
V^{\text{Disc}}(k_{\text{A}}, k_{\text{R}}) &= qp\left[k_{\text{A}}B_{\text{H}}s^*(k_{\text{R}}, B_{\text{H}}) - \frac{s^*(k_{\text{R}}, B_{\text{H}})^2}{2}\right] \\
&\quad + (1-q)p\left[k_{\text{A}}B_{\text{L}}s^*(k_{\text{R}}, B_{\text{L}}) - \frac{s^*(k_{\text{R}}, B_{\text{L}})^2}{2}\right] \\
&\quad + (1-p)\left[k_{\text{A}}\overline{B}s^*(k_{\text{R}}, \overline{B}) - \frac{s^*(k_{\text{R}}, \overline{B})^2}{2}\right] \qquad (12) \\
&= \left[k_{\text{A}}k_{\text{R}} - \frac{k_{\text{R}}^2}{2}\right]\{p[qB_{\text{H}}^2 + (1-q)B_{\text{L}}^2] + (1-p)\overline{B}^2\} \\
&= \left[k_{\text{A}}k_{\text{R}} - \frac{k_{\text{R}}^2}{2}\right][\overline{B}^2 + pq(1-q)(B_{\text{H}} - B_{\text{L}})^2].
\end{aligned}
$$

The agent chooses search effort to satisfy the first-order condition $e = \max\{0, V^{\text{Disc}}(k_{\text{A}}, k_{\text{R}})/C\}$.

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 values of k_{A}, much as in the baseline model, the relevant case is $k_{\text{A}} > k_{\text{R}}$, 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 $k_{\text{A}} < k_{\text{R}}$.

Denote by $\hat{q}$ the reviewer's beliefs in such an equilibrium about the probability that $B = B_{\text{H}}$ in the subgame in which the agent does not disclose, and represent the corresponding expected marginal benefit by $\hat{B}$. The reviewer's equilibrium beliefs are given by Bayes's rule:

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

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

$$
\begin{aligned}
V^{\text{Hide}}(k_{\text{A}}, k_{\text{R}}) &= qp\left[k_{\text{A}}B_{\text{H}}s^*(k_{\text{R}}, B_{\text{H}}) - \frac{s^*(k_{\text{R}}, B_{\text{H}})^2}{2}\right] \\
&\quad + (1-qp)\left[k_{\text{A}}\hat{B}s^*(k_{\text{R}}, \hat{B}) - \frac{s^*(k_{\text{R}}, \hat{B})^2}{2}\right] \qquad (14) \\
&= \left[k_{\text{A}}k_{\text{R}} - \frac{k_{\text{R}}^2}{2}\right][qpB_{\text{H}}^2 + (1-qp)\hat{B}^2] \\
&= \left[k_{\text{A}}k_{\text{R}} - \frac{k_{\text{R}}^2}{2}\right]\left[\overline{B}^2 + \frac{qp(1-q)^2}{1-qp}(B_{\text{H}} - B_{\text{L}})^2\right].
\end{aligned}
$$

The agent chooses search effort $e = \max\{0, V^{\text{Hide}}(k_{\text{A}}, k_{\text{R}})/C\}$.

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. To see this, observe that $V^{\text{Disc}}(k_A, k_R) > V^{\text{Hide}}(k_A, k_R)$. This implies that the agent's search effort is greater 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 $V^{\text{Disc}}(k, k_R)$ and $V^{\text{Hide}}(k, k_R)$ but with k_R and k_P, respectively, substituted for the first argument of the functions. Thus, 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 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.

Proposition 3: Equilibrium under Asymmetric Information.

1) Holding $\overline{B}$ 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 weakly prefers to disclose the marginal benefit for low-marginal-benefit opportunities, $k_A^{**} = \min\{\overline{k}(k_R^{**}), k^{\max}\}$, 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 $\overline{B}$.

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

12. Mixed-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.

3) The principal strictly prefers a regulatory process that includes regulatory review to full delegation to the agent.

Proposition 3.1.a 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 (the value of $B_H - B_L$ is small), then the principal chooses an extremely 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 from 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 chooses bureaucrats who will fail to communicate only when the cost associated with the loss of information is small. If the asymmetric information is important (the value of $B_H - B_L$ is large), then proposition 3.1.b shows that the principal will forgo the extra effort she might get by appointing a very extreme agent and instead chooses an agent more in line with the reviewer in order to guarantee disclosure. Large values of $B_H - B_L$ make 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 level. When the value of $B_H - B_L$ is large, there is a large gap between the preferred stringency level in each state. Second, as the value of $B_H - B_L$ grows, it 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).

Proposition 3.2 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 regarding the agency's domain. Thus, one might think that a Republican president will receive more disclosure from national security agencies, while Democrats will receive more disclosure from the EPA.

Furthermore, the fact that T is increasing in $\overline{B}$, along with proposition 3.1, 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 (5 U.S.C. sec. 551) provides interested parties notice of rule making and an opportunity to comment on proposed rules and requires agencies to state the "basis and purpose" for their decisions 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 analyses with their rules is as a way of forcing the agencies 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 that can reduce the information asymmetry by providing information.

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 strictly prefers not to. The reasons that 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 that 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 its current level of stringency is $s^*(k_O) = k_O B$ (where O is old). We assume that $k_P \neq k_O$, so the principal knows that 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 level 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 $\psi(e) = Ce^2/2$. If the agent finds an opportunity, he can propose a rule to the reviewer, who then chooses stringency level s.[13] This model nests the baseline regulatory review model from Section 2 in the particular case in which $k_O = 0$.

3.2.2. Stringency Level and Search Effort. Beginning with the reviewer's choice of stringency level, we note that if the agent proposes a revision of an existing regulation, the reviewer sets stringency level at $s^*(k_R) = k_R B$. Given this stringency-setting strategy, the agent receives the following incremental payoff over his policy payoff from the status quo from finding and proposing a revision:

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

The agent's policy payoff from uncovering an opportunity to revise regulation is nonnegative as long as the reviewer's preference is closer to his own than the old regulation is; that is, $|k_R - k_O| \leq |k_A - k_O|$. 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 $e = V(k_A, k_R)/C$ and propose a revision when he finds an opportunity. 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:

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

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 4: Equilibrium with Revision of Existing Regulations.

1) If $k_O < k_P$, the equilibrium agent's preference is $k_A^* = k^{\max}$. If $k_O > k_P$, the equilibrium agent's preference is $k_A^* = 0$.

2) The principal appoints a reviewer with preferences strictly between her own and the agent's.

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

3) A more pro-regulatory reviewer is selected as the principal becomes more pro-regulatory or the old regulation becomes less extreme; that is, $(\partial k_R^*/\partial k_P) > 0$, $(\partial k_R^*/\partial k_O) < 0$.

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 in which 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 a trade-off between effort and policy. Proposition 4.3 shows that the preferences of the reviewer respond to changes in the 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 he wants to move policy toward more stringent regulation. Because he 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 him but still biased toward 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 overregulation 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 agent 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 an environment with multiple regulations (or regulatory opportunities), little changes as long as the principal wants to move the stringency level of all regulations in the same direction. The principal still prefers an extreme 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 multitask 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.[14]

This multitask 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).

14. One 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 by-product also produce information about deregulatory opportunities, and the reverse is also true.

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 level 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.

3.3. Other Settings

We now consider how generalizable our results are, focusing on different types of information production and alternative assumptions about the structure of preferences. To do so, 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 $\psi(e) = Ce^2/2$. For simplicity, assume that if the agent finds information, the reviewer learns it as well. 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 $V^{\mathrm{I}}(k_{\mathrm{A}}, k_{\mathrm{R}})$ and $V^{\mathrm{N}}(k_{\mathrm{A}}, k_{\mathrm{R}})$, respectively. The agent chooses effort e to solve

$$\max_{e \in [0,1]} \{V^{\mathrm{N}}(k_{\mathrm{A}}, k_{\mathrm{R}}) + e[V^{\mathrm{I}}(k_{\mathrm{A}}, k_{\mathrm{R}}) - V^{\mathrm{N}}(k_{\mathrm{A}}, k_{\mathrm{R}})] - \psi(e)\}. \tag{17}$$

The solution is $e^*(k_{\mathrm{A}}, k_{\mathrm{R}}) = [V^{\mathrm{I}}(k_{\mathrm{A}}, k_{\mathrm{R}}) - V^{\mathrm{N}}(k_{\mathrm{A}}, k_{\mathrm{R}})]/C$. Thus, the agent's effort is increasing in k_{A} if and only if $V_1^{\mathrm{I}}(k_{\mathrm{A}}, k_{\mathrm{R}}) - V_1^{\mathrm{N}}(k_{\mathrm{A}}, k_{\mathrm{R}}) > 0$. In economic terms, this condition simply requires that agents value information more the greater is their k_{A}.[15]

3.3.1. 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.

For example, suppose that the agency has already identified a regulatory opportunity, but the marginal benefit of the opportunity $B \geq 0$ is unknown and distributed according to the density function $f(B)$ with

15. In our baseline model, we had $V^{\mathrm{N}}(k_{\mathrm{A}}, k_{\mathrm{R}}) = 0$, so this condition amounted to $V_1^{\mathrm{I}}(k_{\mathrm{A}}, k_{\mathrm{R}}) > 0$, which is true in that model.

expected value $E(B)$, and the agent searches for the true marginal benefit B. It can easily be shown that in this model, $V^{N}(k_A, k_R) = (k_A k_R - k_R^2/2)E(B)^2$ and $V^{I}(k_A, k_R) = (k_A k_R - k_R^2/2)[\mathrm{Var}(B) + E(B)^2]$, so $V_1^{I}(k_A, k_R) - V_1^{N}(k_A, k_R) = k_R \mathrm{Var}(B) > 0$. The agent's effort thus increases in k_A in this general setting as well, and it can be shown that all of our results in proposition 2 from the baseline model also hold in this model. And in the case of full delegation, the principal prefers an agent biased toward regulation ($k^* > k_P$), as in the baseline model. Moreover, similar results can be derived for settings in which the agent searches for information about the costs of regulation.[16]

However, there are other environments in which agents who are more pro-regulatory put a lower value on information and hence work less hard to generate it. One important example considered above is for information required to deregulate.[17] To give another example, consider the model of investigating B just introduced, but suppose now that there is an upper bound on stringency level, $s \leq \bar{s}$. Suppose that if the principal were not informed about B, then she would prefer $s = \bar{s}$, but if the principal knew that $B < \bar{s}/k_P$, then the principal would prefer $s = k_P B < \bar{s}$. It can easily be shown in this model that $V_1^{I}(k_A, k_R^*) - V_1^{N}(k_A, k_R^*) < 0$, so agents who put less weight on the gross benefits of regulation will work harder to generate information about B, for reasons similar to the case of deregulatory opportunities considered above.

3.3.2. Other Assumptions about Policy Preferences. In each of these applications, our assumption about the structure of policy preferences is important for our results. To see this, consider the contrasting results of Che and Kartik (2009), who analyze a model of delegation in which, in the case of public information, differences in preferences do not motivate the agent in the absence of a difference in opinion. This result is a consequence of their use of a quadratic loss function to represent policy preferences. In our notation, they assume that $u(s, k, B) = -(s - k - B)^2$. If we assume this structure for policy preferences in our model of investigating B, we get $V^{N}(k_A, k_R) = -(k_R - k_A)^2 - \mathrm{Var}(B)$ and $V^{I}(k_A, k_R) = -(k_R - k_A)^2$, so $V^{I}(k_A, k_R) - V^{N}(k_A, k_R) = \mathrm{Var}(B)$, which is not a

16. For example, suppose that $c(s) = L(s^2/2)$, but L is uncertain and the agent searches for the true L.

17. Mapping that extension into the framework introduced here, we find that $V^{I}(k_A, k_R) - V^{N}(k_A, k_R) = B[(k_A k_R - k_R^2/2) - (k_A k_O - k_O^2/2)]$ and $V_1^{I}(k_A, k_R) - V_1^{N}(k_A, k_R) < 0$ if and only if $k_P < k_O$.

function of k_A, and hence policy bias is not a useful motivational instrument.

The quadratic loss function in Che and Kartik (2009) is a commonly used version of the canonical representation of policy preferences in the political science literature—the spatial model—in which utility is a decreasing function of the distance between the agent's ideal point and the policy outcome, $u(s, k, B) = h(|s - k - B|)$, where $h'(\cdot) < 0$. In the spatial model, the preference parameter k shifts only the location of the ideal point, not the payoff that the agent receives from achieving the ideal policy. One consequence is that in the case of full delegation in the model of investigating B, bias is not a useful incentive instrument, since the agent's effort is independent of the preference parameter k_A.[18] Thus, in that setting the ally principle holds. In the case of regulatory review, a quadratic loss function is a special case in which again agent bias is not a useful motivational instrument.[19] These results illustrate that the spatial model of policy preferences does not generate a general preference of the principal for a biased agent, even when the agent must exert costly effort to generate information about the policy choice.

We take a different approach to modeling the structure of policy preferences. Our view is that some agents put greater value on the ultimate ends of the policy than do others. Environmentalists, for example, put greater value on improvements in the environment, relative to their own consumption, than do nonenvironmentalists. That approach leads to a specification of utility such as ours, $u(s, k, B) = kb(s) - c(s)$ (where we assume the functional forms $b(s) = Bs$ and $c(s) = s^2/2$), in which a shift in the preference parameter k not only shifts agents' ideal point but also changes the weight that they put on policy outcomes relative to other goods, such as the costs of effort. Unlike the spatial model, our approach results in policy bias having useful incentive properties across a broad domain of information production problems and institutional settings.

18. To see the intuition for why, define $\tilde{s} = s - k_A$. The policy choice problem expressed by substituting $\tilde{s}$ for $s - k_A$, when the agent is informed and uninformed, is identical for any agent type k_A. This implies that V^I and V^N are independent of k_A and therefore that agent effort is not a function of k_A.

19. However, a more concave form of spatial preferences, for example, $u(s, k, B) = -(s - k - B)^4$, can generate an incentive effect of policy bias in the case of regulatory review.

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 agents to generate information about regulatory opportunities but will 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 decision-making process—appears in many contexts across public bureaucracies. Our main application is to presidential appointments and centralized regulatory review. To mitigate shirking by a regulatory agency, the president may appoint an official with biased policy preferences to head the agency 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 of 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 bureaucrats, typically thousands, most of whom are career civil servants and some of whom are political appointees. The senior political appointees have a range of tools by which they can shape the preferences of subordinate staff, such as reorganizations and, of course, hiring decisions. Our analysis shows why the heads of agencies will employ agency staff with biased policy preferences to generate information about regulatory opportunities, especially when they can effectively review the policy decisions of such staff—either themselves or through their own intra-agency review office. An implication of our model, then, is that bureaucrats should be more biased toward the mission of the agency as one moves down the agency's hierarchy. The agency head is chosen to be biased relative to the president. The agency head then hires subordinate staff who are biased relative to the agency head, and so forth.

Our analysis also sheds light on the relationship between the legislative branch and the bureaucracy. For example, Congress plays the role of 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.

While we focus on public bureaucracies, the phenomena we analyze appear in other types of organizations, including private firms, as well. For example, consider the research division at a pharmaceutical company that is responsible for identifying new drug candidates. Suppose that the scientists whom the company could hire vary in their social preferences such that some place greater value on the health benefits or scientific novelty of new drugs than do others. If incentive contracts are imperfect, the company may prefer to hire these intrinsically motivated scientists because, ceteris paribus, they will work harder. However, they may propose drug candidates that while socially or scientifically valuable are not particularly profitable for the company. Hence, the company may want to utilize a more profit-focused manager at the research division to filter the drug candidates proposed by intrinsically motivated scientists.[20]

To make our discussion of applications more concrete, we now turn to 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. The FTC example illustrates two special cases of the model, full delegation and delegation without policy commitment, while the EPA example illustrates the complementarity between agency bias and 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 two different principals used to control the FTC: the president 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

20. Henderson (1994) describes this phenomenon in some leading pharmaceutical firms that, in the early 1990s, implemented a science-based approach to drug discovery, hiring publication- and peer-review-oriented scientists in place of drug "hunters." They coupled these intrinsically motivated scientists with pro-publication management that tried to encourage and reward both basic research and drug discovery.

president, then, this entails full delegation to the agency to set policy.[21] In contrast, 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 policy-making authority to them. The chairman, with the other FTC commissioners, retains ultimate authority to set policy on the basis of information generated by the staff, just as in the case of our model in which the principal delegates the search task to an agent but cannot commit to delegating policy authority.

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 it, one by Ralph Nader's students (Cox, Fellmeth, and Schulz 1969)[22] and the other by the American Bar Association (1969), in the fall of that year President Richard Nixon announced that "the time has now come for the reactivation and revitalization of the FTC" (Nixon 1969, p. 887). 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" (p. 887).

Key to Nixon's approach was the appointment of Caspar Weinberger as chairman of the FTC. Nixon reported that Weinberger had "assured me that he intends to initiate a new era of vigorous action" (Nixon 1969, p. 887). Personnel was a major focus during Weinberger's brief

21. The president's appointments to the Federal Trade Commission (FTC) are subject to the advice and consent of the Senate. We do not model these interactions formally, but we could think of $k^{\max}$ (and maybe even a $k_A^{\max} > 0$) as the most extreme candidates who could win confirmation. Congress's incentives to prefer extremist appointments will mirror those analyzed in the case of the president, with a bias-effort trade-off leading to a preferred appointee who is more pro-mission than Congress itself is. Since Congress and the president will, in general, want to strike this trade-off differently, there will likely be some bargaining over the exact equilibrium bias.

22. 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 members of Congress as a source of patronage employment.

6-month chairmanship, during which he discharged 18 of the top 31 staff members at the FTC (Clarkson and Muris 1981, p. 4). To succeed Weinberger, Nixon appointed Miles Kirkpatrick, 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" (Harris and Milkis 1996, p. 167). 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 (U.S. Senate 1977, p. 217).

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 (Harris and Milkis 1996, p. 181). The second stage began in 1975, when Congress expressly delegated to the FTC the authority to issue industrywide rules to regulate unfair and deceptive industry practices (Magnuson-Moss Warranty–FTC Improvement Act of 1975 [15 U.S.C. sec. 2301]). Prior to this legislation, the FTC's legal authority to promulgate industrywide rules was broadly doubted and only sparingly asserted.[23] The FTC used the authority granted it under the statute to propose rules regulating numerous industries, products, and fields, including eyeglasses, franchising, funeral homes, 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" (U.S. Senate 1973, p. 25). 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

23. The first judicial recognition of the FTC's authority to issue industrywide rules came in *National Petroleum Refiners Association 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.

to us that you are going too far. That would be the day, as far as I am concerned" (U.S. Senate 1973, p. 31).

But while the FTC was certainly more activist than it had been prior to the overhaul initiated by Nixon, consumer advocates criticized the FTC for failing to become "a real zealot in terms of consumer affairs" (U.S. Senate 1974, p. 150). That changed in 1977, when President Jimmy 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" (Harris and Milkis 1996, p. 155). Under Pertschuk, the FTC put greater emphasis on hiring committed consumer advocates for the staff. For example, Harris and Milkis (1996, p. 178) quote an FTC attorney as saying, "[W]ho is better, a 4.0 graduated 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 pro-consumer 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 it had been under his predecessors (Harris and Milkis 1996, p. 177). The most controversial rulemaking proceeding in the period entailed a proposal to 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 of agency bias and its interaction with regulatory review. 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 (Marcus 1980, p. 31). 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 (Marcus 1980, pp. 34–37).[24] 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 (Whitaker 1976, p. 55), Nixon adopted the Ash Council's proposal and created the EPA by executive order on July 9, 1970 (Reorganization Plan No. 3, 35 Fed. Reg. 15,623).

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 an 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, "[T]hat is . . . precisely what I would intend to do" (U.S. Senate 1970b, p. 18). Ruckelshaus quickly established his environmentalist bona fides with an aggressive campaign of enforcement of existing environmental laws against polluters (Marcus 1980, pp. 88–90), which led the *New York Times* to refer to him as the "house liberal" in the Nixon administration (*New York Times* 1973, p. 22).

Ruckelshaus'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 better air quality 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 policy making under the act, in May 1971 the director of the Office of Management and Budget (OMB), George Shultz, sent Ruckelshaus a letter informing him that EPA regulations had to be cleared through the OMB and other agencies before being issued (Marcus 1980, p. 125). The following month, the

24. 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" (U.S. Senate 1970a [statement of Senator Edmund Muskie, July 28], p. 40).

OMB asserted that authority by preventing the 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.[25]

Soon after Nixon's reelection in 1972, Nixon decided to retain Ruckelshaus as administrator, but Ruckelshaus stipulated as a condition of staying that Nixon revise the system of OMB review to clarify that the EPA had final authority over regulatory decisions (Quarles 1976, p. 118). While Nixon agreed, the OMB review process continued, and the OMB successfully pressured the EPA to modify rules in response to White House concerns.[26]

Subsequent presidents continued the practice of OMB review of executive agencies' proposed rules. President Ronald Reagan further formalized the process by executive order and expressly authorized the OMB to block regulations by directing agencies to "refrain from publishing" proposed rules until the OMB's review was concluded (Exec. Order 12,291 sec. 3[f], 46 Fed. Reg. 13,193 [February 17, 1981]). Responsibility for coordinating the regulatory review process was lodged in OIRA within the OMB. When President Bill 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 their policy preferences? In the case of Ruckelshaus, 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. Nixon wanted an energetic EPA and achieved it by appointing an agency head who put relatively heavy

25. This sequence of events was recounted at a Senate hearing called over the controversy, at which Ruckelshaus insisted that he had made the final call on the modifications (U.S. Senate 1972).

26. For example, Quarles (1976, pp. 117–42) recounts an episode in which Office of Management and Budget officials successfully pressured the Environmental Protection Agency to delay the target date of new regulations restricting the lead content of gasoline in 1973.

weight on the benefits of environmental regulation. The OMB review process helped keep in check the resulting bias in EPA's specific rule proposals. The fact that centralized regulatory review emerged in tandem with the intentional creation by the president of an activist regulatory agency, led by an appointee who was more pro-regulatory than the president, nicely illustrates the complementarity between politicization and centralization that our theory, in contrast to existing accounts, implies.

5. CONCLUSION

Our goal in this paper has been to provide an account of why bureaucratic principals appoint agents with policy preferences different from their own and subject those agents' policy decisions to review by more aligned bureaucrats. While our primary goal is descriptive, we conclude by 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.[27]

A second justification for OIRA review focuses on the value of presidential control rather than on specifically a need either to check overzealous agencies or to correct for agency capture. Kagan (2000) argues that OIRA review facilitates presidential control over the administrative state. In her view, presidential control of the bureaucracy enhances the democratic legitimacy of bureaucratic decision making 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,

27. But 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.

presidential control results in more rational, cost-effective, and consistent bureaucratic decision making.

But critics of regulatory review argue that the institution in practice has been systematically biased against regulation (Morrison 1985; Percival 1991). For example, they observe that changes made to rules during OIRA review tend to reduce the stringency level of regulations (Bressman and Vandenbergh 2006) and that OIRA reviews only agency decisions to act, not agency decisions not to act (Bagley and Revesz 2010). 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 (Bagley and Revesz 2010).

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 (Livermore and Revesz 2012, p. 13). In addition, presidents and their senior staff only rarely directly intervene in agency rules under OIRA review.[28]

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 rule making can be grouped into two categories: direct presidential policy decisions on specific rules and presidential decisions on the structure of the administrative state and its personnel. In our model of regulatory review, the principal makes no direct policy decisions. Rather, she achieves her policy objectives by choosing which bureaucrats to appoint and by structuring their decision-

28. Livermore 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."

making process. As the model makes clear, in terms of achievement of the principal's policy objectives, a system with regulatory review strictly dominates full delegation. Hence, the frequency with which the president personally intervenes in rule making 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 does strictly worse than one who can commit to delegating the policy decision to a reviewer.

Furthermore, our analysis shows why 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 a very strong deregulatory motive, he 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 leads more often to a reduction rather than to an increase in the stringency level 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 OIRA 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 over 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, Alex, and Charles M. Cameron. 2012. Regulatory Auditing at the Office of Information and Regulatory Affairs. Unpublished manuscript. Princeton University, Department of Politics, Princeton, N.J.

Aghion, Philippe, and Jean Tirole. 1997. Formal and Real Authority in Organizations. *Journal of Political Economy* 105:1–29.

American Bar Association. 1969. *Report of the ABA Commission to Study the Federal Trade Commission*. Chicago: American Bar Association.

Bagley, Nicholas, and Richard Revesz. 2010. Centralized Oversight of the Regulatory State. *Columbia Law Review* 106:1260–1329.

Bendor, Jonathan, Amihai Glazer, and Thomas Hammond. 2001. Theories of Delegation. *Annual Review of Political Science* 4:235–69.

Bendor, Jonathan, and Adam Meirowitz. 2004. Spatial Models of Delegation. *American Political Science Review* 98:293–310.

Bressman, Lisa Schultz, and Michael P. Vandenbergh. 2006. Inside the Administrative State: A Critical Look at the Practice of Presidential Control. *Michigan Law Review* 105:47–99.

Bueno de Mesquita, Ethan, and Matthew C. Stephenson. 2007. Regulatory Quality under Imperfect Oversight. *American Political Science Review* 101:605–20.

Calvert, Randall L., Matthew D. McCubbins, and Barry R. Weingast. 1989. A Theory of Political Control and Agency Discretion. *American Journal of Political Science* 33:588–611.

Che, Yeon-Koo, and Navin Kartik. 2009. Opinions as Incentives. *Journal of Political Economy* 117:815–60.

Clarkson, Kenneth W., and Timothy J. Muris. 1981. *The Federal Trade Commission since 1970: Economic Regulation and Bureaucratic Behavior*. Cambridge: Cambridge University Press.

Cox, Edward Finch, Robert C. Fellmeth, and John E. Schulz. 1969. *The Nader Report on the Federal Trade Commission*. New York: R. W. Baron.

DeMuth, Christopher C., and Douglas H. Ginsburg. 1986. White House Review of Agency Rulemaking. *Harvard Law Review* 99:1075–88.

Dessein, Wouter. 2002. Authority and Communication in Organizations. *Review of Economic Studies* 69:811–38.

Dewatripont, Mathias, and Jean Tirole. 1999. Advocates. *Journal of Political Economy* 107:1–39.

Downs, Anthony. 1967. *Inside Bureaucracy*. Boston: Little, Brown.

Gailmard, Sean, and John W. Patty. 2007. Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise. *American Journal of Political Science* 51:873–89.

Gilligan, Thomas W., and Keith Krehbiel. 1987. Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures. *Journal of Law, Economics, and Organization* 3:287–335.

Harris, Richard A., and Sidney M. Milkis. 1996. *The Politics of Regulatory Change: A Tale of Two Agencies*. 2d ed. New York: Oxford University Press.

Henderson, Rebecca. 1994. The Evolution of Integrative Capability: Innovation in Cardiovascular Drug Discovery. *Industrial and Corporate Change* 3:607–30.

Holmstrom, Bengt, and Paul Milgrom. 1991. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics, and Organization* 7:24–52.

Kagan, Elena. 2000. Presidential Administration. *Harvard Law Review* 114: 2245–2385.

Lewis, David E. 2008. *The Politics of Presidential Appointments*. Princeton, N.J.: Princeton University Press.

Livermore, Michael, and Richard Revesz. 2012. OIRA, Regulatory Review, and Capture. Unpublished manuscript. New York University School of Law, New York.

Marcus, Alfred Allen. 1980. *Promise and Performance: Choosing and Implementing an Environmental Policy.* Contributions in Political Science 39. Westport, Conn.: Greenwood Press.

McCarty, Nolan. 2004. The Appointments Dilemma. *American Journal of Political Science* 48:413–28.

McCubbins, Matthew D., Roger G. Noll, and Barry R. Weingast. 1987. Administrative Procedures as Instruments of Political Control. *Journal of Law, Economics, and Organization* 3:243–77.

Moe, Terry M. 1985. The Politicized Presidency. Pp. 235–72 in *The New Direction in American Politics,* edited by John E. Chubb and Paul E. Peterson. Washington, D.C.: Brookings Institution.

Moe, Terry M., and Scott A. Wilson. 1994. Presidents and the Politics of Structure. *Law and Contemporary Problems* 57(2):1–44.

Morrison, Alan B. 1985. OMB Interference with Agency Rulemaking: The Wrong Way to Write a Regulation. *Harvard Law Review* 99:1059–74.

New York Times. 1973. Government's Pollution Fighter. April 12.

Niskanen, William A. 1975. Bureaucrats and Politicians. *Journal of Law and Economics* 180:617–43.

Nixon, Richard M. 1969. Special Message to the Congress on Consumer Protection. Pp. 883–91 in *Richard Nixon: 1969: Containing the Public Messages, Speeches, and Statements of the President.* Public Papers of the Presidents of the United States. Washington, D.C.: Office of the Federal Register.

Percival, Robert V. 1991. Checks without Balance: Executive Office Oversight of the Environmental Protection Agency. *Law and Contemporary Problems* 54:127–204.

Posner, Eric A. 2001. Controlling Agencies with Cost-Benefit Analysis: A Positive Political Theory Perspective. *University of Chicago Law Review* 68:1137–99.

Prendergast, Canice. 2007. The Motivation and Bias of Bureaucrats. *American Economic Review* 97:180–96.

Quarles, John. 1976. *Cleaning up America: An Insider's View of the Environmental Protection Agency*. Boston: Houghton Mifflin.

Stephenson, Matthew C. 2011. Information Acquisition and Institutional Design. *Harvard Law Review* 124:1422–1829.

Sunstein, Cass. 2011. Letter to Environmental Protection Agency administrator Lisa Jackson. September 2. http://www.whitehouse.gov/sites/default/files/ozone_national_ambient_air_quality_standards_letter.pdf.

U.S. Senate. 1970a. Committee on Government Operations. *Reorganization Plan*

Nos. 3 and 4 of 1970: Hearings before the Subcommittee on Executive Reorganization and Government Research. 91st Cong., 2d Sess. July 28, 29; September 1.

———. 1970b. Committee on Public Works. *Nomination of William D. Ruckelshaus: Hearings before the Committee on Public Works*. 91st Cong., 2d Sess. December 1, 2.

———. 1972. Committee on Public Works. *Implementation of the Clean Air Act Amendments of 1970—Part I (Title I): Hearings before the Subcommittee on Air and Water Pollution*. 92d Cong., 2d Sess. February 16, 17, 18, 23.

———. 1973. Committee on Commerce. *Nomination of Lewis A. Engman, to Be a Commissioner: Federal Trade Commission Hearings*. 93d Cong., 1st Sess.

———. 1974. Committee on Commerce. *Federal Trade Commission Oversight: Hearings before the Committee on Commerce*. 93d Cong., 2d Sess. March 1, 7, 14; May 9.

———. 1977. Committee on Government Operations. *Study on Federal Regulation: Volume 1: The Regulatory Appointments Process*. 95th Cong., 1st Sess. S. Doc. No. 95-25.

Warren, Patrick L. 2012. Allies and Adversaries: The Roles of Appointees in Administrative Policymaking under Separation of Powers. *Journal of Law, Economics, and Organization* 28:407–46.

Whitaker, John C. 1976. *Striking a Balance: Environment and Natural Resources Policy in the Nixon-Ford Years*. Washington, D.C.: American Enterprise Institute for Public Policy Research.