Bureaucratic Decision Costs and Endogenous Agency Expertise *

Matthew C. Stephenson†

March 20, 2006

Abstract

Bureaucratic expertise is typically the endogenous product of an administrative agency's choices, and these choices are influenced by the incentive structures created by legislative, judicial, and executive overseers. This paper explores the implications of this claim by developing a formal model to investigate two related questions. First, how do the costs of adopting a new regulation (*enactment costs*) affect an agency's investment in expertise? Second, what are the *optimal enactment costs* from the perspective of an uninformed overseer, such as a court or legislature? On the first question, the model shows that the effect of enactment costs on expertise depends on what an uninformed agency would do: If the uninformed agency would regulate, increasing enactment costs increases agency expertise, while if the uninformed agency would choose the status quo, increasing enactment costs decrease agency expertise. On the second question, the model demonstrates that the overseer must balance its interest in influencing the agency's policy preferences against its interest in increasing the agency's expertise. The overseer's optimal enactment will reflect both of these competing considerations, and this can lead to counterintuitive predictions that differ from the predictions of extant theories that assume agency expertise is exogenous. The paper explores the substantive implications of these results for an array of topics in institutional design and public law, including judicial and executive review of regulations, structure-and-process theories of congressional oversight, national security, criminal procedure, and constitutional law.

*I am grateful to Richard Fallon, Jack Goldsmith, Howell Jackson, Louis Kaplow, Adriaan Lanni, Daryl Levinson, John Manning, Martha Minow, Rick Pildes, Fred Schauer, and Ken Shepsle for helpful comments on earlier drafts.

†Assistant Professor of Law, Harvard Law School, Cambridge, MA 02138. Email:mstephen@law.harvard.edu

The delegation of substantial policymaking authority to administrative agencies is often both explained and justified by the belief that agencies have more accurate information about the actual impacts of different policy choices. Consider, for example, the decision whether to ban the manufacture and sale of a toxic substance like asbestos. A common argument for delegating this decision to the Environmental Protection Agency (EPA), rather than leaving the decision to Congress, is that the EPA has greater expertise about the likely effects of the proposed ban, including more accurate estimates of the projected health benefits and economic costs. At the same time, delegation entails the risk that agencies will exploit their policy-making discretion to pursue goals that diverge from those of the electorate and its representatives. The EPA, for example, might be more zealous than the median member of Congress, leading the agency to ban asbestos under circumstances where Congress, if fully informed, would not. The informational asymmetry that justified the delegation in the first place makes it difficult for Congress, courts, or other overseers to monitor the agency.

A rich literature in political science, economics, and law considers institutional mechanisms that less-informed overseers, such as politicians and courts, may employ to induce better-informed agencies to make decisions that more closely track the overseer's policy preferences. This literature, however, typically assumes that agency expertise is exogenous—a given characteristic of the agency that is independent of the scope of the delegation, other aspects of the institutional environment, and the agency's own choices. That assumption, though often a useful simplification, is problematic. Though we may say that the EPA has expertise regarding environmental regulation as a general matter, the EPA may only be able to learn about the likely effects of a specific proposal, such as the asbestos ban, by investing scarce resources (e.g., staff, money, time) into data collection, analysis, consultation with outside parties, and similar activities. In turn, the agency's decisions regarding how much effort to devote to such investigative activities may depend on the institutional structures and incentives created by Congress, courts, and other overseers. Agency expertise, on this view, may be endogenous.

This paper contributes to an emerging literature on the implications of endogenous agency expertise for analyses of bureaucratic politics and public law. In particular, the paper develops a stylized formal model to address two related questions. First, how do changes in the costs associated with adopting a new regulation affect an agency's probability of learning more accurate information about the likely effects of that regulation? That is, how do changes in *enactment costs* affect agency expertise? Second, how would an overseer with the power to manipulate the agency's enactment costs (e.g., a court, legislature, or executive oversight agency) exercise this power when agency expertise is endogenous? In other words, what are the overseer's *optimal enactment costs*?

On the first question, the model shows that the effect of enactment costs on agency expertise depends crucially on what the agency would do if its efforts to acquire additional information are unsuccessful. If an uninformed agency would maintain the status quo, then an increase in enactment costs will *decrease* agency expertise. If an uninformed agency would regulate, then an increase in enactment costs will *increase* agency expertise. This follows from the fact that an agency's incentive to acquire information is maximized when the uninformed agency is indifferent between regulating and maintaining the status quo. So, a change in enactment costs that moves the uninformed agency toward this indifference point will increase agency expertise, but a change that moves the uninformed agency further away from indifference will decrease agency expertise.

On the second question, the model demonstrates that the overseer's optimal enactment costs are influenced by two potentially competing goals. First, and consistent with the existing literature on how an overseer can control an agency by manipulating the agency's decision costs, the overseer in this model would prefer to use enactment costs to align the agency's policy preferences more closely with the overseer's. However, the model shows that enactment costs can also affect the overseer's utility indirectly by influencing the agency's expertise. As a result, the overseer's optimal enactment costs must be sensitive to both these concerns. In some circumstances this can lead to counterintuitive predictions. For example,

even an overseer that is more sympathetic to regulation than the agency may prefer to impose enactment costs if this has a sufficiently positive effect on agency expertise. Likewise, even an overseer that is more skeptical of regulation than the agency may, under some circumstances, prefer an "enactment subsidy" (i.e., negative enactment costs) if this induces a sufficiently large increase in the agency's expertise, and, consequently, fewer erroneous decisions from the overseers perspective.

These two results have implications for an array of ongoing debates in bureaucratic politics, including the role and function of judicial review, the impact of regulatory review conducted by the Office of Management and Budget, and the legislature's use of so-called structure-and-process devices to control the bureaucracy. The model also has implications for other issues in public law, including the appropriate degree of congressional or judicial oversight in the context of both national security matters and ordinary criminal investigations, as well as judicial review of legislative enactments under various constitutional provisions.

Agency Expertise in Theories of Bureaucratic Politics

Most contemporary analyses of bureaucratic policymaking assume a principal-agent problem in which a less-informed principal, usually a legislature, delegates some degree of policy discretion to a (potentially) better-informed bureaucratic agent, but tries to structure the delegation and the institutional environment in order to minimize "bureaucratic drift"—the degree to which the agency pursues goals that diverge from those of the principal (McCubbins, Noll and Weingast 1989; Horn and Shepsle 1989; Shepsle 1992). The assumption that the agency has greater expertise (that is, a higher probability of having superior information about the actual effects of various policy choices[1]) is central to these analyses, both because the agency's greater expertise is often used to explain the initial delegation of authority, and

[1]I use the term "expertise" to refer to the probability of acquiring additional relevant information. "Expertise" might be used in at least two other senses, however. First, it might connote actually having additional information. Second, it might indicate that an actor can improve its probability of learning additional information at low cost. All of these characteristics are parameterized in the model. My use of the word "expertise" to describe only the first is an arbitrary expositional choice.

because the informational asymmetry is what makes the oversight problem so interesting and challenging.

Despite the breadth and sophistication of the literature on this topic, most of this literature has typically assumed exogenous agency expertise. There are some important exceptions, however. In one of the first papers to address the endogenous expertise issue explicitly, Bawn (1995) analyzed the trade-off between bureaucratic expertise and political control by assuming that an agency's incentive to acquire expertise is positively correlated with the scope of its discretion, though in her model this correlation is assumed rather than derived. More recently, Bendor and Meirowitz (2004) have demonstrated that, when information is costly to the agency, the legislature may prefer to delegate to a bureaucrat with policy preferences that are relatively far from the legislature's own, because only such a bureaucrat would be willing to invest in information. In a similar vein, Gailmard and Patty (2006) have developed a career incentives model in which a bureaucrat's investment in job-specific expertise is shown to be positively related with the scope of her autonomy to set policy. Another interesting contribution to this literature is due to Szalay (2005), who shows that under certain conditions a principal's best strategy may be to eliminate the agent's authority to adopt "intermediate" options, because forcing the agent to take a relatively extreme position increases its incentive to invest in information.

The analysis developed in this paper differs from these and other papers in that this paper does not focus on the relationship between the agency's expertise and the scope of its discretion. While expanding, contracting, or otherwise limiting agency discretion is one important tool for influencing bureaucratic policymaking, it is hardly the only one, or even the most effective one. Of particular relevance here, an overseer can also influence agency policy by making certain choices more or less costly relative to others. The fact that a legislature can use this "decision cost" strategy instead of or in addition to a discretion-limiting strategy is one of the important insights in the classic contributions of McCubbins, Noll and Weingast (1987, 1989). The decision-cost approach has been further developed

in important work by, among others, Spiller and Tiller (1997), Tiller (1998), and Spence (1999). More recently, Gailmard (2005) has explored the differences between the decision-cost and discretion-limiting approaches to controlling the bureaucracy and shown that the discretion-limiting approach is only preferable under a limited set of special conditions. It is therefore important to inquire not only how agency expertise might vary with the scope of the agency's delegated discretion, but also how it might vary with the relative decision costs of different actions. This paper is intended as a step in that direction.

Analysis

The Model

Consider a simple sequential policy-making game with two players, a decision-maker and an overseer. The decision-maker might be thought of as an administrative agency, executive official, or bureaucratic subordinate. The overseer might be thought of as a court, legislature, bureaucratic superior, or independent oversight agency. The decision-maker, which I will refer to as the agency, has been charged with making some binary decision, such as whether or not to ban the manufacture of asbestos products, or whether or not to authorize commercial development of a wilderness area. This decision is denoted by $\rho \in \{0, 1\}$, where $\rho = 0$ represents the decision to retain the status quo and $\rho = 1$ represents the decision to take the proposed action.

The proposed action has some net impact, b, which is a random variable drawn from a continuous distribution F with support on $\mathbb{R}$. The density of the distribution is f and the mean is μ. The parameter b may be thought of as a reduced form expression of all the decision-relevant empirical effects associated with the proposed action. In the asbestos example, the b parameter might capture the annual number of cancer cases and other adverse health effects that would be prevented, the economic burden on affected industries and consumers, and so forth. If lives and dollars were the only relevant considerations, b might

be interpreted as the some function that is increasing in the number of lives saved and decreasing in the amount of economic cost.

The preferences of the agency and the overseer are correlated in that, for both of them, the payoff of regulation is increasing in b. Though this assumption may not always hold, it is often sensible. For example, in the asbestos case, if b is the ratio of lives saved per dollar spent, it is reasonable to suppose that extreme liberals, extreme conservatives, and everyone in between would agree that high b values are better than low b values (cf. Stephenson 2006). The agency and the overseer may nonetheless have substantially different views about when a proposed regulation is cost-justified. To capture this preference divergence formally, assume that the utility payoff to the agency from the enactment of new regulation is b, while the utility payoff to the overseer from the new regulation is $b - s$, where $s \in \mathbb{R}$ measures the degree to which the overseer is more skeptical of, or hostile to, the proposed regulation than is the agency.[2] If $s > 0$, the overseer is more skeptical of regulation than the agency, while if $s < 0$, the overseer is more pro-regulation than the agency. In the asbestos example, we might say that the s parameter measures how much more conservative the overseer is relative to the agency, with higher s values indicating greater conservatism. In a different example, though, the ideological connotations of s might differ: If the decision is whether to open a wilderness area to commercial development, greater skepticism toward altering the status quo (a higher s) looks more liberal, while sympathy for the proposed change (a lower s) looks more conservative.

Initially, both the agency and the overseer know the distribution F, but neither knows the true realization of b. The agency, but not the overseer, can attempt to learn b by investing in costly research. Specifically, before the agency chooses ρ, it chooses a level of expertise $\pi \in [0, 1]$ and pays research cost $c(\pi)$, where $c(0) = 0$, $c(1) = \infty$, $c' > 0$, and $c'' > 0$. The overseer observes neither π nor c. After choosing π, the agency observes private signal σ,[3]

[2]The model assumes, for simplicity, that s is constant and common knowledge, though both these assumptions could be relaxed in future research.

[3]In this model, the only effect of the research investment c is to increase the probability that the agency learns more accurate information about the true value of b. Agency research may have other effects, however.

where:

$$\sigma = \begin{cases} b & \text{with probability } \pi; \\ \emptyset & \text{with probability } 1 - \pi. \end{cases}$$

The overseer has the power to make the agency's decision to adopt a new regulation more or less costly relative to a decision to retain the status quo. For example, the overseer might mandate that, before the agency adopts a new regulation, it must comply with onerous procedures or build an elaborate record defending its decision. Alternatively, the overseer might make the decision to initiate new regulation less costly relative to the status quo, perhaps by threatening political retaliation for inaction or by imposing a statutory presumption that action is necessary (e.g., a "hammer" provision) and requiring the agency to comply with burdensome requirements in order to justify inaction. The model captures this power by allowing the overseer at the beginning of the game to select an enactment cost $k \in \mathbb{R}$, which the agency incurs if it decides to adopt a new regulation rather than to retain the status quo.[4] The agency observes k before deciding how much expertise to acquire.[5]

To summarize, the order of play is as follows:

- Step 0: Nature chooses the benefit of regulation b from distribution F;

For example, investment in research may lead the agency to discover alternative regulatory approaches that could achieve higher benefits at lower cost. In other words, research might increase b. One simple way to model this is to assume that the payoff of regulation is higher when the agency is informed than when it is ignorant, i.e. that $(b|\sigma = 1) \geq (b|\sigma = 0)$. That formulation is entirely consistent with the model presented in this paper: One need only redefine F as the distribution of b conditional on $\sigma = 1$ and $\mu = E(b|\sigma = 0)$. It would, of course, also possible to model other relationships between agency research spending and the regulatory payoffs (e.g. Bueno de Mesquita and Stephenson 2006), but I do not pursue those possibilities here.

[4]Note that this framework allows the overseer to make the policy decision itself, rather than delegating this decision to the agency, by selecting $k = \infty$ or $-\infty$.

[5]It is important to highlight two characteristics of enactment costs in this model. First, in contrast to related models of bureaucratic oversight (cf. Gailmard 2005), in this model enactment costs or subsidies do not affect the overseer's utility directly. So, it would be inapt to think of enactment costs in this model as transfers. Rather, they are better though of as levers the overseer can manipulate to make the agency's life easier or harder under different conditions. The imposition of procedural or explanatory requirements would probably be consistent with this assumption, but a change in the agency's budget probably would not be. Second, the model assumes that the overseer can credibly commit to k at the beginning of the game. This assumption may be substantively plausible in some circumstances. It also establishes a baseline case against which other cases, involving imperfect or no credible commitment, might be compared.

- Step 1: The overseer chooses enactment cost k;
- Step 2: The agency chooses a level of expertise π;
- Step 3: After observing signal σ, the agency chooses action ρ, and both players receive their final utility payoffs.

The final utility payoffs to the agency and the overseer are, respectively:

$$U_A = \rho(b - k) - c(\pi); \text{ and}$$

$$U_O = \rho(b - s).$$

Results

The Effect of Enactment Costs on Agency Expertise

The first question to address is how marginal changes in enactment cost k affect the agency's equilibrium level of expertise, π^*. The answer to this question follows straightforwardly from the following proposition:

Proposition 1 *The agency's preferred level of expertise, π^*, is decreasing in $|\mu - k|$, the difference between the enactment cost and the proposed regulation's* ex ante *expected benefit.*[6]

The intuition behind this result is straightforward and grounded in well-known principles of statistical decision theory (Raiffa 1997; Raiffa and Schlaifer 1961). Additional information is valuable to the agency only if it causes the agency to do something different from what it would have done had it remained uninformed. Information is therefore most valuable when the agency is most uncertain *ex ante* as to its best course of action (i.e., when $\mu - k = 0$). If the agency starts out thinking that the benefits of regulation are likely very high relative to the enactment cost ($\mu >> k$), then the agency's investment in research will only improve its

[6] All proofs are in the Appendix.

payoff if the agency discovers that the benefit of the regulation is actually much lower than expected. But the agency considers this possibility unlikely *ex ante*. Similarly, if an agency starts out believing the benefits of regulation, net of enactment costs, are very negative ($\mu << k$), then investing in research helps the agency only in the unlikely event that the true payoff of regulation turns out to be much higher than expected. When the expected net benefit of regulation is close to zero, however, the potential gains from additional information are large: In this case there is a substantial probability that new information will reveal to the agency that its initial hunch about the best course of action turned out to be wrong.

It therefore follows that the effect of marginal changes in enactment costs on agency expertise depends crucially on what the agency would do if it remained ignorant, i.e. observed $\sigma = \emptyset$. When an ignorant agency would regulate (that is, when $\mu > k$), marginal increases in the enactment cost k will increase the agency's expertise, π^*, because such increases reduce the distance between μ and k. On the other hand, when an ignorant agency would choose not to regulate (that is, when $\mu < k$), increases in the enactment cost will decrease the agency's expertise, because in this case increasing k increases the distance between μ and k. Agency expertise is maximized when $k = \mu$.[7]

The Overseer's Optimal Enactment Costs

The next question concerns the optimal enactment cost from the overseer's perspective. The enactment cost affects the overseer's utility in two ways. First, the enactment cost may improve the overseer's utility by bringing the agency's policy preferences into closer alignment with the overseer's. In this way, the overseer can get the agency to make choices that more closely track the choices the overseer itself would have made if it had the same information as the agency. This use of enactment costs is consistent with the existing literature on the manipulation of decision costs as a technique of political control (Spiller and Tiller 1997; Tiller 1998). Therefore, if preference alignment were the overseer's only concern, its optimal

[7] At this point, the indifferent agency could choose $\rho = 1$ with any probability between 0 and 1, the choice of which would be arbitrary and would not affect the expected payoffs of the overseer or the agency.

k, denoted k^*, would be equal to s.

However, Proposition 1 demonstrates that the enactment cost can have a second effect on the overseer's utility. Changes in enactment costs can increase or decrease the agency's expertise, and the overseer benefits from higher levels of agency expertise because greater expertise reduces the number of cases in which an uninformed agency makes a decision that the agency and the overseer would both consider an error. Furthermore, the overseer does not bear the costs associated with increasing agency expertise.[8] Hence, all else equal, the overseer would like the agency's expertise to be as high as possible. Recall from Proposition 1 that agency expertise is maximized when $k = \mu$.

The problem for the overseer is that, except in the special case where $s = \mu$, the overseer's interest in eliminating agency bias and its interest in maximizing agency expertise will conflict. The overseer's optimal choice of k^* will reflect these competing interests, as demonstrated in the following proposition:

Proposition 2 *The overseer's preferred enactment cost, k^*, lies between s (the degree to which the overseer is more skeptical of regulation than the agency) and μ (the expected benefit of regulation to the ignorant agency) That is:*

- $s = \mu \Rightarrow k^* = s = \mu$
- $s > \mu \Rightarrow s > k^* \geq \mu$
- $s < \mu \Rightarrow s < k^* \leq \mu$

Recall that, if agency expertise were exogenous, as most of the extant literature typically assumes, then the overseer's optimal k^* would be equal to s. Qualitatively, this means that

[8]This assumption might be problematic if the costs the agency bears for increasing its expertise are opportunity costs associated with forgoing activities that might also benefit the overseer. While the marginal costs of increased expertise would still most likely be higher for the agency than the overseer, the latter cost would not be zero. The model could be extended to include this possibility by adding an additional term in the overseer's utility function. For this paper, however, I adopt the simplifying assumption that the costs to the overseer of additional research by the agency are sufficiently small that they can be treated as nonexistent.

if the overseer is more skeptical of regulation than the agency, the overseer would prefer a positive enactment cost, while if the overseer is more zealous than the agency (that is, more sympathetic to the proposed regulation), the overseer will prefer a negative enactment cost (a status quo cost or enactment subsidy). Furthermore, the size of this enactment cost or subsidy should correspond as closely as possible to the size of the ideological distance between the agency and the overseer. If the overseer and the agency have the same policy preferences, though, then the overseer would prefer not to impose any enactment cost or subsidy.[9]

Proposition 2 demonstrates how these results change if agency expertise is endogenous. First, the overseer will generally prefer a nonzero enactment cost even when the overseer and the agency have identical policy preferences ($s = 0$). Specifically, Proposition 2 predicts that, if the agency and the overseer are both zealous (that is, if $\mu > s = 0$), the overseer will prefer a positive enactment cost, while if they are both skeptical ($\mu < s = 0$), the overseer will prefer an enactment subsidy. The reason for this is that the overseer and the agency disagree over how much the agency should invest in information. This disagreement arises because the agency incurs the costs associated with higher levels of expertise, while the overseer does not.

What about circumstances in which the agency and the overseer have divergent policy preferences ($s \neq 0$)? Here, there are several cases to consider. Suppose first that the agency and the overseer are both zealous, but the agency is more zealous than the overseer ($\mu > s > 0$). In this case, the overseer will prefer a positive enactment cost, as the traditional theory would predict. But, the magnitude of the optimal enactment cost will be *greater* than s. In this case, the endogeneity of expertise leads the overseer to prefer more substantial enactment costs than would be optimal if agency expertise were exogenous. Similarly, if both the agency and the overseer are skeptical, but the agency is more skeptical than the

[9]These claims are related to the hypothesis known as the "ally principle," which posits that a principal will confer more discretion on an agent with preferences similar to the principal's own (Bendor and Meirowitz 2004).

overseer ($\mu < s < 0$), then the overseer will prefer an enactment subsidy that is larger (a k^* that is more negative) than what the traditional decision cost theory would predict.

The next case to consider is one in which the agency is zealous, but the overseer is skeptical ($s > \mu > 0$). In this case, the overseer will prefer a positive enactment cost, but the magnitude of this cost will be *smaller* than what would be needed to align the agency's policy preferences with those of the overseer. In other words, the overseer would prefer an enactment cost that appears insufficiently large if the endogeneity of agency expertise is ignored. Likewise, in the case where the agency is skeptical and the overseer is zealous ($s < \mu < 0$), the overseer prefers an enactment subsidy, but one that is too small to bring the policy preferences of the agency and overseer into alignment.

Finally, suppose that the agency and the overseer are both zealous, but the overseer is more zealous than the agency ($\mu > 0 > s$). If agency expertise were exogenous, the overseer would prefer an enactment subsidy (in particular, a subsidy $k^* = s < 0$). But if agency expertise is endogenous, we can no longer be certain even of the sign on k^*. It is possible that in this case the overseer would prefer an enactment cost rather than an enactment subsidy, even though the overseer is more zealous than the agency. A similar logic applies to the case where the agency and the overseer are both skeptical, but the overseer is more skeptical than the agency ($s > 0 > \mu$). If agency expertise were exogenous, the overseer would prefer an enactment cost equal to $s > 0$, but when expertise is endogenous the preferred enactment cost will be smaller than s, and may even be negative.

Taken together, these results demonstrate that the predictions regarding overseer preferences, and the influence of enactment costs on agency behavior and regulatory outcomes, may be quite different when agency expertise is endogenous than when it is exogenous. This central insight, and the model's more specific predictions, may have implications for ongoing debates about regulatory oversight and related issues in institutional design and public law.

Implications

Administrative Law and Procedure

Consider the implications of the foregoing analysis for three of the most widely discussed and controversial mechanisms of bureaucratic oversight: "hard look" judicial review of agency decisions, regulatory review by the Office of Management and Budget (OMB), and legislative use of "structure-and-process" control mechanisms.

"Hard Look" Judicial Review

Under §706 of the Administrative Procedure Act (APA), federal courts are empowered to "hold unlawful and set aside agency action, findings, and conclusions found to be arbitrary, capricious, [or] an abuse of discretion." Though this statutory language might have been interpreted as conferring on courts only a limited mandate to strike down plainly irrational agency decisions, courts have interpreted the "arbitrary and capricious" standard to require that an agency demonstrate that it has "examine[d] the relevant data and articulate[d] a satisfactory explanation for its action, including a rational connection between the facts found and the choice made." (State Farm v. Motor Vehicles Manufacturers' Association [463 U.S. 29 (1983)]). This approach is typically referred to as "hard look" judicial review.[10]

Scholars dispute whether hard look review is effective in providing courts with useful information or filtering out unreasonable agency decisions (McGarity 1992; Seidenfeld 1997). One clear effect of hard look review, though, is to make certain actions—usually decisions to alter the status quo[11]—more costly. Critics charge that this leads to the "ossification" of agency rulemaking, deterring socially desirable regulation (McGarity 1992; Pierce 1995). Defenders of hard look review typically argue that the ossification problem is overstated and

[10]Originally the term "hard look" referred to the agency's obligation to take a hard look at relevant considerations before making its decision. Over time, though, the term has come to connote the court's obligation to take a hard look at the agency's reasoning and explanation.

[11]This is not necessarily the same as the decision to impose new regulatory requirements. For example, in *State Farm*, the leading Supreme Court case on hard look review, the agency action under review was a decision to eliminate a regulatory requirement that the agency had adopted in an earlier proceeding.

outweighed by the benefits of hard look review (Jordan 2000; Seidenfeld 1997). Others have suggested an upside to ossification: The costs associated with hard look review may ensure that agencies only pursue policies with sufficiently large benefits (Stephenson 2006).

In the language of the model, hard look review imposes enactment costs on decisions to change the status quo. Courts can manipulate these costs by subjecting certain kinds of agency decisions to greater or lesser scrutiny under the hard look standard. Though in theory the standard is supposed to apply evenhandedly, in practice many observers would agree that "whether the court will dig deeply or bow cursorily depends . . . on whether the judge agrees with the result of the administrative decision" (Rodgers 1981). Social scientists have formalized this intuition and shown how courts can induce agencies to pursue policies that more closely track the courts' regulatory preferences by manipulating agency decision costs in this way (Tiller 1998; Stephenson 2006). Empirical evidence on judicial decision-making, though hardly conclusive, generally supports the view that judges practice "selective deference" in applying the hard look standard (Revesz 1997; Tiller 1998).

How might our understanding of hard look review change if we consider the possibility that agency expertise might be endogenous? First, the prediction that courts will use the stringency of hard look review to bring agency policy preferences into line with those of the judiciary must be qualified along the lines indicated by Proposition 2. Though agency expertise is not valuable to courts in and of itself, higher levels of agency expertise increase the judiciary's utility indirectly if agency and judicial preferences are positively correlated. The court's interest in increasing the agency's expertise will cause the court's optimal level of stringency to diverge from what one would predict if the court's only concern were influencing the agency's policy preferences.

Second, the model suggests a mechanism by which hard look judicial review might affect agency expertise that is different from the mechanisms generally discussed in the existing literature. Proponents of aggressive hard look review have argued that it increases agency expertise because agencies must demonstrate such expertise in order to survive judicial scrutiny

(Seidenfeld 1997; Sunstein 1984). Critics, on the other hand, have asserted that hard look review actually reduces agency expertise because it causes agencies to divert resources away from activities that enhance expertise, such as technical research, and toward those that do not, such as lawyer-dominated *post hoc* record building (Pierce 1995; Shapiro 1988). The model developed here does not engage these competing claims directly, because it does not incorporate the possibility that enactment costs may be a function of expertise or information.[12] Nonetheless, because the model demonstrates that enactment costs *qua* costs can affect agency expertise, it has implications for this debate.

Suppose, for example, that neither the critics nor the defenders of hard look review are correct in their arguments as to how hard look review might affect agency expertise. That is, suppose that the amount the agency invests in expertise neither increases nor decreases the cost to the agency of producing a record sufficient to survive judicial review. Does this mean that the stringency of such review will have no effect on agency expertise? The model indicates that the answer is no. If the ignorant agency would be inclined to regulate, then more stringent hard look review will increase agency expertise even if courts are not able to distinguish between informed and uninformed agencies. On the other hand, if the ignorant agency would retain the status quo, then more stringent hard look review will decrease agency expertise, even if devoting resources to defending regulation in court does not itself erode the agency's technical capabilities.

We can push the point further. Suppose that the hard look defenders are correct that it is easier for an agency to survive judicial review if the agency really knows what it is talking about. Even so, hard look review may still entail enactment costs that are independent of the agency's expertise. If this is correct, then in the case where the ignorant

[12]This assumption is made not for substantive reasons, but rather to simplify the exposition. A useful extension of the model might incorporate the possibility that enactment costs are affected by investment in expertise (i.e., that k is a function of π) or that enactment costs are a function of the agency's private information (i.e., that k is a function of σ). Intuitively, it appears likely that, compared to the baseline model developed in this paper, the agency's expertise would be higher if enactment costs are negatively correlated with the agency's expertise or the accuracy of its private information, and lower if if enactment costs are positively correlated with these things. However, I defer full consideration of this possibility to future research.

agency would regulate, hard look review has two countervailing effects on agency expertise. On one hand, investment in research will make it easier to survive judicial review, which would increase the agency's incentive to acquire expertise. On the other hand, though, the enactment costs associated with hard look review will decrease the agency's incentive to acquire expertise. Without more specific information, it is impossible to say which of these effects will predominate.

Alternatively, suppose the critics are correct that the activities an agency must engage in to survive hard look review not only do not increase agency expertise, but compete for resources with activities that do. Even under this assumption, it is not necessarily the case that more stringent hard look review reduces agency expertise. If an ignorant agency would be wiling to regulate, and if hard look review increases enactment costs, then the model shows that more stringent review can increase the agency's incentive to acquire expertise. This effect may or may not be stronger than the decrease in the agency's expertise created by the shift of resources from technical research to lawyer-dominated record building.

OMB Review

Courts are not the only oversight body that may subject agency regulations to a "hard look." Under executive orders promulgated by Presidents Reagan and Clinton, agencies must submit proposals for major federal regulation to OMB's Office of Information and Regulatory Affairs (OIRA) for review and consultation. Though OIRA cannot formally veto a regulatory proposal, OIRA review can significantly delay or derail regulatory initiatives. Defenders of OIRA oversight argue that it improves the quality of regulation, correcting for agency "tunnel vision," overzealousness, or other failings (DeMuth and Ginsburg 1986; Seidenfeld 2001). Critics have argued that OIRA's ability to impose delays and require additional justification for regulatory proposals creates a regulatory "black hole" that swallows up many worthwhile initiatives, often for political reasons but sometimes as an unintended effect of the overly cumbersome review process (McGarity 1992; Morrison 1986). While a great deal

has been and could be said about the pros and cons of OIRA review, for purposes of the present analysis the salient question is how it might affect an agency's incentives to invest in expertise. Like judicial review, OIRA review may sometimes be effective in identifying cases where the policymaking agency did or did not have relevant information, but it also imposes enactment costs. It is therefore sensible to ask how the existence of these costs might affect the agency's expertise, independent of the substantive efficacy of OIRA review.

Because the analogy between judicial review and OIRA review is so close, it is unnecessary to belabor the model's predictions for the effects of changes in the enactment costs associated with OIRA review. It is worth briefly restating and highlighting two of these results, however. First, if an uninformed agency would be inclined to regulate, then the costs and delays associated with the OIRA process can improve agency expertise even if the review process does nothing directly to improve the quality of the regulation. As a result, OIRA may prefer to impose such costs and delays even if it has exactly the same policy preferences as the agency. Indeed, OIRA may sometimes want to make enactment of regulation costly even if OIRA is *more* enthusiastic about the proposed regulation than is the agency. Second, if an uninformed agency would not regulate, then OIRA review costs will reduce the agency's expertise. So, even if the defenders of OIRA review are correct that it improves agency expertise by forcing the agency to produce more high-quality studies of a proposed regulation's effects, the costs that OIRA review imposes can sometimes generate a countervailing effect. Without more information, one cannot determine which of these effects will predominate, and hence one cannot determine whether more stringent OIRA review will meet its goal of increasing agency expertise.

"Structure-and-Process" Control of Agency Decision-Making

The model developed here does not apply directly to many of the most widely discussed forms of legislative oversight of the bureaucracy, such as budgetary control (Carpenter 1996; Ting 2001) or the threat of statutory override (Cameron and Rosendorff 1993; Ferejohn

and Shipan 1990). That said, the model may apply to a subset of so-called structure-and-process techniques for extending legislative influence over bureaucratic policymaking (McCubbins, Noll and Weingast 1987, 1989). While the structure-and-process category includes a wide variety of control mechanisms, certain of these mechanisms operate primarily by manipulating an agency's decision costs, making some courses of action relatively more or less costly by altering procedural requirements. Empirical research has suggested that these forms of structure-and-process control may be among the most effective and important. Spence (1999), for example, finds in the context of federal hydroelectric licensing decisions that "those procedures that were specifically tailored to increase the transaction costs of a particular decision ... were more effective than more general, facially neutral procedures." The model directly applies to this form of legislative control over bureaucratic policy.

As noted earlier, most of the existing literature on these mechanisms assumes that Congress would prefer to impose costs that eliminate agency "bias" or "drift," aligning agency policy preferences with those of the legislature (Bendor, Taylor and Van Gaalen 1987; McCubbins, Noll and Weingast 1987, 1989; Spiller and Tiller 1997). To avoid repetition, I will not restate the model's main conclusions in this context, except to point out that if the model captures something important about the impact of these mechanisms on Congress's utility, then congressional incentives regarding the appropriate design of these structure-and-process mechanisms may differ from the predictions of the conventional theory. If that is true, then the model both suggests new avenues for theoretical and empirical research on structure-and-process control, and also suggests that many of the existing empirical tests of the structure and process hypothesis may be misspecified.

Other Public Law Applications

Though the analysis so far has focused on delegation and oversight in the context of regulatory policymaking, the model's basic insight is relevant in other contexts in which a decision-maker can acquire expertise by engaging in costly effort, while a relatively unin-

formed overseer has the power to influence the decision-maker's choices by making certain decisions relatively more or less costly. In this section, I discuss how the model might apply in three other public law contexts: congressional and judicial oversight of executive decision-making on national security issues; magistrate screening of police search warrant applications; and judicial review of legislative enactments.

National Security

Debates involving the tension between the importance of executive branch expertise and the perceived need for judicial and legislative oversight are not limited to regulatory policy. More recently, these debates have assumed particular salience in the context of national security. To what extent should executive decisions regarding national security—whether to undertake military action, employ coercive interrogation techniques, authorize wiretaps, and the like—be subject to procedural safeguards or other forms of congressional or judicial oversight? Many have argued that oversight of executive decisions is essential to preserving meaningful checks and balances and to preventing abuses of power (Cole 2003, 2004; Stone 2004). Others have countered that because the executive which has greater expertise and access to relevant information, and must often act quickly and decisively in times of war or emergency, burdensome procedures or intrusive oversight can endanger national security (Posner and Vermeule 2003, 2005; Sunstein 2005). This set of questions obviously involves myriad political, legal, and moral concerns well beyond the scope of this paper. That said, the analysis presented here may shed light on one important aspect of this problem that has received comparatively little attention: the effect of burdensome oversight mechanisms on the executive's expertise regarding the national security implications of different courses of action.

As a stylized illustrative example, consider the question whether the executive may authorize otherwise impermissible coercive interrogation techniques to be used against suspected terrorists or terrorist sympathizers. Before the responsible official decides whether to au-

thorize such tactics, she may attempt to acquire additional information about the extent of a particular suspect's likely knowledge of terrorist activities. This costly pre-interrogation investigation may or may not uncover additional useful information, however. How does the incentive to pursue pre-interrogation investigation change as the decision to use coercive techniques becomes more costly for the responsible executive official? This question is important because these decision costs may increase if external actors, such as Congress or the courts, impose more procedural or substantive requirements on the decision to employ coercion.

The model demonstrates that the impact of burdensome oversight on executive expertise depends on what the executive would do if attempts to acquire more pre-interrogation information are unsuccessful. If the responsible official would not use coercive interrogation if she fails to uncover additional information indicating that the suspect has information critical to national security, then procedural requirements that make coercive interrogation more costly to the executive will decrease pre-interrogation investigation. As a result, a rational overseer might prefer to refrain from imposing burdensome procedures, and might even prefer to make coercive interrogation a relatively more desirable option. This may be the case even if the overseer is more strongly opposed to coercive interrogation than is the executive.

On the other hand, though, consider the alternative case in which the responsible official would go ahead with a coercive interrogation absent additional concrete information about the suspect's likely knowledge of terrorist threats. In this case, increasing the procedural burdens associated with the authorization of coercive techniques will increase the official's *ex ante* investment in pre-interrogation investigation. This prediction implies that a legislative or judicial overseer might prefer to impose such burdens, even if the overseer and the executive official have exactly the same views on the circumstances under which coercive interrogation is justified. In other words, the case for some degree of burdensome procedural oversight does not depend on the belief that the executive has the wrong policy preferences, e.g. that it

cares too little about civil liberties. In some circumstances, as the example above illustrates, burdensome oversight may be appropriate because the executive might otherwise invest too little in expertise from a social point of view, even though the executive has the right policy preferences.

Criminal Investigations

The same basic argument applies to more garden-variety forms of criminal law enforcement as well. Consider, as another example consistent with the model, police applications for search warrants. Police officers are likely better-informed, relative to the magistrate judges who review warrant applications, as to whether a search is justified by its expected benefits. But, magistrates may lack the expertise to evaluate the information contained in warrant applications, and in practice warrant applications are rarely denied. Even if substantive review of warrant applications is minimal, however, the application process itself can affect police officers' behavior by making the acquisition of a warrant more costly to police, and this cost can screen out searches the police view as low value *ex ante* (Dripps 1986; Stephenson 2006; Stuntz 2002). If police officers' expertise—their probability of learning additional information about the likely benefits of a given search—is exogenous, then the rational thing for magistrate judges or legislatures to do is to set application costs such that the preferences of the police are aligned with those of the relevant overseer. For example, if the legislature believes the police are too eager to search in marginal cases, the optimal solution would be to impose an application cost high enough to eliminate this pro-search bias.

But officers usually can learn more about the likely value of a particular search only by investing scarce resources in pre-search investigation. If this is so, then the resources the police will devote to such investigations will be influenced by the costs associated with applying for the warrant. If the police, absent additional information, would prefer to search, then increasing warrant application costs will induce police to do more pre-application investigation. On the other hand, when the police would forgo a given search unless they learn

more definitive information, increasing warrant application costs will reduce pre-application investigation. Even if the legislature thinks that the police are applying exactly the right standard to determine when they should search, the legislature may still want to make warrant applications costly if the police are often inclined to search even when their pre-search investigation turns up little information. Conversely, the fact that the legislature is generally more skeptical of the police about the value of searches does not necessarily justify the imposition of significant warrant application costs.

Judicial Review of Legislative Enactments

The focus so far has been on decisions made by the executive branch, as this is the public policy context in which the asymmetric information problem is most frequently discussed. That said, there is a potentially analogous set of arguments with respect to judicial review of legislative decisions, particularly in constitutional cases. Though we often think of such cases as involving only pure questions of law, the inquiry into whether a particular legislative enactment violates a constitutional provision often involves a disputed empirical claim or prediction. For example, deciding whether an exercise of the eminent domain power in the service of economic development satisfies the Public Use Clause may turn on an empirical judgment as to the likely economic benefits of the proposed taking (Kelo v. City of New London [125 S.Ct. 2655 (2005)]). Similarly, deciding whether a restriction on speech serves a legitimate state interest may entail an assessment of the government's claims regarding the likely consequences of permitting the speech in question (Turner Broadcasting v. FCC [512 U.S. 622 (1994)]). Also, the decision whether a given federal law is a legitimate exercise of Congress's power under Section Five of the Fourteenth Amendment turns, as a matter of law, on the "congruence and proportionality" between the Congressional mandate and the constitutional harm it seeks to redress (City of Boerne v. Flores [521 U.S. 507 (1997)]). This inquiry necessarily involves a factual assessment of both the existing state of the world and the likely impact of the challenged statute (cf. Nevada Dep't of Human Resources v. Hibbs

[538 U.S. 721 (2003)]).

The question therefore arises: How deferentially or aggressively courts should scrutinize legislative findings (or predictions) of constitutionally significant facts? The answer varies by doctrinal area, but at least in some cases the Supreme Court has used rhetoric quite similar to what one observes in the administrative context. Although the Supreme Court has emphasized that the legislature has greater expertise on these factual questions and is therefore deserving of deference, in at least a few cases the Court has suggested that it will look closely at the record to make sure that Congress gave adequate consideration to the question and made satisfactory factual findings (Turner Broadcasting v. FCC [512 U.S. 622 (1994)], University of Alabama v. Garrett [531 U.S. 356 (2001)]). This approach, in the view of some observers, suggests an emerging form of "hard look" review for legislative enactments (Bryant and Simone 2001; Frickey and Smith 2002). One effect of such review in this context, as in the administrative context, would be to raise the relative costs to Congress of enacting particular kinds of statutes (Stephenson 2006).

If something like hard look review of legislative enactments does indeed exist, or if reviewing courts have at their disposal other means by which they can alter the relative costs to Congress of different courses of action, for example through clear statement rules or other techniques of statutory interpretation (Eskridge and Frickey 1992), then the analysis developed here may offer some insights into how these mechanisms influence legislative acquisition of relevant factual information. If legislative expertise, like agency expertise, is endogenous, then assessments of both the level of legislative expertise and of optimal enactment costs from the judiciary's perspective must take this fact into account.

Conclusion

This paper has developed a stylized formal model to investigate two questions. First, how does a decision-maker's incentive to acquire expertise changes as the enactment costs associ-

ated with adopting a new policy change? Second, given this effect, what level of enactment costs would an overseer prefer to set?

The model demonstrates that the answer to the first question depends on what the decision-maker would do if it remains uninformed. The incentive to acquire expertise is strongest when the uninformed decision-maker is least sure of its best course of action. Therefore, when an uninformed decision-maker would retain the status quo, increases in enactment costs will decrease expertise, but when an uninformed decision-maker would adopt the new policy, increases in enactment costs will increase expertise.

On the second question, the model shows that the overseer's optimal enactment costs must balance two interests: On one hand, the overseer would like to align the decision-maker's policy preferences with its own, but on the other hand, the overseer would like to increase the agency's incentive to acquire expertise. These interests may compete, which leads to predictions that can be quite different from what one would predict if the decision-maker's expertise were assumed to be exogenous.

These conclusions have applications to an array of important problems in public law and institutional design. Most obviously, they imply that the study of various forms of bureaucratic oversight must take into account the impact these oversight mechanisms may have on agency expertise. The analysis may also apply to debates over executive power in times of national emergency, legislative and judicial oversight of criminal investigations, and constitutional review of legislation.

The model developed here, of course, is stylized and incomplete. The basic framework might be usefully extended to incorporate, for example, the overseer's commitment problem, multiple agencies or overseers, a larger set of possible actions, or the possibility that enactment costs might correlate with expertise or information. Work along these and related lines may enrich our understanding of the complex interplay between the policymaking process and the acquisition of policy-relevant information.

Appendix

Proof of Proposition 1

When the agency must choose ρ at Step 3, research costs c are sunk, so the agency will choose $\rho = 1$ if and only if $E(b|\sigma) - k > 0$. Therefore, the agency will choose $\rho = 1$ if (1) if it observes $\sigma = b > k$, or if (2) if $\sigma = \emptyset$ and $\mu > k$.

Consider first the case where $\mu < k$. In this case:

$$\frac{d|\mu - k|}{dk} > 0 \tag{1}$$

The agency will choose $\rho = 1$ if and only if the agency observes $\sigma = b > k$. Therefore, the agency's expected utility at Step 2 is:

$$\begin{aligned} E(U_A|\mu < k) &= \pi Pr(b > k)[E(b|b > k) - k] - c(\pi) \\ &= \pi \left(\int_k^\infty (b - k) f(b)\, db \right) - c(\pi) \end{aligned} \tag{2}$$

At Step 2, the agency chooses π^* to solve:

$$\frac{d}{d\pi} E(U_A|\mu < k) = \int_k^\infty (b - k) f(b)\, db - c'(\pi) = 0 \tag{3}$$

By the implicit function theorem, the effect of changes in the enactment cost k on the agency's preferred level of expertise π^* is:

$$\frac{d\pi^*}{dk} = -\frac{\frac{d}{dk}(\int_k^\infty (b - k) f(b)\, db - c'(\pi^*))}{\frac{d}{d\pi^*}(\int_k^\infty (b - k) f(b)\, db - c'(\pi^*))} = -\frac{1 - F(k)}{c''(\pi^*)} < 0 \tag{4}$$

From equations 1 and 4 it follows immediately that:

$$\frac{d\pi^*}{d|\mu - k|} < 0 \text{ when } \mu < k \tag{5}$$

Next, consider the case where $\mu > k$. In this case:

$$\frac{d|\mu - k|}{dk} < 0 \tag{6}$$

In this case, the agency will choose $\rho = 1$ if and only if it observes either $\sigma = b > k$ or $\sigma = \emptyset$. Therefore, the agency's expected utility at Step 2 is:

$$\begin{aligned} E(U_A|\mu > k) &= \pi Pr(b > k)[Pr(b > k)E(b|b > k) - k] + (1 - \pi)(\mu - k) - c(\pi) \\ &= \pi \left(\int_k^\infty (b - k) f(b)\, db \right) + (1 - \pi)(\mu - k) - c(\pi) \end{aligned} \tag{7}$$

At Step 2, the agency chooses π^* to solve:

$$\frac{d}{d\pi} E(U_A|\mu > k) = \int_k^\infty (b - k) f(b)\, db - (\mu - k) - c'(\pi^*) \tag{8}$$

By the implicit function theorem, the effect of changes in the enactment cost k on the agency's preferred level of expertise π^* is:

$$\frac{d\pi^*}{dk} = -\frac{\frac{d}{dk}(\int_k^\infty (b - k) f(b) db - (\mu - k) - c'(\pi^*))}{\frac{d}{d\pi^*}(\int_k^\infty (b - k) f(b) db - (\mu - k) - c'(\pi^*))} = \frac{F(k)}{c''(\pi^*)} > 0 \tag{9}$$

From equations 6 and 9 it follows immediately that:

$$\frac{d\pi^*}{d|\mu - k|} < 0 \text{ when } \mu > k \tag{10}$$

Equations 5 and 10 are sufficient to establish the proposition.

Proof of Proposition 2

First, consider the case where $s = \mu$. In this case, $k = s = \mu$ both maximizes the agency's expertise (because $k = \mu$) and minimizes the distance between the agency's policy preferences and the overseer's policy preferences (because $k = s$). Because, holding other factors

constant, the overseer's utility is increasing in agency expertise and decreasing in the distance between k and s, it follows immediately that:

$$s = \mu \Rightarrow k^* = s = \mu \tag{11}$$

Next, consider the case where $s > \mu$. First suppose, consistent with the proposition, that $\mu < k^*$. In this case, the ignorant agency not regulate, so at Step 1 of the game the overseer must have maximized:

$$E(U_O|\mu < k) = \pi^* \left(\int_k^\infty (b - s) f(b)\, db \right) \tag{12}$$

Therefore, this k^* must solve:

$$\frac{d\pi^*}{dk} \left(\int_k^\infty (b - s) f(b)\, db \right) = \pi^* f(k)(k - s) \tag{13}$$

Because $\mu < k^*$, $\frac{d\pi^*}{dk} < 0$ (see equation 4 in the proof of Proposition 1). From this, and the facts that $\int_{k^*}^\infty (b - s) f(b)\, db$, π^*, and $f(k^*)$ are all positive, it follows that $k^* < s$.

Now suppose that, inconsistent with the proposition, $s > \mu$ but $\mu > k^*$. In this case, the ignorant agency would regulate, so in order to select this k^* the overseer at Step 1 of the game must have maximized:

$$E(U_O|\mu > k) = \pi^* \left(\int_k^\infty (b - s) f(b)\, db \right) + (1 - \pi^*)(\mu - s) \tag{14}$$

Therefore, this k^* must solve:

$$\frac{d\pi^*}{dk} \left(\int_k^\infty (b - s) f(b)\, db + (s - \mu) \right) = \pi^* f(k)(k - s) \tag{15}$$

Because $\mu > k^*$, $\frac{d\pi^*}{dk} > 0$ (see equation 9 in the proof of Proposition 1). From this, and the facts that $\int_k^\infty (b - s) f(b)\, db$, π^*, $f(k^*)$, and $(s - \mu)$ are all positive, it follows that $k^* > s$. But this contradicts the assumption that $s > \mu > k^*$. It therefore follows that:

$$s > \mu \Rightarrow s > k^* > \mu \tag{16}$$

Next, consider the case where $s < \mu$. First suppose, consistent with the proposition, that $\mu > k^*$. In this case, the ignorant agency would regulate, so this k^* must satisfy equation 15, which implies that $\mu > k^* > s$, consistent with the proposition. Now suppose, contrary to the proposition, that $\mu < k^*$. Because the ignorant agency would not regulate in this case, this k^* must satisfy equation 13. But that implies that $k^* < s$, which contradicts the assumption that $s < \mu < k^*$. Therefore:

$$\mu > s \Rightarrow \mu > k^* > s \tag{17}$$

Equations 11, 16, and 17 are sufficient to establish the proposition.

■

References

Bawn, Kathleen. 1995. "Political Control versus Expertise: Congressional Choices About Administrative Procedures." *American Political Science Review* 89(1):62–73.

Bendor, Jonathan and Adam Meirowitz. 2004. "Spatial Models of Delegation." *American Political Science Review* 98(2):293–310.

Bendor, Jonathan, Serge Taylor and Roland Van Gaalen. 1987. "Stacking the Deck: Bureaucratic Missions and Policy Design." *American Political Science Review* 81(3):873–896.

Bryant, A. Christopher and Timothy J. Simone. 2001. "Remanding to Congress: The Supreme Court's New 'On the Record' Constitutional Review of Federal Statutes." *Cornell Law Review* 86:328–.

Bueno de Mesquita, Ethan and Matthew C. Stephenson. 2006. "Regulatory Review with Observable and Unobservable Effort." Unpublished Manuscript, Harvard University.

Cameron, Charles M. and B. Peter Rosendorff. 1993. "A Signaling Theory of Congressional Oversight." *Games and Economic Behavior* 5:44–70.

Carpenter, Daniel P. 1996. "Adaptive Signal Processing, Hierarchy, and Budgetary Control in Federal Regulation." *American Political Science Review* 90:283–302.

Cole, David. 2003. "Judging the Next Emergency: Judicial Review and Individual Rights in Times of Crisis." *Michigan Law Review* 101:2565–2595.

Cole, David. 2004. "The Priority of Morality: The Emergency Constitution's Blind Spot." *Yale Law Journal* 113:1753–1800.

DeMuth, Christopher C. and Douglas H. Ginsburg. 1986. "White House Review of Agency Rulemaking." *Harvard Law Review* 99:1075–1088.

Dripps, Donald. 1986. "Living with *Leon*." *Yale Law Journal* 95:906–.

Eskridge, William N. and Philip P. Frickey. 1992. "Quasi-Constitutional Law: Clear Statement Rules as Constitutional Lawmaking." *Vanderbilt Law Review* 45:593–646.

Ferejohn, John and Charles R. Shipan. 1990. "Congressional Influence on Bureaucracy." *Journal of Law, Economics, and Organization* 6:1–43.

Frickey, Philip P. and Steven S. Smith. 2002. "Judicial Review, the Congressional Process, and the Federalism Cases: An Interdisciplinary Critique." *Yale Law Journal* 111:1707–.

Gailmard, Sean. 2005. "Discretion Rather that Rules: Choice of Instruments to Constrain Bureaucratic Policy-Making." Unpublished manuscript, Northwestern University.

Gailmard, Sean and John W. Patty. 2006. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." Unpublished manuscript, Northwestern University.

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

Jordan, William S. 2000. "Ossification Revisited: Does Arbitrary and Capricious Review Significantly Interfere With Agency Ability to Achieve Regulatory Goals Through Informal Rulemaking." *Northwestern Law Review* 94:393–450.

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(2):243–277.

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* 75:431–482.

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

Morrison, Alan B. 1986. "OMB Interference with Agency Rulemaking: The Wrong Way to Write a Regulation." *Harvard Law Review* 99:1059–1074.

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

Posner, Eric A. and Adrian Vermeule. 2003. "Accommodating Emergencies." *Stanford Law Review* 56:605–644.

Posner, Eric A. and Adrian Vermeule. 2005. "Emergencies and Democratic Failure." University of Chicago Public Law Working Paper No. 104.

Raiffa, Howard. 1997. *Decision Analysis*. New York: McGraw-Hill.

Raiffa, Howard and Robert Schlaifer. 1961. *Applied Statistical Decision Theory*. Boston: Division of Research, Graduate School of Business Administration.

Revesz, Richard L. 1997. "Environmental Regulation, Ideology, and the D.C. Circuit." *Virginia Law Review* 83:1717–1772.

Rodgers, William H. 1981. "Judicial Review of Risk Assessments: The Role of Decision Theory in Unscrambling the *Benzene* Decision." *Environmental Law* 11:301–.

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

Seidenfeld, Mark. 2001. "The Psychology of Accountability and Political Review of Agency Rules." *Duke Law Journal* 18:1051–1095.

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(1):111–118.

Spence, David B. 1999. "Managing Delegation Ex Ante: Using Law to Steer Administrative Agencies." *Journal of Legal Studies* 28:413–459.

Spiller, Pablo T. and Emerson H. Tiller. 1997. "Decision Costs and the Strategic Design of Administrative Process and Judicial Review." *Journal of Legal Studies* 26:347–370.

Stephenson, Matthew C. 2006. "A Costly Signaling Theory of 'Hard Look' Judicial Review." Olin Center Discussion Paper.

Stone, Geoffrey R. 2004. *Perilous Times*. W. W. Norton.

Stuntz, William J. 2002. "Local Policing After the Terror." *Yale Law Journal* 111:2137–.

Sunstein, Cass. 1984. "In Defense of the Hard Look: Judicial Activism and Administrative Law." *Harvard Journal of Law and Public Policy* 7:51–59.

Sunstein, Cass. 2005. "Administrative Law Goes to War." *Harvard Law Review* 118:2663–2672.

Szalay, Dezso. 2005. "The Economics of Clear Advice and Extreme Options." *Review of Economic Studies* 72(4):1178–1198.

Tiller, Emerson H. 1998. "Controlling Policy by Controlling Process: Judicial Influence on Regulatory Decision-Making." *Journal of Law, Economics, and Organization* 14(1):114–135.

Ting, Michael M. 2001. "The 'Power of the Purse' and its Implications for Bureaucratic Policy-Making." *Public Choice* 106:243–274.