DOI: 10.1017/S0003055407070359

Regulatory Quality Under Imperfect Oversight

ETHAN BUENO DE MESQUITA *Washington University*
MATTHEW C. STEPHENSON *Harvard Law School*

We analyze the positive and normative implications of regulatory oversight when the policymaking agency can improve the quality of regulation through effort, but only some kinds of effort are observable by the overseer, and the overseer's only power is the ability to veto new regulation. Such oversight can increase the quality of agency regulation, but it also introduces inefficiencies—the agency underinvests in unobservable effort and overinvests in observable effort. Agencies have no incentive to conceal their activities from the overseer; the reforms that are likely to reduce inefficiency are therefore those that improve overseer expertise or lower the costs of agency disclosure, not those that compel disclosure. The normative implications depend on the relative severity of bureaucratic drift and slack problems. When slack is paramount, an overseer that is more anti-regulation than the agency or society improves social welfare, as long as it does not deter the agency from regulating entirely. When drift is paramount, oversight improves social welfare only when it deters regulation. In this case, regulatory oversight is weakly dominated by one of two alternatives: eliminating oversight or banning regulation.

There is a deep and persistent ambivalence about bureaucratic government in American politics. On the one hand, the perceived need for government regulation of private social and economic activity, coupled with the demand for technocratic expertise in designing such regulation, fueled a dramatic expansion in the size and power of the federal bureaucracy in the twentieth century (Spence and Cross 2000). On the other hand, some reject the premise that widespread government regulation is desirable (Epstein 1995; Hayek 1944), whereas others attack the constitutional and democratic legitimacy of vesting regulatory authority with bureaucrats rather than with Congress (Lawson 1994; Lowi 1979). Even if one puts these objections aside and accepts in principle the desirability of bureaucratic governance, a number of more practical concerns remain. Two of the most serious of these are "bureaucratic drift" and "bureaucratic slack."

Bureaucratic drift, as we use the phrase, is an umbrella term for a diverse set of phenomena that lead an administrative agency to pursue policies whose consequences diverge from social and/or legislative goals (Horn and Shepsle 1989; Shepsle 1992). Drift may result, for example, from agency capture by regulated entities or interest groups (Niskansen 1971; Stigler 1971), the *ex ante* selection practices and *ex post* career ambitions of bureaucrats (Eckert 1981; Heclo 1988), or cognitive or institutional biases that afflict agency decision-making (Rachlinski and Farina 2002; Seidenfeld 2002). Bureaucratic slack, in contrast, is the tendency of agencies to pursue their mandates with insufficient effort, diverting resources from producing higher quality regulations to socially unproductive activities, such as leisure time for bureaucrats (Moe 1990).[1]

The drift and slack problems mean that, even if one accepts that bureaucratic government may be desirable in principle, it will often produce suboptimal results in practice. One potential solution to these problems is to subject agencies to oversight by another entity (DeMuth and Ginsburg 1986; Sunstein 1984). Although there are many forms of bureaucratic oversight in the American political system, the formal literature has focused primarily on oversight by legislatures and those parts of the executive branch with the ability to manipulate budgets, restructure bureaucratic organizations, and alter the scope of delegated authority.[2] However, there are other important forms of bureaucratic oversight that rely on different, and more blunt, instruments of control.

For example, under the the Administrative Procedure Act (APA), courts are supposed to take a "hard look" at agency regulations and reject rules that are "arbitrary and capricious" (Breyer 1986; McGarity 1992).[3] Courts, however, generally lack the power to propose

Ethan Bueno de Mesquita is Assistant Professor, Department of Political Science, CB 1063, Washington University, One Brookings Dr., St. Louis, MO 63130 (ebuenode@artsci.wustl.edu).

Matthew C. Stephenson is Assistant Professor, Harvard Law School, Griswold 509, Cambridge, MA 02138 (mstephen@law.harvard.edu).

We have received valuable comments from Scott Ashworth, Bruce Bueno de Mesquita, Amanda Friedenberg, Dimitri Landa, Gary Miller, Hong Min Park, Cass Sunstein, Adrian Vermeule, and Alan Wiseman.

[1] "Bureaucratic drift" is sometimes used to include the phenomenon we have called bureaucratic slack (Hopenhayn and Lohmann 1996; McCubbins, Noll, and Weingast 1987). Moreover, what we refer to as slack is alternatively referred to as shirking (Moe 1990). The terminology is not important as long as the conceptual distinction is clear.

[2] For instance, scholars have studied the initial decision to delegate authority to a bureaucratic agent (Aranson, Gellhorn, and Robinson 1983; Epstein and O'Halloran 1999) and the assignment of delegated power among different potential delegates (Bendor and Meirowitz 2004; Boehmke, Gailmard, and Patty 2006; Stephenson 2006b). Other work has explored how the legislature might structure agency decision-making processes (Bawn 1995; McCubbins, Noll, and Weingast 1989) or offer various incentive schemes (Gailmard 2006; Ting 2001; Weingast and Moran 1983) in order to align the agency's preferences more closely with those of the legislature, to induce the agency to reveal private information (Banks and Weingast 1992; de Figueiredo, Spiller, and Urbiztondo 1999; Stephenson 2006a), or to give the agency incentives to invest in expertise (Gailmard and Patty N.d.; Stephenson 2007).

[3] Under the Supreme Court's decision in *Vermont Yankee v. Natural Resources Defense Council*, 435 U.S. 519 (1978), federal courts may not require agencies to comply with procedural requirements other than those mandated by statute or by the Constitution. However,

new rules,[4] or to influence an agency's budget or structure. Another important form of bureaucratic oversight is performed by the Office of Information and Regulatory Affairs (OIRA), a division of the Office of Management and Budget (OMB; Wiseman 2007). Under executive orders promulgated by Presidents Reagan and Clinton, agencies must notify OIRA of any major proposed regulation, and OIRA is empowered to return proposals to the agency "for further consideration." Though OIRA cannot formally veto a regulatory proposal, in practice the review process gives OIRA the ability to delay indefinitely regulations it finds unsatisfactory (Cooper and West 1988; Morrison 1986). OIRA, like the courts, cannot enact regulations on its own, nor can it alter agency budgets or structure.[5]

Although courts, OIRA, and similar oversight institutions cannot control agency budgets, make delegation decisions, or alter institutional rules, they have the power effectively to veto regulations, and this authority can have a significant impact on bureaucratic politics (Gordon and Hafer 2005). These and other forms of oversight are meant to improve the quality of regulatory output by forcing the policymaking agency to produce only those regulations that are of sufficiently high quality (relative to the status quo) that the overseer would approve them. A key strategic problem faced by this sort of overseer, however, is its inability to evaluate directly the quality of a proposed regulation. Instead, the overseer must attempt to infer regulatory quality from other indicia, such as the effort devoted by the agency to various quality-improving activities. But not all of the relevant quality-improving activities are observable by the overseer.

For example, courts may be relatively good at observing whether agencies have satisfied procedural requirements, compiled a detailed record, and provided reasonable-sounding responses to comments and objections lodged by interested parties. All of these agency activities may be positively correlated with the expected quality of the final regulation that the agency adopts (Seidenfeld 1997; Sunstein 1984). Courts, however, may not be very good at assessing other correlates of regulatory quality, such as the soundness of an agency's scientific or economic analysis (McGarity 1992; Shapiro 1988). Moreover, much of the actual analytical work that goes into developing a regulation may occur in informal contexts that the courts cannot observe at all (Cross 2000; Shapiro 1988). Similarly, the OIRA staff who evaluate agency regulations are likely to be effective in assessing the quality and sophistication of an agency's cost-benefit analysis (CBA), at least with respect to those costs and benefits that are easily susceptible to quantification on a common scale (DeMuth and Ginsburg 1986). But OIRA might be quite bad at evaluating other correlates of regulatory quality, including the agency's assessment of costs and benefits that do not fit as well within a traditional CBA framework, such as certain types of environmental harm, distributive justice considerations, and social attitudes toward risk and regulation (Ackerman and Heinzerling 2002; Wagner 2003). As a general matter, the observability of agency effort is likely to vary by policy area and by overseer type. We do not attempt to specify the conditions under which particular types of agency effort are observable. Our objective, instead, is to consider the implications of regulatory oversight when at least some relevant forms of agency effort cannot be observed by the overseer.[6]

We analyze the positive and normative implications of regulatory oversight when the policymaking agency can improve the quality of the regulations it produces by investing in costly activities (e.g., research, analysis, consultation, hearings, old-fashioned hard work); the overseer can observe some but not all of these activities; and the overseer is empowered only to accept or reject the proposed regulation. The existence of observable and unobservable effort creates a multitask moral hazard problem between the agency and the overseer (Holmström and Milgrom 1991), leading to inefficient effort allocations. The model has implications for the quality and frequency of regulation, agency incentives to disclose information to overseers, and the impact of oversight on social welfare. When oversight affects agency behavior, it increases the quality of regulation but also introduces inefficient distortions in agency effort allocation. As a result, sometimes Pareto improving regulation will not occur. Therefore, agencies have no incentive to keep their

the Court's interpretation of the APA's "arbitrary and capricious" standard in cases such as *State Farm v. Motor Vehicles Manufacturers' Association*, 463 U.S. 29 (1983), requires that agencies demonstrate "reasoned" decision making, which in practice often induces agencies to adopt additional procedures (Pierce 1995).

[4] This statement must be qualified somewhat, given that private parties can sue to compel certain regulatory actions. Successful action-forcing suits are rare, and even when the court orders an agency to act, the agency almost always retains a great deal of discretion.

[5] Of course, OMB is not an independent or purely technocratic institution. Rather, OMB seeks to advance the President's regulatory agenda, and the President has considerable weapons at his disposal to influence agency decisions. Indeed, this is one of the reasons that OIRA may wield a de facto veto over agency decisions, even though as a formal legal matter agencies do not have to accept OIRA recommendations. The fact that the President may have other means by which to control agencies does not, however, alter the fact that this particular form of Presidential oversight—OIRA review—involves the power to block regulation but not to initiate regulation or to manipulate agency budgets or structure.

[6] Of course, overseers try to acquire information about those aspects of regulatory proposals that they cannot directly observe or evaluate. Thus, President Clinton's executive order specifically instructed OIRA to consider "qualitative measures of costs and benefits that are difficult to quantify, but nevertheless essential to consider," and OIRA has recently upgraded its ability to evaluate scientific issues by hiring more staffers with scientific training. If a regulation implicates complex technical questions, overseers might also solicit input from outside experts. And, because different types of effort may be observable to different overseers, overall oversight might be improved with multiple overseers. In practice, however, these strategies are unlikely to solve the observability problem completely. Hiring experts or soliciting their opinions is limited by institutional and cost considerations, and may create an agency problem within the oversight entity. Adding layers of oversight is also costly and will typically leave some gaps in overseer expertise and capacity. Thus, even after overseers have adopted all cost-justified means to learn more information about the agency's decision, uncertainty is likely to persist.

activities unobservable, and the reforms that are likely to improve regulatory quality are those that improve overseer expertise or lower the costs of agency information disclosure, not those that compel disclosure. Finally, the normative evaluation of different oversight arrangements depends on the relative size of the bureaucratic drift and slack problems. When slack is a significant problem but drift is not, oversight by an entity that is more anti-regulation than society itself improves social welfare, as long as the overseer's demands do not deter regulation altogether. When drift is a significant problem but slack is not, oversight is weakly dominated by banning either oversight or regulation.

THE MODEL

Actions

Consider a game with two players: an overseer (referred to as C for "court") and an agency (A). The order of play is as follows. First the agency chooses whether or not to initiate new regulation. We will denote this choice by $\rho \in \{r, nr\}$, where $\rho = r$ represents the decision to regulate. If the agency decides to regulate, it can improve the quality of the regulation it ultimately promulgates by exerting costly effort. The agency can allocate its effort to two different sets of activities: those that are observable by the overseer ($a_O \in \mathbb{R}^+$) and those that are unobservable by the overseer ($a_U \in \mathbb{R}^+$). A strategy for the agency is a triple, (ρ, a_O, a_U).

If the agency does not regulate, then the overseer does not have a decision to make. If the agency does regulate, then the overseer, after seeing the agency's choice of observable effort, decides whether to uphold the regulation. We denote this decision by $\sigma \in \{0, 1\}$, where $\sigma = 1$ represents the decision to uphold the regulation. If the regulation is upheld, it is implemented. If it is struck down then the status quo remains in effect. We treat a decision by the overseer to remand the decision to the agency as equivalent to a decision to strike down the regulation. The difference between a remand and an outright reversal, though significant in some contexts, does not matter in our model, as we discuss in more detail later. A strategy for the overseer is a mapping $\sigma : \{r, nr\} \times \mathbb{R}^+ \rightarrow \{0, 1\}$, which says whether the overseer will uphold a regulation for any given level of observable effort by the agency, should new regulation be proposed.[7]

Payoffs

We assume that the agency and the overseer have different preferences with respect to regulatory policy. We model this by assuming that each player, $i \in \{C, A\}$, has a payoff from the existing status quo given by q_i. The distance $|q_C - q_A|$ can be interpreted as the extent to which the preferences of the agency and the overseer differ with respect to the value of new regulation. If, for example, the regulation under consideration targets air pollution discharged by electric utilities, q_i may incorporate the relative values that player i places on environmental protection and low electricity prices; the more player i cares about the former goal relative to the latter, the lower q_i will be. So, for example, if the agency places more weight on environmental protection relative to low electricity prices than does the overseer, q_A will be lower than q_C.

Although we assume that the agency and the overseer have divergent preferences, we also assume that their interests are positively correlated, in that agency effort increases the quality of the regulation from the perspective of both the agency and the overseer. For example, suppose that the agency and the overseer place different weights on the values of clean air and low electricity costs, but that the agency can, through costly effort, design its regulation so that it achieves greater pollution reductions at lower cost. Both the agency and the overseer would view this as an improvement in regulatory quality, even though they will continue to disagree about how high-quality the regulation must be to make it preferable to the status quo.

Our approach is most applicable to circumstances in which the agency and the overseer share the same basic values (e.g., health, safety, efficiency), even though they may disagree sharply over how to trade off these values when they conflict (Stephenson 2007). Of course, there may also be forms of agency effort that increase regulatory quality from the agency's perspective but decrease regulatory quality from the overseer's perspective. We do not model those forms of effort directly, although they are captured to some extent by the difference between q_A and q_C. (If, for example, the overseer either observes or can infer that the agency invested resources in activities that reduce the value of regulation to the overseer, the overseer would presumably place a higher relative value on retaining the status quo, q_C.) Our analysis may be less applicable in cases where an agency is focused exclusively or primarily on goals that the overseer finds objectionable. Also, although our model allows the preferences of the agency and the overseer to diverge—because each may attach a different payoff to retaining the status

[7] Although our model builds on the multitask moral hazard idea first suggested by Holmström and Milgrom (1991), it differs in several ways. In the canonical model, the agent likes wages and dislikes effort, but shares no interests with the principal. In our model, both the agency and the overseer prefer higher quality regulation, all else equal. This difference in preference structure gives rise to both the information disclosure result in Proposition 4, which need not hold in the standard model, and the social welfare analysis, which has no analogy in the canonical model. Second, in our model the overseer has only a blunt tool—its ability to veto regulation—which provides less flexibility than the schedule of wages an employer can offer a worker in the standard multitask model. Third, the order of play in our model conforms to the order of moves in real world oversight—the agency proposes regulation, then the overseer evaluates it. In Holmström and Milgrom, the order of play is the reverse, as is standard in models of contracting. This difference is important because it gives rise to a commitment problem for our overseer that does not exist in the standard contracting model. In particular, the overseer would like to be able to credibly threaten to reject regulation with higher levels of observable spending than it does in equilibrium.

quo—the model does not directly cover a case in which the overseer considers certain types of effort more or less productive of regulatory quality than does the agency.[8]

Formally, we assume that, given effort allocation (a_O, a_U), the quality of the regulation is $f(a_O, a_U)$, where $f(\cdot, \cdot)$ is common knowledge, strictly increasing, and concave in both of its arguments. That is, the quality of regulation is increasing at a decreasing rate in both observable and unobservable effort.

We further assume that $f_{OU} = 0$ (where f_{OU} is the cross-partial with respect to a_O and a_U), meaning that the marginal product of one form of effort is not affected by the other form of effort. This assumption greatly simplifies the analysis and allows us to focus on substantive intuitions by eliminating some indirect effects. The assumption would be reasonable in cases where the observable and unobservable forms of agency effort addressed different aspects of regulatory quality. In most cases, however, observable and unobservable forms of effort are likely to be partial substitutes or complements ($f_{OU} \neq 0$). Observable and unobservable effort might be substitutes ($f_{OU} < 0$) if they both improve the quality of the regulation along a similar dimension. For example, an agency might be able to improve the cost-effectiveness of regulation by conducting (observable) cost-benefit studies and by engaging in (unobservable) consultations with industry and other interested parties.[9] The marginal benefit of consultation might be higher if the agency has not invested much in the studies, and vice versa, because both forms of effort provide similar information on compliance costs. One can also imagine situations in which observable and unobservable efforts are complements ($f_{OU} > 0$). Suppose, for instance, that the agency can engage in observable effort to improve the efficiency of a regulation, and can engage in unobservable effort to improve the enforceability of the regulation. The marginal product of improving enforceability is higher when the regulation is more efficient, and the marginal product of improving efficiency is higher when the regulation is widely enforced. Because many real-world situations may violate our simplifying assumption that $f_{OU} = 0$, we include a discussion of the implications of relaxing this assumption immediately following our presentation of the basic model. We show that our results are robust to substitution effects and complementarities as long as these effects are not too strong.

Finally, we assume that $\lim_{a_i \to 0} f_i(\cdot, \cdot) = \infty$ for $i \in \{O, U\}$ and all a_i. This assumption ensures that, if the agency prefers to exert any effort, it will choose positive levels of both types of effort.

Although the agency prefers higher quality regulation, it also bears costs for effort. These costs might be thought of as forgone leisure time or the diversion of budgetary surplus away from bureaucratic perks. This interpretation is consistent with the view that the agency suffers from what we have termed bureaucratic slack. Alternatively, agencies might perceive effort as costly because the resources devoted to improving the quality of a given regulation must be diverted from the pursuit of other valuable policy goals. Under that interpretation, effort is costly to the agency not because of forgone slack, but simply because of the opportunity costs associated with improving the quality of any particular regulation.

The agency's payoff from proposing new regulation and allocating efforts (a_O, a_U) is:

$$V_A(r, a_O, a_U, \sigma) = \sigma f(a_O, a_U) + (1 - \sigma) q_A - c(a_O + a_U),$$

where $c(\cdot)$ is the cost of effort. This expected payoff function says that if regulation is upheld ($\sigma = 1$), then the payoff to the agency is $f(a_O, a_U) - c(a_O + a_U)$, and if the regulation is struck down ($\sigma = 0$), then the payoff to the agency is $q_A - c(a_O + a_U)$. We assume that $c(\cdot)$ satisfies $\lim_{x \to \infty} c'(x) = \infty$, that it is strictly increasing ($c' > 0$), convex ($c'' > 0$), and that $c''' \geq 0$. All of these assumption are satisfied, for example, by quadratic costs ($c(a_0 + a_U) = (a_O + a_U)^2$).

The agency's payoff if it does not propose new regulation is

$$V_A(nr, a_O, a_U, \sigma) = q_A.$$

The overseer also prefers higher quality regulation. In contrast to the agency, the overseer does not bear the costs associated with improving regulatory quality. Importantly, although the overseer observes the level of observable effort (a_O) and knows the functional form of $f(\cdot, \cdot)$, it observes neither the level of unobservable effort (a_U) nor the quality of the regulation $(f(a_O, a_U))$ at the time it makes its decision. Hence, the overseer's expected payoff, if regulation is proposed, is

$$V_C(r, a_O, a_U, \sigma) = \sigma \mathbb{E}[f(a_O, a_U)] + (1 - \sigma) q_C.$$

As noted earlier, we do not directly model the possibility that the overseer might remand the decision to the agency with more specific instructions. Given our assumptions about preferences and information, such an alternative is superfluous. As we show below, in equilibrium the agency never proposes a regulation that the overseer would not uphold. And, because the overseer cannot observe regulatory quality or unobservable effort, the overseer could not credibly insist, in a remand order, that the agency improve

[8] Allowing for the agency and overseer to have different, but positively correlated, views of the effectiveness of the two types of effort would not qualitatively change equilibrium behavior. The overseer would demand a level of effort that made it believe the quality of regulation was sufficiently high, and the agency would only be willing to provide such effort if the resulting quality of regulation, from the agency's point of view, was high enough relative to the status to offset the costs. However, introducing this additional complexity would make the evaluation of social welfare more complicated, as we would need to specify society's view of the effectiveness of the two types of effort.

[9] The choice of which type of effort is observable in this example is arbitrary.

regulatory quality through increased unobservable effort.

Finally, if no regulation is proposed, the overseer's payoff is simply[10]

$$V_C(nr, a_O, a_U, \sigma) = q_C.$$

EQUILIBRIUM

Our equilibrium concept is perfect Bayesian equilibrium (PBE), which requires the following. At each information set, the overseer observes the level of observable effort, forms beliefs about unobservable effort, and upholds regulation if and only if the regulation yields higher expected payoffs than the status quo given these beliefs. Moreover, these beliefs must be correct in equilibrium. The agency chooses whether to regulate and its effort allocation to maximize its expected utility, given the action these choices will induce from the overseer.

We further restrict attention to those PBE in which the overseer has what we argue are reasonable beliefs about unobservable effort. In particular, off the equilibrium path, the overseer believes that if the agency chose to regulate ($\rho = 1$) and chose a positive level of observable effort, then the agency must have thought that its regulation would be upheld when choosing its unobservable effort. Such a belief is reasonable because it is irrational for the agency to exert any effort if it believes regulation will not be upheld.[11] We refer to the unique PBE satisfying this restriction simply as an *equilibrium*.

Unobservable Effort

The first step in solving the model is to determine what level of unobservable effort the agency will choose, given a level of observable effort. The following result makes it easy to characterize the agency's problem.

Lemma 1. *The agency will propose and invest in regulation only if it will be upheld.*[12]

Given Lemma 1 and the fact that unobservable effort cannot affect the overseer's decision, we examine how much unobservable effort the agency will expend, conditional on regulation passing. Of course, if the overseer would not uphold the regulation given the level of observable effort, the agency will not invest in any unobservable effort. However, if the agency anticipates that regulation will be upheld, then it chooses a level of unobservable effort to solve

$$\max_{a_U} f(a_O, a_U) - c(a_O + a_U).$$

Given the fact that the objective function is concave, and the assumptions on the limits of f_U and c', the first-order condition defines the unique optimal level of unobservable effort as a function of observable effort, $a_U^*(a_O)$:

$$f_U(a_O, a_U^*) = c'(a_O + a_U^*). \quad \textbf{(1)}$$

This condition says that, for a fixed level of observable effort and assuming regulation will be upheld, the agency chooses a level of unobservable effort such that the marginal benefit, in terms of increased regulatory quality, equals the marginal cost. Two facts can be deduced from this result that will be useful later in the analysis.

Lemma 2. *The optimal level of unobservable effort* (a_U^*) *is characterized by equation* (1) *and is decreasing and concave in observable effort* (a_O), *so long as the level of observable effort induces beliefs for the overseer that ensure that the regulation will be upheld.*

The Overseer's Decision

The overseer will only uphold regulation if the regulation's expected payoff, given the overseer's beliefs, is greater than the overseer's status quo payoff. The overseer uses the level of observable effort to infer the level of unobservable effort and, thus, the overall quality of the regulation.

Let $\tilde{a}_U(a_O) : \mathbb{R}^+ \to \mathbb{R}^+$ denote the overseer's beliefs about the level of unobservable effort, given the level of observable effort. These beliefs are formed by inverting the agency's optimization problem (thus, $\tilde{a}_U(a_O) = a_U^*(a_O)$). Given this, on observing the level of observable effort, the overseer will uphold the regulation if and only if

$$f(a_O, a_U^*(a_O)) \geq q_C. \quad \textbf{(2)}$$

We now have the following result.

Lemma 3. *The overseer upholds regulation if inequality* (2) *holds. As* a_O *increases*, $f(a_O, a_U^*(a_O))$ *is either decreasing everywhere*, *increasing everywhere*, *or has a single interior peak. If there exists any level of*

[10] Although we assume that the overseer's payoff if it strikes down regulation, q_C, is equal to its payoff if no regulation is ever proposed, relaxing the assumption does not alter the results. If the agency does not propose regulation, there is nothing the overseer can do to alter its payoffs. Thus, our results hold even if the overseer's payoff from striking down an agency regulation is arbitrarily higher or lower than the payoff the overseer receives if no regulation is ever proposed.

[11] This restriction rules out equilibria of the following form. The overseer (unreasonably) believes that, for some level of observable effort, the agency will choose a level of unobservable effort less than what the agency would actually choose if that information set were reached. Given this belief, the overseer would strike down the regulation, even though the overseer would uphold the regulation given the agency's utility maximizing level of unobservable effort. Consequently, the agency never chooses that level of observable effort. The overseer's beliefs are consistent, because the relevant information set is never reached. However, the equilibrium is only sustainable because the overseer's unreasonable beliefs deter the agency from taking an action it would have liked to if the overseer had reasonable beliefs.

[12] All proofs are in the Appendix.

FIGURE 1. The overseer upholds regulation if and only if $f(a_O, a_U^*(a_0)) \geq q_C$. Since $f(a_O, a_U^*(a_0))$ is single peaked, the set of acceptable levels of observable spending (Φ) takes one of the three forms illustrated in the figure, as long as the set is not empty

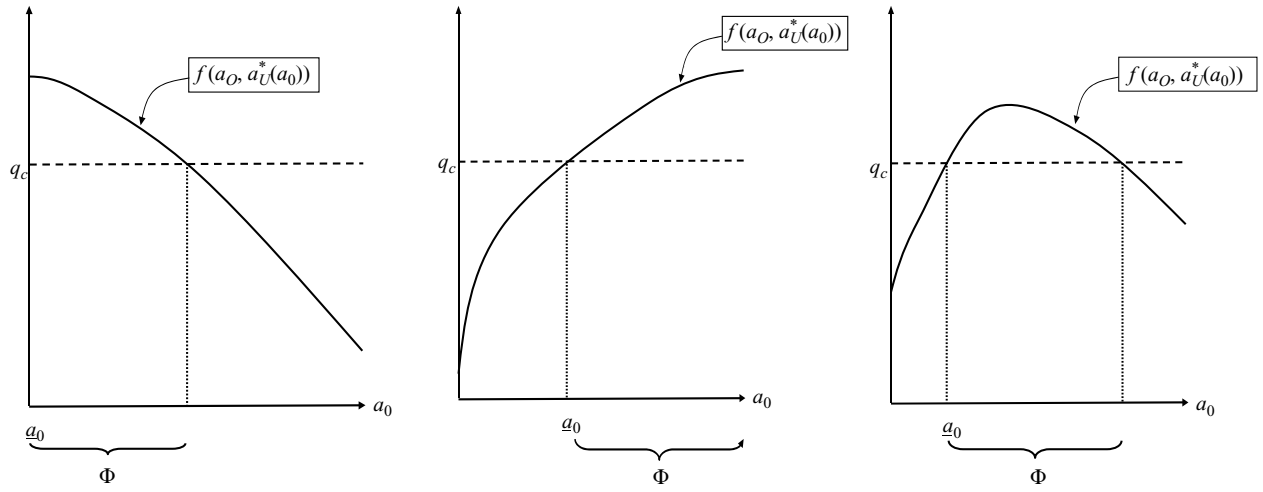

observable spending that will lead the overseer to uphold regulation (*i.e.*, *an* a_O *that satisfies inequality* (2)), *then there is a connected subset*, Φ, *of the real line such that the overseer upholds regulation if and only if the level of observable effort is in that subset. Moreover*, *there is always a lowest level of observable effort that will lead the agency to uphold regulation* (*i.e.*, Φ *is closed on the left*).

Lemma 3 shows how the overseer decides whether to uphold regulation based solely on the level of observable effort.[13] There are two facts to notice from the Lemma.

First, the overseer can only commit to rejecting regulation if the regulation is inferior to the status quo from the overseer's perspective. This limits the overseer's power. If the overseer could commit to any rule for upholding regulation, it could threaten to strike down regulation unless the agency chose the overseer's most preferred level of observable effort (subject to the agency being willing to propose regulation in the first place). This would potentially allow the overseer to extract significantly greater effort from the agency. However, such a threat is not credible. The fact that the overseer has an apparent incentive to commit itself to a more stringent oversight strategy suggests a potentially interesting direction for future research into the institutional mechanisms that may facilitate such commitment, but we do not explore that issue here. Because the overseer cannot tie its hands in our model, it will only strike down those regulations whose quality does not meet the less stringent requirement of dominating the overseer's status quo payoff.[14]

A second fact to notice about Lemma 3 is that because the level of unobservable spending is decreasing in the level of observable spending (Lemma 2), the quality of regulation need not be everywhere increasing in the level of observable effort. Lemma 3 establishes that, if there is any level of observable effort that would lead to regulation being upheld, the range of acceptable observable effort levels takes one of the three forms illustrated in Figure 1. As will become clear, the important fact is that the set of acceptable levels of observable effort (Φ) always includes its lower bound, so that there is a lowest acceptable level of observable effort (labeled $\underline{a}_O$).

Allocation of Observable Effort

Because the agency will only choose to allocate effort if it expects its regulation to be upheld (Lemma 1), the overseer's equilibrium decision rule sets a constraint for the agency. The agency must exert an amount of observable effort in the acceptable range (Φ) if it wants the overseer to approve the regulation.

In the absence of this constraint, the agency would choose observable effort to solve

$$a_O^* = \arg\max_{a_O} f(a_O, a_U^*(a_O)) - c(a_O + a_U^*(a_O)).$$

[13] Although adding the possibility that the overseer could remand a regulation to the agency with specific instructions would be superfluous in our model, Lemma 3 suggests a natural substantive interpretation of remands in our framework. If a_O is outside of Φ, the overseer could remand the decision, informing the agency that the agency's decision can be upheld only if it selects a new $a_O \in \Phi$.

[14] This highlights a difference between our model of agency oversight and multitask models of elections. In an election with identical candidates, voters are indifferent between candidates and are therefore able to commit to a more stringent retrospective voting rule, threatening not to reelect an incumbent who does not choose the voters' most preferred level of effort (e.g., Bueno de Mesquita 2007). This gives voters in the electoral setting more leverage to solve the moral hazard problem than the overseer has in our model of bureaucratic politics.

FIGURE 2. In the first panel, the agency is unaffected by the overseer because its most preferred allocation (a_O^*, a_U^*) is sufficient to induce the overseer to uphold regulation. In the second panel the agency is constrained by the overseer to choose a higher level of unobservable effort ($\underline{a}_O$). In the third panel the agency is dissuaded from regulating because choosing a level of observable effort that would induce the overseer to uphold regulation would given the agency a payoff lower than the status quo

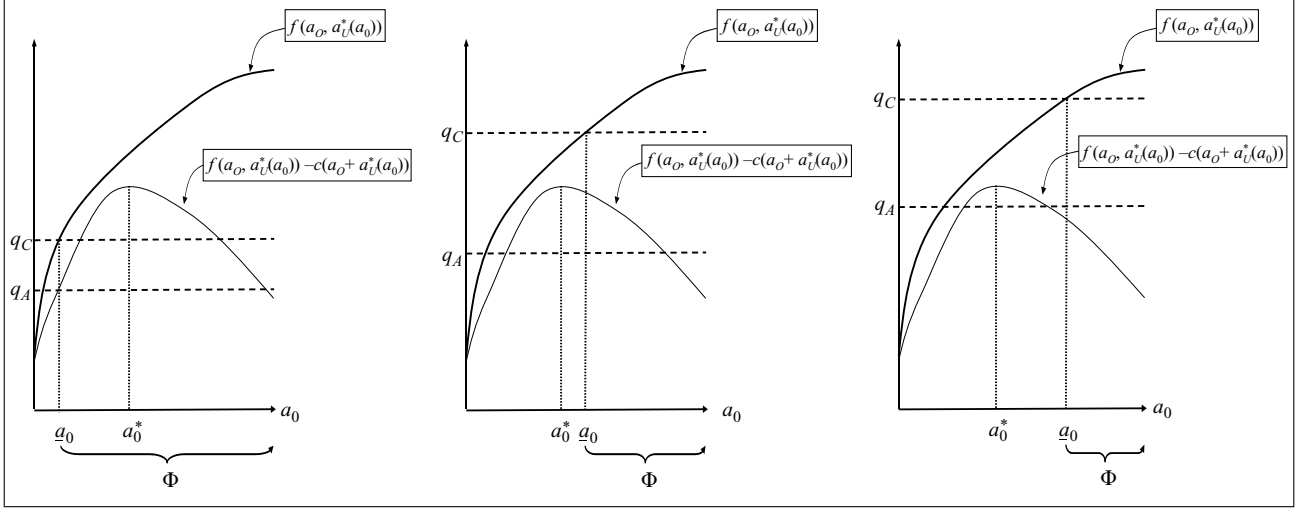

The unique unconstrained optimal level of observable effort is given by

$$f_O(a_O^*, a_U^*(a_O^*)) + f_U(a_O^*, a_U^*(a_O^*))\frac{\partial a_U^*}{\partial a_O} - c'(a_O^* + a_U^*(a_O^*))\left(1 + \frac{\partial a_U^*}{\partial a_O}\right) = 0. \quad \textbf{(3)}$$

This first-order condition reveals several effects of increasing observable effort. First, there are two direct effects. Increasing observable effort increases the quality of regulation ($f_O > 0$), which is a marginal benefit from the agency's perspective, but requires effort ($-c' < 0$) which is a marginal cost from the agency's perspective. Second, there are two indirect effects. When observable effort increases, unobservable effort decreases. This diminishes the quality of regulation ($f_U \frac{\partial a_U^*}{\partial a_O} < 0$), which is a marginal cost from the agency's perspective, but this indirect decrease in unobservable effort also lowers effort costs ($-c' \frac{\partial a_U^*}{\partial a_O} > 0$), which is a marginal benefit from the agency's perspective. The first-order condition equates these marginal costs and benefits. This gives rise to the following:

Lemma 4. *If there is an upper bound on the level of observable effort that will lead to the overseer upholding the regulation* ($\bar{a}_O$)*, then* $a_O^* \leq \bar{a}_O$.

This Lemma implies that the agency never wants to engage in more observable spending than the overseer wants it to. Given this, we can characterize the agency's choice of observable effort.

Lemma 5. *In equilibrium, the agency will choose*

$$a_O = \begin{cases} a_O^* & \text{if } V_A(a_O^*, a_U^*(a_O^*)) \geq q_A \quad \text{and} \quad a_O^* \in \Phi \\ \underline{a}_O & \text{if } V_A(\underline{a}_O, a_U^*(\underline{a}_O) \geq q_A \quad \text{and} \quad a_O^* \notin \Phi \\ 0 & \text{else,} \end{cases}$$

where a_O^* *is implicitly defined by equation* (3).

Lemma 5 shows that there are three possible outcomes associated with bureaucratic oversight in equilibrium. In the first case, the agency wants to propose regulation, and its most preferred level of observable effort is sufficient to induce the overseer to uphold the regulation. Hence, the agency chooses its most preferred level of observable effort. This case is illustrated in the first panel of Figure 2.

In the second case, the agency wants to propose regulation, but its most preferred level of observable effort is not sufficient to convince the overseer to uphold the regulation. The agency will then consider exerting the minimum level of observable effort sufficient to convince the overseer to uphold regulation ($\underline{a}_O$). If the payoff associated with having regulation approved at this level of effort is greater than the payoff to the agency of retaining the status quo, it will propose regulation and choose this level of observable effort. In this scenario, we will refer to the agency as *constrained* by the overseer. This case is illustrated in the second panel of Figure 2.

Definition 1. The agency is *constrained* by the overseer if:

1. $a_O^* \notin \Phi$
2. $f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\underline{a}_O + a_U^*(\underline{a}_O)) \geq q_A$

In the third case, the agency is unwilling to propose regulation at the level of observable effort required for approval by the overseer. In this case the agency will not propose regulation and will exert no effort. If the agency would have proposed regulation in the absence of oversight, but chooses not to due to the overseer's demands, we will say that the agency is *dissuaded* by the overseer. This case is illustrated in the third panel of Figure 2.

Definition 2. The agency is *dissuaded* by the overseer if:

1. $a_O^* \notin \Phi$
2. $f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\underline{a}_O + a_U^*(\underline{a}_O)) < q_A \leq f(a_O^*, a_U^*(a_O^*)) - c(a_O^* + a_U^*(a_O^*))$

Substitutability and Complementarity of Observable and Unobservable Effort

Our analysis employed the simplifying assumption that $f_{OU} = 0$. In particular, three of the results in our equilibrium analysis—Lemmata 2, 3, and 5—made use of this assumption. However, as we noted earlier, the assumption that $f_{OU} = 0$ often will not hold in the real-world applications of greatest interest. An important question, therefore, is how robust the results are to relaxing this assumption.

Lemmata 2 and 3 rely on the assumption that there are not very strong complementarities between observable and unobservable effort. In the case of Lemma 2, strong complementarities (in particular, $f_{OU} > c''$) would create a situation in which unobservable effort is increasing in observable effort. The reason is that an increase in observable effort would increase the marginal benefit (in terms of increased quality) of unobservable effort more than it would increase its marginal cost. If, however, $f_{OU} \leq c''$, then unobservable effort is decreasing in observable effort even if $f_{OU} \neq 0$.

In the case of Lemma 3, if there were very strong complementarities between unobservable and observable effort, then the equilibrium quality of regulation might not be concave in observable effort. This is because as observable effort increased, unobservable effort might also increase fast enough to make quality increase at an increasing rate. This would make it possible for Φ to be made up of disjoint subsets of the real line. This would not undermine our analysis, but it would introduce a technical complication that does not add additional insight.

For Lemma 5, if there were very strong substitution effects between observable and unobservable effort—that is, if f_{OU} were strongly negative—the agency might choose an effort level greater than the lowest acceptable level of effort ($\underline{a}_O$). The reason is that if $f_{OU} < 0$, an increase in the level of observable effort will decrease the efficacy of unobservable effort dramatically. If this effect is very large, then the increase in observable effort demanded by overseer's constraint might make unobservable effort so unproductive that the agency would substitute even further away from unobservable effort and toward observable effort. The result would be that the overseer's demand for a minimum level of observable effort would result in a level of observable effort above that minimum threshold. Again, though, this effect only obtains if the substitution effect is very powerful. For more moderate substitution effects, Lemma 5 will still hold.

As long as these three Lemmata hold, the overall equilibrium analysis remains unchanged. Thus, the equilibrium derived above and the results discussed below are robust to a relaxation of the assumption that $f_{OU} = 0$. In particular, observable and unobservable effort can be complements, as long as the complementarities are not so strong that either Lemma 2 or Lemma 3 no longer holds. Similarly, they can be substitutes, as long as the substitutabilities are not so strong that Lemma 5 does not hold.

RESULTS

Oversight and the Efficiency of Agency Effort Allocation

The existence of unobservable effort can create situations in which the agency's allocation of resources is Pareto inefficient. To see this, consider the case where the agency is constrained by the overseer, so that it chooses an effort allocation given by $(\underline{a}_O, a_U^*(\underline{a}_O))$. Label the total amount of effort exerted $a = \underline{a}_O + a_U^*(\underline{a}_O)$. Holding total effort fixed at a, there is always a new division between observable and unobservable effort that would improve regulatory quality.

The intuition for why this is true is as follows. Because the agency cares about the quality of regulation, it wants to allocate resources efficiently. That is, holding the total level of effort constant, the agency wants the highest quality regulation possible. To achieve this, the agency must allocate its effort such that the marginal product of observable effort ($f_O(a_O, a_U)$) equals the marginal product of unobservable effort ($f_U(a_O, a_U)$).

To satisfy the overseer, the agency is forced to choose a level of observable effort higher than its most preferred level. But the overseer cannot control the agency's unobservable effort. As demonstrated in Lemma 2, the agency responds to an increase in observable effort with a decrease in unobservable effort. Thus, the fact that the overseer requires a level of effort higher than the agency's optimum results in an inefficient overall allocation (the marginal product of observable effort is lower than the marginal product of unobservable effort in equilibrium). This distortion means that, even though both the agency and the overseer prefer an efficient effort allocation, the constraint imposed by the overseer leads the agency to choose an inefficient mix of observable and unobservable effort.

Proposition 1. *If the agency chooses the effort allocation* $(\underline{a}_O, a_U^*(\underline{a}_O))$, *then there is a mix of observable and unobservable effort that Pareto dominates* $(\underline{a}_O, a_U^*(\underline{a}_O))$.

Another way to interpret this result is that the type of oversight we study forces the agency to overemphasize observable forms of effort and underemphasize unobservable forms of effort. Thus, judicial oversight may induce agencies to produce records full of the kind of lawyerly rationality that impresses courts at the expense of other productive activities (McGarity 1992; Shapiro 1988). Similarly, OMB review, and cost-benefit analysis more generally, may distort agency priorities in the direction of those goals that can be easily quantified (Ackerman and Heinzerling 2002; Gillette and Krier 1990).

Oversight and the Quality of Regulation

Importantly, despite the inefficiency that oversight induces when it constrains the agency, the overseer is not acting irrationally by imposing this constraint. Even though the oversight constraint induces an inefficient allocation of agency effort, it nonetheless improves regulatory quality in those cases where the agency decides to regulate. The reason is that oversight compels the agency to increase total effort. For oversight to constrain the agency, it must be the case that the agency's optimal effort, though efficiently allocated, does not produce sufficiently high quality for the overseer to approve the regulation. The quality of constrained regulation is higher than it would have been in the absence of oversight; otherwise the overseer would not approve it. Thus, the overseer demands and receives higher-quality regulation, though the cost of that increased quality is the introduction of inefficiency in the agency's allocation of effort.

Proposition 2. *If the agency is constrained, then the quality of regulation is higher than it would have been in the absence of an overseer.*

Oversight and the Frequency of Regulation

Although oversight may lead to higher quality regulations when the agency is constrained, the overseer's demands may also dissuade the agency from regulating. This implies that agencies initiate regulation less often when their decisions are subject to oversight. Most interestingly, there exist situations in which oversight will dissuade the agency from regulating even though there is an effort allocation that would make both the agency and the overseer better off with new regulation than with the status quo.

As shown in Proposition 1, the threat of oversight can force the agency to choose an inefficient effort allocation, overemphasizing observable effort at the expense of unobservable effort. If this inefficient allocation leaves the agency with a payoff less than its status quo payoff, it will not regulate. However, if the agency could commit to an efficient allocation, it could achieve the same level of quality at lower total cost. If the magnitude of the cost savings gained by increased efficiency were sufficiently large, the agency would be willing to propose regulation, thereby making both players better off. Thus, the inefficiency caused by unobservability sometimes leads to no regulation being proposed, even though it is possible for the agency to produce regulation that both it and the overseer prefer to the status quo.

Proposition 3. *There exist situations in which no regulation is proposed even though there exist effort profiles that make both the agency and the overseer prefer regulation to the status quo.*

It is important to note, however, that although there is a Pareto improving regulation and effort pair, this can never be achieved in equilibrium, even in the absence of oversight. The overseer blocks regulation in these situations precisely because the agency cannot be trusted to choose a high enough level of effort to make regulation attractive to the overseer.

Information Disclosure

A standard intuition is that agencies have an incentive to conceal information from overseers in order to generate bureaucratic slack and minimize the level of effort they are compelled to exert (Banks and Weingast 1992). It might seem that because our model demonstrates that unobservability introduces inefficiency, it offers further support for the idea that agencies should be compelled to disclose otherwise unobservable information. In fact, however, our analysis suggests that agencies have an incentive to disclose information even without additional policy interventions.

The reason is that if some forms of effort are unobservable, the agency has to demonstrate the quality of its proposed regulation to the overseer in an inefficient way, through inflated levels of observable effort. If all effort were observable, the agency could demonstrate sufficient regulatory quality to the overseer at lower cost, because the agency's effort allocation would be efficient. Because the overseer cannot commit to requiring the agency to expend more effort than what is needed for the overseer to prefer regulation to the status quo, the agency has no incentive to conceal information. Thus, when oversight is of the form we model, unobservability is likely to exist because certain types of effort are unobservable in principle or because the cost of making them observable is greater than the efficiency gains associated with doing so. The policy interventions our model would support are therefore not those that target agency incentives to conceal information; rather they are those that improve the ability of overseers to observe effort or that lower the costs to agencies of making effort observable.

Proposition 4. *The agency is sometimes better off and never worse if the effort that is unobservable in our model were made observable.*

Social Welfare

The social welfare implications of the model are contingent on the interpretation given to various components of the payoff functions. For the purposes of building

intuition, we focus on two dimensions of social welfare. Each of these dimensions can be interpreted in two ways, with different implications for social welfare.

The first dimension is the relative value of regulation versus the status quo, about which the overseer and the agency disagree ($q_C \neq q_A$). Under one interpretation, the overseer and agency disagree because the agency is more democratically accountable than the overseer. On this view, it makes sense to assume that society's valuation of the status quo (q_S) is the same as the agency's but different from the overseer's ($q_S = q_A \neq q_C$). A second interpretation argues the opposite. Agencies may be more subject to capture by special interests, or to other forms of bureaucratic drift, than are more insulated overseers such as courts. On this view, the overseer may be more representative of society's underlying values ($q_S = q_C \neq q_A$).

The second dimension is the cost associated with improving the quality of regulation. Under one interpretation, this cost is a social cost, perhaps because costly effort necessitates either additional taxation or diversion of resources away from other socially important regulatory activities. On this view, social welfare should internalize the costs of effort directed at improving regulation. Thus, the social payoff from regulation is $V_S = f(a_O, a_U) - c(a_O + a_U)$. Alternatively, one might think that the agency's effort cost is simply forgone bureaucratic slack. Under this interpretation, these costs should not enter into the social welfare, so the social payoff from regulation is $V_S = f(a_O, a_U)$.

Different interpretations of social welfare are appropriate in different circumstances. Hence, rather than choosing a particular social welfare function, we explore the implications of several candidates. The question of social welfare is uninteresting if society fully agrees with either the agency (in which case oversight is clearly bad) or the overseer (in which case oversight is clearly good). The interesting cases are those where society is aligned with the agency on one dimension and the overseer on the other—that is, when drift is a problem but slack is not, or when slack is a problem but drift is not. We start by analyzing these cases and then consider situations in which society's regulatory preferences do not align perfectly with either the agency or the overseer.

Slack Is a Problem; Drift Is Not. In this case, we assume that society has the same preferences as the agency over the relative value of regulation versus the status quo, but that agency effort is not socially costly (e.g., it is forgone slack). Thus, the social payoff of regulation is $V_S(a_O, a_U) = f(a_O, a_U)$ and the social payoff of the status quo is $q_S = q_A$. One can think of this as a situation in which society is concerned about bureaucratic slack but not bureaucratic drift.

Oversight only has an effect on social welfare if the agency is constrained or dissuaded. As shown in Proposition 2, when the agency is constrained, oversight improves the quality of regulation by increasing total effort. Because social welfare includes the quality of regulation but does not internalize the costs of effort, in this case oversight leads to an improvement in social welfare.

Oversight, however, can also create situations in which the agency is dissuaded from proposing regulation. In this case, social welfare is made worse by oversight. To see why, notice that the fact that the agency would be willing to regulate without oversight means that $f(a_O^*, a_U^*(a_O^*)) - c(a_O^* + a_U^*(a_O^*)) \geq q_A$. The payoff from new regulation with the agency's most preferred allocation of effort, net of the costs of effort, must yield a payoff at least as high as the agency's payoff from the status quo. Because society agrees with the agency on the value of the status quo and does not bear the costs of effort, it certainly would prefer this regulation to the status quo ($f(a_O^*, a_U^*(a_O^*)) > q_S = q_A$).

Thus, in this case there are two possible effects of oversight on social welfare. If the agency is constrained to exert a greater level of effort, oversight improves social welfare. However, if the demands of the overseer are so stringent that the agency is dissuaded from regulating, then oversight diminishes social welfare.

These two possibilities provide some insight into the kind of overseer society would like to have when it is concerned about slack but not drift. The agency will choose not to regulate (i.e., will be dissuaded) if $f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\underline{a}_O + a_U^*(\underline{a}_O)) < q_A$; otherwise the agency will comply with the overseer's demands (i.e., the agency will be constrained). When the agency is constrained, the social payoff is equal to $V_S = f(\underline{a}_O, a_U^*(\underline{a}_O)) = q_C$. Thus, the more skeptical the overseer is of new regulation (the higher q_C), the higher the social welfare, as long as the agency remains willing to regulate. However, as the overseer becomes more stringent, it demands greater levels of observable spending, reducing the net value of regulation from the agency's perspective. If the skeptical overseer's demands become so extreme that the agency is dissuaded from regulating altogether, oversight reduces social welfare.

This intuition is illustrated in Figure 3 and is formalized in the following result.

Proposition 5. *Consider the case where social welfare is given by $V_S(a_O, a_U) = f(a_O, a_U)$ and $q_S = q_A$. Further, suppose there is some effort pair $(a_O, a_U^*(a_O))$ such that society would prefer regulation to the status quo. There is a critical value $\hat{q}_C$ such that social welfare is weakly increasing as the overseer becomes more anti-regulation up to the critical value (i.e., for all $q_C < \hat{q}_C$), but beyond that critical value (i.e., for all $q_C > \hat{q}_C$), social welfare is weakly decreasing as the regulator becomes more anti-regulation.*

A further implication of Proposition 5 is that if society were empowered to choose the overseer directly, it would want an overseer that is more skeptical of new regulation than is society itself. Indeed, society would want the most stringent overseer possible, subject to the constraint that the agency not be dissuaded from regulating ($q_C \leq \hat{q}_C$). In choosing such an overseer, society compels the agency to exert a high level of effort (the costs of which society does not bear), thereby ensuring itself the highest quality regulation possible.

FIGURE 3. As the overseer becomes more stringent (higher q_C) the quality of constrained regulation ($f(\underline{a}_O(q_C), a^*_U(\underline{a}_O(q_C))) = q_C$) increases. Moreover, the social welfare maximizing level of overseer stringency is $q_C = \hat{q}_C$, which is greater than society's level of stringency ($q_S = q_A$). However, if the overseer becomes sufficiently stringent ($q_C > \hat{q}_C$), it demands so much observable spending that the agency chooses not to regulate. Thus, the social welfare is increasing in q_C up to $\hat{q}_C$ but then drops for higher levels of q_C

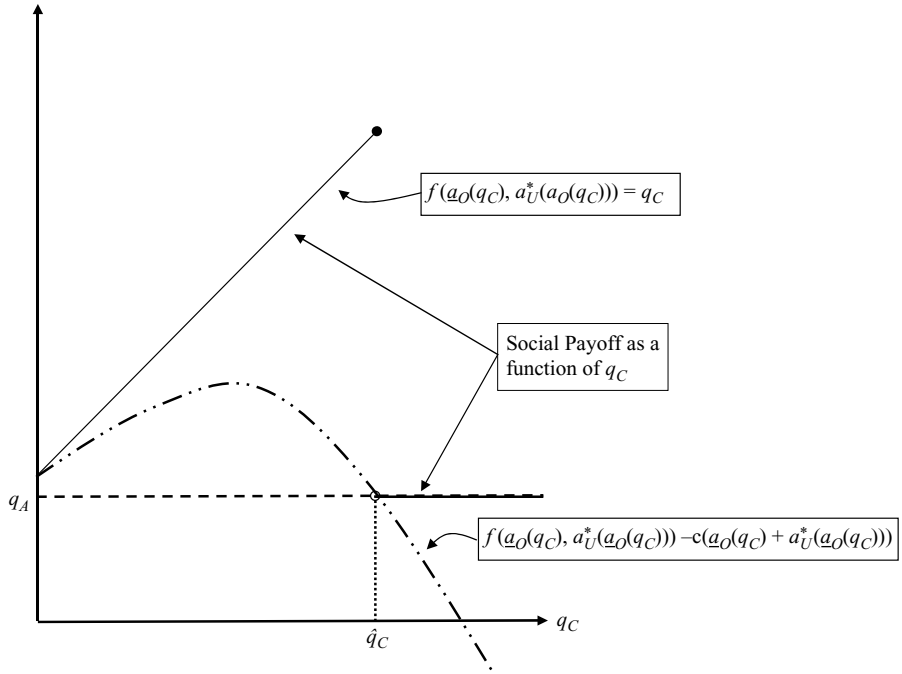

This intuition is formalized in the following result and is illustrated by the fact that the social welfare is maximized at $\hat{q}_C$ in Figure 3.

Proposition 6. *Consider the case where social welfare is given by $V_S(a_O, a_U) = f(a_O, a_U)$ and $q_S = q_A$. Further, suppose there is some effort pair $(a_O, a^*_U(a_O))$ such that society would prefer regulation to the status quo. Social welfare is maximized when the overseer values the status quo more than does society—in particular when $q_C = \hat{q}_C > q_S = q_A$.*

Drift Is a Problem; Slack Is Not. In this case we assume that society has the same preferences as the overseer regarding the relative value of regulation versus the status quo, but that society, like the agency, internalizes the costs of effort. That is, society is concerned with bureaucratic drift but not with bureaucratic slack. The social welfare implications are quite different under this interpretation of the model. In the previous case oversight improved social welfare when it constrained the agency but diminished social welfare when it dissuaded the agency. Here the opposite is true.

Although society shares the overseer's preferences regarding the relative value of new regulation versus the status quo, oversight diminishes social welfare whenever it constrains the agency to choose a higher level of effort. The reason is that the overseer's requirements raise the quality of regulation to a level that makes the overseer exactly indifferent between the regulation and the status quo ($f(\underline{a}_O, a^*_U(\underline{a}_O)) = q_C$). But society, unlike the overseer, internalizes the costs of effort. So, the net social payoff of adopting a regulation that just satisfies the overseer's minimum requirements ($f(\underline{a}_O, a^*_U(\underline{a}_O)) - c(\underline{a}_O + a^*_U(\underline{a}_O) = q_C - c(\underline{a}_O + a^*_U(\underline{a}_O))$ is always lower than the social payoff of the status quo ($q_S = q_C$).

Oversight increases social welfare when it dissuades the agency from proposing regulation, however. If the agency is dissuaded by the overseer, it must be that the social payoff from the status quo ($q_S = q_C$) is less than the social payoff associated with the agency's unconstrained optimum ($f(a^*_O, a^*_U(a^*_O)) - c(a^*_O + a^*_U(a^*_O))$). And, as we have already seen, the social payoff from the status quo is higher than at the constrained choice. Hence, given that the overseer is interested in altering the agency's behavior, society wants the overseer to force the agency not to regulate at all.

This has the added implication that the results from Propositions 5 and 6 are reversed here. Those Propositions establish that, when society is concerned about slack but not drift, the more conservative (i.e., skeptical of regulation) the overseer is, the better off society is, so long as oversight constrains but does not dissuade the agency. In contrast, where society is concerned about drift but not slack, a more conservative overseer makes society worse off, unless the overseer is so demanding that it dissuades the agency from regulating altogether.

The combination of these two cases suggests that oversight of the sort we model is a valuable tool when society is concerned primarily about slack, but not when society is concerned primarily about drift. The stringency of the overseer can have a positive effect on social welfare when slack is at issue, because an overseer can compel the agency to exert more effort. However, when drift is the issue, oversight is weakly dominated, from the social perspective, either by allowing the agency free rein or by banning regulation altogether; constraining the agency is always bad in this case.

Society's Preferences Located between the Agency and the Overseer. Of course, society's regulatory preferences need not conform exactly to those of the agency or the overseer. Consider, then, the possibility that society's preference for the status quo is somewhere between the agency's and the overseer's (i.e., $q_A < q_S < q_C$).[15] The intuitions built in the previous sections apply here in only slightly altered form.

Suppose, first, that society does not internalize the costs of agency effort (i.e., these costs are forgone slack). Just as in the case where society was concerned about drift but not slack, social welfare here includes the quality of regulation but not the costs of effort. As such, when the agency is constrained by the overseer, social welfare is improved because the quality of regulation is improved.

However, the effect on social welfare of the overseer dissuading the agency is slightly different than it was when society had the same regulatory preferences as the agency. In that earlier case, dissuasion was always bad for society. Now, it is possible for dissuasion to be good for society. Recall that, for the agency to be willing to regulate in the absence of oversight, it must be the case that $f(a_O^*, a_U^*(a_O^*)) - c(a_O^* + a_U^*(a_O^*)) \geq q_A$. Society would prefer the agency's most preferred regulation to the status quo only if $f(a_O^*, a_U^*(a_O^*)) \geq q_S$. Because society likes the status quo more than the agency does ($q_S > q_A$), it is possible for the agency to be willing to regulate in the absence of oversight even if society prefers the status quo to the agency's optimal regulation. This occurs if

$$q_s \geq f(a_O^*, a_U^*(a_O^*)) > f(a_O^*, a_U^*(a_O^*)) - c(a_O^* + a_U^*(a_O^*)) \geq q_A.$$

From this, we can see that society benefits from dissuasion when it values the status quo relatively highly ($q_S \geq f(a_O^*, a_U^*(a_O^*))$). This difference notwithstanding, the results in Propositions 5 and 6 continue to hold here. Society always prefers higher quality regulation and, regardless of the value of q_S, society always prefers the constrained outcome to the status quo. Hence, society still wants as conservative an overseer as possible, up to the point where the agency is dissuaded from regulating.

Now suppose that society does internalize the costs of effort (i.e., these costs are social opportunity costs). Here the social welfare results are similar to those when society is concerned about drift but not slack. Because society internalizes the costs of effort, the net payoff to society from regulation is maximized when the agency's effort allocation is unconstrained; the payoff to both society and the agency from this regulation is $f(a_O^*, a_U^*(a_O^*)) - c(a_O^* + a_U^*(a_O^*))$. Hence, constraint is always bad for society, just as in the earlier case. The difference is that oversight need not increase social welfare when it dissuades the agency. If the agency is dissuaded, the overseer's payoff from the status quo (q_C) must be greater than the social payoff associated with the agency's unconstrained optimum ($f(a_O^*, a_U^*(a_O^*)) - c(a_O^* + a_U^*(a_O^*))$). However, society values the status quo less than the overseer does ($q_S < q_C$). If q_S is sufficiently low, then society would have preferred the agency not to be dissuaded. That said, the central insight of the case when drift is a problem but slack is not continues to hold even if $q_S \neq q_C$: oversight is weakly dominated either by eliminating oversight or by banning regulation.

CONCLUSION

We have developed a model in which a policymaking agency must get the approval of an overseer with the power to reject regulation (e.g., a court or OMB) when the overseer can observe some, but not all, of the activities that the policymaking agency may employ to improve regulatory quality. This framework is consistent, for example, with the idea that a reviewing court can observe adherence to lawful procedures but not the soundness of an agency's technical analysis, or with the claim that OMB oversight is effective at assessing an agency's analysis of easily quantifiable costs and benefits but ineffective at evaluating an agency's attention to other important aspects of regulatory design.

Under these conditions, oversight increases the quality of proposed regulations, reduces the frequency of regulation, and distorts the policymaking agency's effort allocation toward those tasks that the overseer can observe. This last effect introduces an inefficiency that both the agency and the overseer would prefer to eliminate. This, coupled with the overseer's inability to credibly demand regulation whose quality provides a payoff strictly higher than the status quo, suggests that in many circumstances agencies have an incentive to disclose as much information as possible about the agency's activities, even though the agency values bureaucratic slack. Eliminating this source of regulatory inefficiency may therefore depend not on compelling agencies to disclose more information, but on improving agencies' ability to communicate such information credibly and cost-effectively to overseers.

[15] The case where $q_A > q_C$ is trivial, because oversight never has any impact on agency behavior in that case. Furthermore, the results for a q_S located outside of $[q_A, q_C]$ follow straightforwardly from the earlier cases. If the society does not internalize the cost of agency effort and $q_S < q_A$, then society benefits from oversight that constrains the agency, but society is harmed if oversight dissuades the agency. Society thus wants as conservative an overseer as possible, as long as the agency is not dissuaded. If $q_S > q_C$ and society internalizes the cost of effort, then society strictly prefers dissuasion to constraint. The discussion in the text focuses on the more interesting cases where $q_A < q_S < q_C$.

We further show that the social welfare implications of bureaucratic oversight depend on context—in particular, on the relative magnitudes of the bureaucratic drift and slack problems. Thus, when agency preferences regarding the benefits of regulation track social preferences but agency effort is not socially costly (i.e., there is bureaucratic slack but not bureaucratic drift), introducing regulatory oversight by an overseer that is more biased against regulation than society can improve social welfare, but only if the overseer is not so skeptical of regulation that its demands dissuade the agency from acting at all. In contrast, when society's policy preferences align with those of the overseer but agency effort is socially costly (i.e., there is drift but no slack), then regulatory oversight is weakly dominated by banning either regulation or oversight altogether.

These preliminary findings and the basic framework we develop suggest directions for further research into the effect of imperfect oversight by courts, OMB, or similar institutions on bureaucratic performance. Future work could, for example, explore the possibility for credible commitment by the overseer to a more stringent oversight strategy; the use of multiple overseers; differences in the marginal impact of different agency activities on the respective regulatory payoffs to the agency and the overseer; and a more extensive analysis of the social welfare implications of various forms of oversight. These and other avenues of inquiry may further enrich our understanding of the effects and limitations of what, we hope we have shown, is an interesting and comparatively understudied aspect of bureaucratic governance.

APPENDIX

Proof of Lemma 1

If the agency proposes regulation that does not pass, its payoff is $q_A - c(a_O + a_U)$, which is less than or equal to q_A, the payoff from not proposing regulation. ■

Proof of Lemma 2

The choice of unobservable effort has no effect on whether beliefs are such that regulation is upheld. Thus, if beliefs are such that regulation will be upheld, then the agency will choose the utility maximizing level of unobservable effort. This is characterized by the first-order condition in equation (1).

Applying the implicit function theorem to the first-order condition, and using the fact that $f_{OU} = 0$, we have

$$\frac{\partial a_U^*}{\partial a_O} = \frac{c''}{f_{UU} - c''} < 0, \tag{4}$$

where the inequality follows from the concavity of f and the convexity of c. Differentiating this, and using the fact that $f_{OU} = 0$, we have

$$\begin{aligned}\frac{\partial^2 a_U^*}{\partial a_O^2} &= \frac{c'''}{f_{UU} - c''} + \frac{c''c'''}{(f_{UU} - c'')^2} \\ &= \frac{c'''}{f_{UU} - c''}\left(1 - \frac{c''}{c'' - f_{UU}}\right) \le 0,\end{aligned}$$

where the inequality follows from $c''' \ge 0$, $c'' > 0$, and the concavity of f. ■

Proof of Lemma 3

The overseer upholds if inequality (2) holds. The right-hand side is not a function of a_O. Taking the derivative of the left-hand side with respect to a_O yields

$$\frac{\partial f(a_O, a_U^*(a_O))}{\partial a_O} = f_O + f_U \frac{\partial a_U^*}{\partial a_O} = f_O + f_U \frac{c''}{f_{UU} - c''}. \tag{5}$$

Differentiating again yields

$$\begin{aligned}\frac{\partial^2 f(a_O, a_U^*(a_O))}{\partial a_O^2} &= f_{OO} + 2f_{OU}\frac{\partial a_U^*}{\partial a_O} + f_{UU}\left(\frac{\partial a_U^*}{\partial a_O}\right)^2 + \frac{\partial^2 a_U^*}{\partial a_O^2} f_U \\ &= f_{OO} + f_{UU}\left(\frac{\partial a_U^*}{\partial a_O}\right)^2 + \frac{\partial^2 a_U^*}{\partial a_O^2} f_U < 0,\end{aligned}$$

where the second line uses the fact that $f_{OU} = 0$, and the inequality follows from the concavity of f and Lemma 2. This implies that $f(a_O, a_U^*(a_O))$ is either (1) decreasing everywhere, (2) increasing everywhere, or (3) single-peaked.

A lower bound other than 0 is established when $f(a_O, a_U^*(a_O))$ crosses q_C from below. If $f(a_O, a_U^*(a_O))$ is everywhere decreasing, then it never crosses from below. Thus, if $f(a_O, a_U^*(a_O)) \ge q_C$ the lower bound is 0; otherwise there is no a_O that satisfies equation (2). If $f(a_O, a_U^*(a_O))$ is increasing everywhere or not monotonic, but single peaked, then $f(0, a_U^*(0))$ crosses from below at most once. If it never crosses from below, then if $f(0, a_U^*(0)) \ge q_C$ the lower bound is 0; otherwise there is no a_O that satisfies equation (2). If it does cross from below, then $\underline{a}_O$ is the point where it crosses. ■

Proof of Lemma 4

The proof is by contradiction. Suppose $a_O^* > \bar{a}_O$.

Using the derivative from equation (5) and the fact, from Lemma 3, that if there is an upper bound, $f(a_O, a_U^*(a_O))$ is decreasing to its right, the following holds:

$$f_O(\bar{a}_O, a_U^*(\bar{a}_O)) < -f_U(\bar{a}_O, a_U^*(\bar{a}_O))\frac{\partial a_U^*}{\partial a_O}.$$

By concavity, $a_O^* > \bar{a}_O$ implies that $f_O(a_O^*, a_U^*(a_O^*)) < f_O(\bar{a}_O, a_U^*(\bar{a}_O))$.

Using the fact (from equation (4)) that $\frac{\partial a_U^*}{\partial a_O} < 0$, $a_O^* > \bar{a}_O$ implies that $a_U^*(a_O^*) < a_U^*(\bar{a}_O)$. Concavity then implies that $-f_U(a_O^*, a_U^*(a_O^*)) < -f_U(\bar{a}_O, a_U^*(\bar{a}_O))$. Further, the fact (from Lemma 2) that $\frac{\partial^2 a_U^*}{\partial a_O^2} \le 0$ implies that $\frac{\partial a_U^*(a_O^*)}{\partial a_O} < \frac{\partial a_U^*(\bar{a}_O)}{\partial a_O}$. Putting these two together we have $-f_U(a_O^*, a_U^*(a_O^*))\frac{\partial a_U^*(a_O^*)}{\partial a_O} > -f_U(\bar{a}_O, a_U^*(\bar{a}_O))\frac{\partial a_U^*(\bar{a}_O)}{\partial a_O}$.

Finally, since it is clear that $-1 < \frac{c''}{f_{UU} - c''} = \frac{\partial a_U^*}{\partial a_O} < 0$, and $c' > 0$, we have $\left(1 + \frac{\partial a_U^*}{\partial a_O}\right)c' > 0$.

Putting all of this together implies

$$\begin{aligned}&f_O(a_O^*, a_U^*(a_O^*)) < f_O(\bar{a}_O, a_U^*(\bar{a}_O)) < \\ &-f_U(\bar{a}_O, a_U^*(\bar{a}_O))\frac{\partial a_U^*(\bar{a}_O)}{\partial a_O} < -f_U(a_O^*, a_U^*(a_O^*))\frac{\partial a_U^*(a_O^*)}{\partial a_O} < \\ &-f_U(a_O^*, a_U^*(a_O^*))\frac{\partial a_U^*(a_O^*)}{\partial a_O} + \left(1 + \frac{\partial a_U^*}{\partial a_O}\right)c'(a_O^* + a_U^*(a_O^*)).\end{aligned}$$

But, from equation (3) we know that

$$f_O(a_O^*, a_U^*(a_O^*)) = -f_U(a_O^*, a_U^*(a_O^*))\frac{\partial a_U^*(a_O^*)}{\partial a_O} + \left(1 + \frac{\partial a_U^*}{\partial a_O}\right) c'(a_O^* + a_U^*(a_O^*)),$$

a contradiction. ■

Proof of Lemma 5

The first and third cases are immediate from unconstrained optimization in the subgame. In the middle case, the agency's most preferred choice would not be upheld by the overseer, so the optimization is constrained. Lemma 4 implies that if $a_O^* \notin \Phi$, then $a_O^* < \underline{a}_O$. To see that $\underline{a}_O$ is the constrained optimum, we use the following claim.

Claim 1. *For all* $a_O > a_O^*$, $\frac{\partial V(a_O, a_U^*(a_O))}{\partial a_O} < 0$.

Given the claim and the fact that $a_O^* < \underline{a}_O$, the agency's utility is strictly decreasing as it increases its level of observable effort from a_O^*. This implies that when $a_O^* \notin \Phi$, the agency's utility is decreasing over Φ, so if the agency chooses a positive a_O it will choose the lower bound of Φ, which is $\underline{a}_O$. This establishes the result. All that remains is to prove the claim.

Proof of Claim

We know that at a_O^*, $\frac{\partial V(a_O, a_U^*(a_O))}{\partial a_O} = 0$. Thus, it is sufficient to prove that $V(a_O, a_U^*(a_O))$ is globally concave. To see this, note that the second derivative is

$$\frac{\partial^2 V(a_O, a_U^*(a_O))}{\partial a_O^2} = f_{OO} + f_{OU}\frac{\partial a_U^*}{\partial a_O} + \left(f_{UU}\frac{\partial a_U^*}{\partial a_O} + f_{OU}\right)\frac{\partial a_U^*}{\partial a_O} + f_U\frac{\partial^2 a_U^*}{\partial a_O^2} - c''\left(1 + \frac{\partial a_U^*}{\partial a_O}\right)^2 - c'\frac{\partial^2 a_U^*}{\partial a_O^2},$$

which, using the facts that for any a_O, $f_U(a_O, a_U^*) = c'(a_O + a_U^*)$ and $f_{OU} = 0$, can be rewritten

$$f_{OO} + f_{UU}\left(\frac{\partial a_U^*}{\partial a_O}\right)^2 - c''\left(1 + \frac{\partial a_U^*}{\partial a_O}\right)^2 < 0,$$

where the inequality follows from the concavity of f and the convexity of c. ■

Proof of Proposition 1

We make use of the following claim.

Claim 2. *For any level of spending, a, the quality of regulation is maximized for the allocation (a_O, a_U) such that $f_O(a_O, a_U) = f_U(a_O, a_U)$.*

If the agency chooses $\underline{a}_O$, then $\underline{a}_O > a_O^*$. Lemma 2, then, implies that $a_U^*(\underline{a}) < a_U^*(a_O^*)$. Optimality implies that $f_O(a_O^*, a_U^*(a_O^*)) = f_U(a_O^*, a_U^*(a_O^*))$. These and concavity imply that $f_O(\underline{a}_O, a_U^*(\underline{a}_O)) < f_U(\underline{a}_O, a_U^*(\underline{a}_O))$. But then the Claim implies that there is a redivision of the same level of resources that would improve the quality of regulation, establishing the result.

All that remains is to prove the claim

Proof of Claim

Holding expenditures fixed implies the following budget constraint $a_O + a_U = a$. The constrained optimum can be found from the following Lagrangian:

$$\max_{a_O, a_U} f(a_O, a_U) - \lambda(a_O + a_U - a).$$

The Kuhn–Tucker conditions are $f_i(a_O, a_U) = \lambda$, for $i \in \{O, U\}$, which implies that at the optimum $f_O(a_O, a_U) = f_U(a_O, a_U)$. ■

Proof of Proposition 2

If the agency is constrained, it chooses $\underline{a}_O$. By the definition of constraint, we know that $f(a_O^*, a_U^*(a_O^*)) < q_C$. Moreover, by the definition of $\underline{a}_O$, the quality of constrained regulation is $f(\underline{a}_O, a_U^*(\underline{a}_O)) = q_C$. Thus, $f(a_O^*, a_U^*(a_O^*)) < f(\underline{a}_O, a_U^*(\underline{a}_O))$. ■

Proof of Proposition 3

Let $\hat{a} = \underline{a}_O + a_U^*(\underline{a}_O)$ such that $\underline{a}_O > a_O^*$. By Proposition 1, there exists an efficient allocation $(\hat{a}_O, \hat{a}_U)$ such that $f_O(\hat{a}_O, \hat{a}_U) = f_U(\hat{a}_O, \hat{a}_U)$ and $\hat{a}_O + \hat{a}_U = \hat{a}$. By the definition of efficiency, $f(\hat{a}_O, \hat{a}_U) > f(\underline{a}_O, a_U^*(\underline{a}_O))$. Further, by the definition of $\hat{a}$, $c(\hat{a}_O + \hat{a}_U) = c(\hat{a}) = c(\underline{a}_O + a_U^*(\underline{a}_O))$. Thus, there exist values of q_C and q_A such that $f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\hat{a}) < q_A \leq f(\hat{a}_O, \hat{a}_U) - c(\hat{a})$ and $f(\underline{a}_O, a_U^*(\underline{a}_O)) = q_C < f(\hat{a}_O, \hat{a}_U)$. ■

Proof of Proposition 4

There are three cases to consider. If the agency is unconstrained, then there is no difference between the two games.

If the agency is constrained, then with unobservable effort it choose the allocation $(\underline{a}_O, a_U^*(\underline{a}_O))$, and its payoff is $f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\hat{a})$, where $\hat{a}$ is as defined in the proof of Proposition 3. When the unobservable effort is observable, the agency can choose the allocation $f(\hat{a}_O, \hat{a}_U) > q_C$ as defined in the proof of Proposition 3 and still have regulation upheld. Its payoff then is $f(\hat{a}_O, \hat{a}_U) - c(\hat{a}) > f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\hat{a})$. Moreover, it can even improve on this by decreasing total effort, since $f(\hat{a}_O, \hat{a}_U) > q_C$.

If the agency is dissuaded, then its payoff with unobservable effort is q_A. When unobservable effort is observable, either the agency can choose an efficient allocation that induces the overseer to uphold or it will still accept the status quo. It will only choose the efficient allocation if the payoff is at least as good as the status quo. ■

Proof of Proposition 5

According to Lemma 3, there are three cases to consider.

Case 1: $f(a_O, a_U^*(a_O))$ is everywhere decreasing in $\underline{a}_O$.

In this case it is clear that if there is any pair that will satisfy the overseer, then $\underline{a}_O = 0$. This creates two subcases.

Subcase A.1: $f(0, a_U^*(0)) \geq 0$ and $f(0, a_U^*(0)) - c(0 + a_U^*(0)) \geq q_A$.

In this case for $q_C \in [0, f(0, a_U^*(0))]$, the agency is unconstrained. For $q_C > f(0, a_U^*(0))$, the overseer will veto any regulation. Thus, for $q_C \leq f(0, a_U^*(0)) \equiv \hat{q}_C$, social welfare is $f(0, a_U^*(0)) > q_A$ and for $q_C > \hat{q}_C$, social welfare is q_A. Thus, social welfare is weakly increasing up to $\hat{q}_C$ and weakly decreasing for $q_C > \hat{q}_C$.

Subcase A.2: Either of the two conditions from A.1 do not hold.

In this case, then regulation will never be passed and q_C has no effect on social welfare.

Case 2: $f(a_O, a_U^*(a_O))$ is increasing in a_O up to some $\overline{\overline{a}}_O$ and decreasing in a_O for $a_O \geq \overline{\overline{a}}_O$.

There are two subcases to consider. We will use the following notation. Let $\underline{q}_C$ be the q_C such that $\underline{a}_O(\underline{q}_C) = a_O^*$, so that for all $q_C < \underline{q}_C$, the agency is unconstrained. We also note that, since by hypothesis there exist effort pairs $(a_O, a_U^*(a_O))$ that would make society want regulation, it must be that $f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O)) \geq q_A = q_S$.

Subcase B.1: $f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O)) - c(\overline{\bar{a}}_O + a_U^*(\overline{\bar{a}}_O)) \geq q_A$.

In this case, the agency is never dissuaded. For all $q_C < \underline{q}_C$, the agency chooses a_O^*. For $q_C \in [\underline{q}_C, f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O))]$, the agency chooses $a_O = \underline{a}_O(q_C)$ and so social welfare is equal to q_C. For $q_C > f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O))$ no regulation is ever accepted, so social welfare is q_A.

Let $\hat{q}_C = f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O))$. Thus, for $q_C \leq \underline{q}_C$, social welfare is flat in q_C. For $q_C \in [\underline{q}_C, \hat{q}_C]$, social welfare is strictly increasing in q_C. For $q_C > \hat{q}_C$ social welfare is constant at $q_A = f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O)) - c(\overline{\bar{a}}_O + a_U^*(\overline{\bar{a}}_O))$ which is obviously less than $f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O))$, which was the social payoff at $q_C = \hat{q}_C$.

Subcase B.2: $f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O)) - c(\overline{\bar{a}}_O + a_U^*(\overline{\bar{a}}_O)) < q_A$.

We know from hypothesis that there exists a q_C' such that $f(\underline{a}_O(q_C'), a_U^*(\underline{a}_O(q_C'))) - c(\underline{a}_O(q_C') + a_U^*(\underline{a}_O(q_C'))) \geq q_A$. If this were not true, there would be no pair $(a_O, a_U^*(a_O))$ such that society wanted regulation. Let $\overline{q}_C'$ be the q_C such that $f(\underline{a}_O(\overline{q}_C'), a_U^*(\underline{a}_O(\overline{q}_C'))) - c(\underline{a}_O(\overline{q}_C') + a_U^*(\underline{a}_O(\overline{q}_C'))) = q_A$. The proof of Lemma 5 shows that $f(a_O, a_U^*(a_O)) - c(a_O + a_U^*(a_O))$ is globally concave. Moreover, by the definition of constraint, $\underline{a}_O > a_O^*$, so evaluated at $\underline{a}_O(\overline{q}_C')$, $f(a_O, a_U^*(a_O)) - c(a_O + a_U^*(a_O))$ is decreasing in a_O. This implies that for all $a_O > \underline{a}_O(\overline{q}_C')$, the agency is dissuaded.

Let $\overline{q}_C' = \hat{q}_C$. Then for $q_C \leq \underline{q}_C$, the agency is unconstrained and social welfare is fixed at $f(a_O^*, a_U^*(a_O^*))$. For $q_C \in [\underline{q}_C, \hat{q}_C]$, the agency is constrained and social welfare is q_C. For $q_C > \hat{q}_C$, the agency is dissuaded and social welfare is q_A.

Thus, for $q_C \leq \underline{q}_C$ social welfare is unaffected by q_C. For $q_C \in [\underline{q}_C, \hat{q}_C]$, social welfare is strictly increasing in q_C. For $q_C > \hat{q}_C$ social welfare is constant at $q_A = f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C))) - c(\underline{a}_O(\hat{q}_C) + a_U^*(\underline{a}_O(\hat{q}_C)))$, which is obviously less than $f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C)))$, which was the social payoff at $q_C = \hat{q}_C$.

Case 3: $f(a_O, a_U^*(a_O))$ is increasing for all a_O.

Since $f(a_O, a_U^*(a_O))$ is increasing in a_O and $f(\underline{a}_O, a_U^*(\underline{a}_O)) = q_C$, $\underline{a}_O$ is increasing in q_C. Moreover, since $\underline{a}_O > a_O^*$, by Lemma 5, $f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\underline{a}_O + a_U^*(\underline{a}_O))$ is decreasing in $\underline{a}_O$, which means it is decreasing in q_C. Thus, since $f(\underline{a}_O, a_U^*(\underline{a}_O)) - c(\underline{a}_O + a_U^*(\underline{a}_O))$ is continuous, then as q_C increases from $\underline{q}_C$ to infinity, the intermediate value theorem implies that there exists a $\hat{q}_C$ such that $f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C))) - c(\underline{a}_O + a_U^*(\underline{a}_O(\hat{q}_C)) = q_A$.

For $q_C \leq \underline{q}_C$, the agency is unconstrained and social welfare is $f(a_O^*, a_U^*(a_O^*))$. For $q_C \in [\underline{q}_C, \hat{q}_C]$, the agency is constrained and the social welfare is q_C. For $q_C > \hat{q}_C$, the agency is dissuaded and the social welfare is q_A.

Thus, for $q_C \leq \underline{q}_C$ is flat in q_C. For $q_C \in [\underline{q}_C, \hat{q}_C]$, social welfare is strictly increasing in q_C. For $q_C > \hat{q}_C$ social welfare is constant at $q_A = f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C))) - c(\underline{a}_O(\hat{q}_C) + a_U^*(\underline{a}_O(\hat{q}_C)))$, which is obviously less than $f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C)))$, which was the social payoff at $q_C = \hat{q}_C$. ■

Proof of Proposition 6

Consider the cases from the proof of Proposition 5.

Case 1: In case 1, the agency is never constrained.

Case 2: There are 2 subcases.

B.1 $\hat{q}_C = f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O)) > f(\overline{\bar{a}}_O, a_U^*(\overline{\bar{a}}_O)) - c(\overline{\bar{a}}_O + a_U^*(\overline{\bar{a}}_O)) \geq q_A$.

B.2 $\hat{q}_C = f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C))) > f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C))) - c(\underline{a}_O(\hat{q}_C) + a_U^*(\underline{a}_O(\hat{q}_C))) = q_A$.

Case 3: $\hat{q}_C = f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C))) > f(\underline{a}_O(\hat{q}_C), a_U^*(\underline{a}_O(\hat{q}_C))) - c(\underline{a}_O(\hat{q}_C) + a_U^*(\underline{a}_O(\hat{q}_C))) = q_A$. ■

REFERENCES

Ackerman, Frank, and Lisa Heinzerling. 2002. "Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection." *University of Pennsylvania Law Review* 150 (May): 1553–84.

Aranson, Peter H., Ernest Gellhorn, and Glen O. Robinson. 1983. "A Theory of Legislative Delegation." *Cornell Law Review* 68 (November): l–67.

Banks, Jeffrey S., and Barry R. Weingast. 1992. "The Political Control of Bureaucracies Under Asymmetric Information." *American Journal of Political Science* 36 (May): 509–24.

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

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

Boehmke, Fred, Sean Gailmard, and John Patty. 2006. "Whose Ear to Bend? Information Sources and Venue Choice in Policy-Making." *Quarterly Journal of Political Science* l (March): 139–69.

Breyer, Stephen. 1986. "Judicial Review of Questions of Law and Policy." *Administrative Law Review* 38 (Fall): 363–98.

Bueno de Mesquita, Ethan. 2007. "Politics and the Suboptimal Provision of Counterterror." *International Organization.* 61 (Winter): 9–36.

Cooper, Joseph, and Williams F. West. 1988. "Presidential Power and Republican Government: The Theory and Practice of OMB Review of Agency Rules." *Journal of Politics* 50 (November): 864–95.

Cross, Frank B. 2000. "Pragmatic Pathologies of Judicial Review of Administrative Rulemaking." *North Carolina Law Review* 78 (April): 1013–69.

de Figueiredo, Rui J. P., Pablo T. Spiller, and Santiago Urbiztondo. 1999. "An Informational Perspective on Administrative Procedures." *Journal of Law, Economics, and Organization* 15 (April): 283–305.

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

Eckert, Ross D. 1981. "The Life Cycle of Regulatory Commissioners." *Journal of Law and Economics* 24 (April): 113–20.

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

Epstein, Richard A. 1995. *Simple Rules for a Complex World.* Cambridge: Harvard University Press.

Gailmard, Sean. 2006. "Discretion Rather that Rules: Choice of Instruments to Constrain Bureaucratic Policy-Making." Unpublished manuscript, Northwestern University.

Gailmard, Sean, and John W. Patty. 2007. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51 (October). Forthcoming.

Gillette, Clayton P., and James E. Krier. 1990. "Risk, Courts, and Agencies." *University of Pennsylvania Law Review* 138 (April): 1027–09.

Gordon, Sanford C., and Catherine Hafer. 2005. "Flexing Muscle: Corporate Political Expenditures as Signals to the Bureaucracy." *American Political Science Review* 99 (May): 245–61.

Hayek, Friedrich A. 1944. *The Road to Serfdom.* Chicago: University of Chicago Press.

Heclo, Hugh. 1988. "The In-and-Outer System: A Critical Assessment." *Political Science Quarterly* 103 (Spring): 37–56.

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

Hopenhayn, Hugo, and Susanne Lohmann. 1996. "Fire Alarm Signals and the Political Oversight of Regulatory Agencies." *Journal of Law, Economics, and Organization* 12 (April): 196–213.

Horn, Murray J., and Kenneth A. Shepsle. 1989. "Administrative Process and Organizational Form as Legislative Responses to Agency Costs." *Virginia Law Review* 75 (March): 499–508.

Lawson, Gary. 1994. "The Rise and Rise of the Administrative State." *Harvard Law Review* 107 (April): 1231–54.

Lowi, Theodore J. 1979. *The End of Liberalism: The Second Republic of the United States*. New York: Norton.

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

McCubbins, Mathew D., Roger G. Noll, and Barry R. Weingast. 1989. "Structure and Process, Politics and Policy: Administrative Arrangements and the Political Control of Agencies." *Virginia Law Review* 75 (March): 431–82.

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

Moe, Terry M. 1990. "Political Institutions: The Neglected Side of the Story." *Journal of Law, Economics, and Organization* 6 (SI): 213–54.

Morrison, Alan B. 1986. "OMB Interference with Agency Rulemaking: The Wrong Way to Write a Regulation." *Harvard Law Review* 99 (March): 1059–74.

Niskansen, William A. 1971. *Bureaucracy and Representative Government.* Chicago: Aldine, Atherton.

Pierce, Richard J., Jr. 1995. "Seven Ways to Deossify Agency Rulemaking." *Administrative Law Review* 47 (Winter): 59–95.

Rachlinski, Jeffrey J., and Cynthia R. Farina. 2002. "Cognitive Psychology and Optimal Government Design." *Cornell Law Review* 87 (January): 549–615.

Seidenfeld, Mark. 1997. "Demystifying Deossification: Rethinking Recent Proposals to Modify Judicial Review of Notice and Comment Rulemaking." *Texas Law Review* 75 (February): 483–524.

Seidenfeld, Mark. 2002. "The Psychology of Accountability and Political Review of Agency Rules." *Duke Law Journal* 18 (December): 1051–95.

Shapiro, Martin. 1988. *Who Guards the Guardians?: Judicial Control of Administration.* Athens: University of Georgia Press.

Shepsle, Kenneth A. 1992. "Bureaucratic Drift, Coalitional Drift, and Time Consistency." *Journal of Law, Economics, and Organization* 8 (March): 111–18.

Spence, David B., and Frank Cross. 2000. "A Public Choice Case for the Administrative State." *Georgetown Law Journal* 89 (November): 97–142.

Stephenson, Matthew C. 2006a. "A Costly Signaling Theory of 'Hard Look' Judicial Review." *Administrative Law Review* 58 (Fall): 753–813.

Stephenson, Matthew C. 2006b. "Legislative Allocation of Delegated Power: Uncertainty, Risk, and the Choice Between Agencies and Courts." *Harvard Law Review* 119 (February): 1035–70.

Stephenson, Matthew C. 2007. "Bureaucratic Decision Costs and Endogenous Agency Expertise." *Journal of Law, Economics, and Organization*. Forthcoming.

Stigler, George J. 1971. "The Theory of Economic Regulation." *Bell Journal of Economics and Management Science* 2 (Spring): 3–21.

Sunstein, Cass. 1984. "In Defense of the Hard Look: Judicial Activism and Administrative Law." *Harvard Journal of Law and Public Policy* 7 (Winter): 51–9.

Ting, Michael M. 2001. "The 'Power of the Purse' and its Implications for Bureaucratic Policy-Making." *Public Choice* 106 (March): 243–74.

Wagner, Wendy E. 2003. "The 'Bad Science' Fiction: Reclaiming the Debate Over the Role of Science in Public Health and Environmental Regulation." *Law and Contemporary Problems* 66 (Autumn): 63–133.

Weingast, Barry R., and Mark J. Moran. 1983. "Bureaucratic Discretion or Congressional Control? Regulatory Policymaking by the Federal Trade Commission." *Journal of Political Economy* 91 (October): 765–800.

Wiseman, Alan. 2007. "Delegation and Bureaucratic Policymaking with Competing Oversight." Unpublished Manuscript.