ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS
FELLOWS' DISCUSSION PAPER SERIES

A REVERSE RATIONALE FOR RELIANCE ON REGULATORS

Laurence Tai

Discussion Paper No. 51

06/2013

Harvard Law School
Cambridge, MA 02138

A Reverse Rationale for Reliance on Regulators

Laurence Tai*

Abstract

The standard justification for agencies is that they can obtain information that political leaders are unable to access. This paper develops a model illustrating the opposite logic, that agencies may prevent leaders from acting on or receiving information that leaders would otherwise obtain and use for decision-making. This dynamic can occur when the key information comes from an outside entity and when the agent, based on his policy preferences, is able to induce this outsider to produce higher quality information than the leader can. The model suggests that agencies may be designed to mitigate regulatory capture focused on elected officials, and that disclosure patterns that might appear to indicate agency capture instead facilitate this design. Also, when the agency is at risk of capture, stronger ethics rules may be more effective than increased transparency in combating capture. The FDA drug approval process is offered as a practical application of the model.

JEL Classes: D72, D73, K23

---

*Ph.D. Candidate in Public Policy, Harvard Kennedy School, and J.D. Candidate, Harvard Law School. The author is a Terence M. Considine Fellow in Law and Economics at Harvard Law School and acknowledges support from the Considine Family Foundation and the School's John M. Olin Center for Law, Economics, and Business. This paper has received valuable feedback from Jim Alt, Dan Carpenter, Oliver Hart, Ken Shepsle, Matthew Stephenson, Yuki Takagi, Richard Zeckhauser, and Galina Zudenkova. Various versions of this paper have been presented at the Political Economy and Contracts and Organizational Economics seminars at Harvard, at the 2013 MPSA Annual Conference, and at the 2013 ALEA Annual Meeting.

A Reverse Rationale for Reliance on Regulators

1 Introduction

Scholarship dating at least from Weber's (1922/1978) and Landis' (1938) seminal works on bureaucracy has identified expertise as a fundamental rationale for agencies in the modern administrative state. A recent review of models in this area observes that "[t]he bureaucratic agent in these models typically possesses (or may come to possess) some information that the leader would like to extract to make a decision" (Gailmard and Patty 2012, 354). That elected officials rely on agencies because the latter have better access to information than the former is a foundational principle of bureaucracy studies that does not appear to have been seriously challenged. Instead, because bureaucrats are unelected, a large body of literature focuses on how to address the risk that an agency might use its informational advantage to pursue policies that differ from what political leaders would prefer if they had the same information as the agency (e.g., Epstein and O'Halloran 1999, Huber and Shipan 2002).

More recent work has problematized bureaucratic agents' ability and willingness to secure the specialized information that justifies their role. First, if information is costly to acquire, a political leader may have to incentivize an agent to gather higher quality information, and this task may conflict with her desire for agency decision-making that conforms to her policy preferences (e.g., Gailmard and Patty 2013a). Second, agencies may need to gather information from outside parties like firms to effectively implement policy and may have difficulties doing so (e.g., McCarty 2013). However, works focusing on these adjustments to the theory of bureaucratic expertise seem implicitly to affirm that agencies serve the purpose of gather information that leaders cannot themselves obtain.

This paper considers a reverse rationale for reliance on agencies, that they may benefit a leader by preventing her from receiving or acting on information that she is perfectly capable of obtaining. One half of this logic, that a leader is capable of gaining information, requires only that she can understand information that others have produced. For example, even if she cannot conduct scientific experiments, she may be able to assess the quality of these experiments and interpret the results, either directly or through trusted staff. Steps to obtain passive expertise in practice have occurred in both Congress and the White House: Congress engages in oversight through the Government Accountability Office and the Congressional Research Service (Beermann 2006, 127–30), which may allow its members to rely on in-house expertise. Meanwhile, during the George W. Bush Administration, the Office of Information and Regulatory Affairs added scientific experts to its personnel (Graham, Noe, and Branch 2006).

The other half of this logic, that an agency would helpfully prevent the leader from receiving information, implies minimally that some outside party is generating the information. Much information for agency policymaking does actually comes from firms. In the area of industrial regulation, firms are likely to have private knowledge about their manufacturing processes and the costs of potential regulations (Coglianese, Zeckhauser, and Parson 2004). Other examples are tests of new chemicals under the Toxic Substances and Control Act and project proposals for government contracts. Thus, there exists a variety of cases in which regulated parties, rather than agencies, act as the primary researchers for decisions involving the executive branch.

Although a typical problem is outside parties' reluctance to provide their information (see, e.g., Coglianese, Zeckhauser, and Parson 2004, McCarty 2013), they may, in other circumstances, be too willing to provide information. This scenario arises when the leader and agent can observe their information. This situation can harm a leader if they produce lower quality information as a result. The area of drug approvals serves as a useful illustration: Congress generally allows Food and Drug Administration (FDA) to investigate pharmaceutical firms' clinical trials, rather than itself examining the results. This mode of operation may be desirable for Congress even if it is perfectly capable of reviewing the data and has no time constraints. The reason is that, if Congress were to directly review a firm's application, the company could create evidence just barely favorable enough to warrant approval. FDA might require stronger evidence, perhaps because, as an unelected bureaucracy, it is less swayed by constituencies interested in new treatments for diseases. Congress might be able to induce the drug manufacturer to produce higher quality information by forcing it to communicate only with FDA or by delegating the final decision to the agency.

Therefore, instead of securing information from outsiders that the leader is unable to access or process, an agent may have the opposite function of preventing leaders from acting on or receiving outsiders' information. Generalizing from the above scenario, this paper presents a model in which the leader can do better if the outside party cannot directly convey information to her and motivate her to act but must instead work through an agent. Her payoff can improve if the agent initially disfavors the third party's preferred policy and requires more accurate information than the leader does to favor that policy. More accurate information, in turn, allows the selection of policy that is more likely to be correct.

This rationale for agencies has several implications for understanding regulatory capture: First, institutional arrangements to have agencies make decisions or filter information may mitigate this phenomenon rather than enhance it. Second, it relates to work on how capture of agencies may (not) be inferred from their behavior (e.g., Carpenter 2004). In particular, seemingly discomforting disclosure patterns, such as (1) release only of information that favors interest groups and (2) release of no information at all, are not necessarily signs of capture but instead may further public interest. Finally, to the extent that interest groups might attempt to influence agencies' policy preferences, the model suggests that mandatory disclosure is less promising than ethics rules for agents in combating capture. The difference in the desirability of different measures against capture is especially important since the Obama Administration has initiated policies relating both to transparency and ethics (Coglianese 2009, Thurber 2011).

The rest of this paper proceeds as follows: Section 2 presents a simple model of agency in which policy may be based on an interested outside party's research. Section 3 compares equilibrium outcomes in which the leader can receive the third party's information and select the policy, those in which the agent selects the policy, and those in which the leader retains decision-making authority but allows the third party to communicate only with the agent. Section 4 discusses the implications for capture from the baseline model and an extension. Section 5 suggests that the FDA drug approval process is a plausible example of the reverse rationale operating in practice. Section 6 concludes.

2 The Model

The game features a political leader (L, she), whose preferences, until Section 4, are identical to those of some "public," represented as a passive principal (P); an agent or agency (A, he), who proposes policy; and an outsider or researcher (R, it), which generates new information about policy. The goal is to structure information channels and decision-making authority to benefit the public.

2.1 Basic Elements

There are two policies, $x \in X = \{0, 1\}$, from which to choose, and the leader may either make the final selection or delegate that decision to the agent. The state of the world, $w \in \{0, 1\}$, is unknown, but the probability of each state w is $q_w \in (0, 1)$. The outsider can conduct research to reduce the initial uncertainty by expending effort, $e \in E = \mathbb{R}_+$. If its effort is zero, it can be said to be doing no research. Lack of research may be understood as literally no effort or as an idealization of doing some minimum amount of work to produce a report, but not any real work that contributes to the players' understanding of the situation. Any effort, even zero, generates a signal, $s \in S = \{0, 1\}$. $\Pr(s = w|e) \equiv g(e)$ is an accuracy function, which is a continuous, increasing, and concave function, with $g(0) = 1/2$ and $\lim_{e\to\infty} g(e) = 1$. The cost of research effort, $c(e)$, is continuous, increasing, and convex, with $c(0) = c'(0) = 0$. Together, the accuracy and cost functions will be referred to as the outsider's research technology, which is the same as that in Prendergast (2007).

For each player $i \in \{L, P, A, R\}$, the benefits when $x \neq w$ are normalized to 0 in each state, so that net benefits to player i in state w when $x = w$ are denoted by b^i_w. For now, the leader has the same preferences as the principal, so whatever benefits the leader benefits the public, and the principal will not mentioned again until Section 4. The player's preferences are restricted as follows: First, the leader and agent strictly prefer $x = w$ in each state, i.e., $b^i_w > 0, i \in \{L, A\}, w \in \{0, 1\}$, which means that they agree on which policy is better in both states of the world. However, because the benefits for matching the state can differ for each state and each player, there remains scope for disagreement about what policy to pursue in the face of uncertainty about the state. The outsider has $b^R_1 > 0$, but b^R_0 may be positive or negative, which means that it may prefer policy 1 in both states. Interest groups might be thought to have preferences of this sort. However, so that the researcher may end up exerting effort in equilibrium, its preferences are bounded so that $q_0 b_0^R + q_1 b_1^R > 0$. Additionally, to avoid borderline cases, it will be assumed that $q_0 b_0^i \neq q_1 b_1^i$ for each player i. Finally, without loss of generality, the researcher will have $q_1 b_1^R > q_0 b_0^R$.

Of the above elements, the realized effort level and signal begin as the third party's private information, but the other players may be able to learn them. The others are common knowledge, including the prior distribution on the states of the world, the players' preferences and the outsider's research technology. In the end, one of the players makes a final decision. If it is the agent, then $x^A \in X^A = \{0, 1\}$ represents his decision. Otherwise, x^A represents his cheap-talk proposal to the leader, followed by her policy selection, $x^L \in \{0, 1\}$.

2.2 Communication, Authority, and Game Stages

The main focus of information transmission will be on the outsider's effort level and signal. These items are observable, which means that they can be truthfully conveyed or withheld, but not faked. However, these items of information will not be verifiable in the sense that a court could review effort, so contracts cannot be based on them. In making information observable, the model follows Ting (2008) and Gailmard and Patty (2013b). In particular, the model will assume that the leader and agent are equally capable of comprehending any information they receive, although the discussion will also consider the implications of a less capable leader.

Although communications between any two players will generally be unregulated, the rules for decision-making reduce the scope of communications that need to be analyzed. The game will take one of three forms: First, the leader can place the decision under her direct *administration*, so that she can communicate with the outsider and select the final policy, just as Congress can seek input from various interest groups in the legislative process. Second, she can commit to *delegation*, so that agent instead of the leader makes the final policy choice. Finally, she can engage in *oversight*, in which she retains decision-making authority but cannot communicate with the outsider. Instead, she can only receive the researcher's information if it transmits it to the agent and he relays it to her. The second two modes of the game apply to scenarios in which the leader lacks the time or ability to actively monitor the player exerting effort (see Tirole 1986, Aghion and Tirole 1997). With a large volume of regulatory policy under consideration each year, it is quite plausible that the leader will end up not directly observing many activities of outsiders (McCubbins, Noll, and Weingast 1987).

The three important directions for communicating information are from the researcher to the agent, from the researcher to the leader, and from the agent to the leader. In particular, since the outsider will always prefer to provide both items of information to the decision-maker, communications to the researcher are nugatory. Also, because researcher can communicate freely with the agent, information transmission from the leader to the agent can be ignored. Then, using ϵ_j^i and σ_j^i to denote player i's decision of whether to transmit to player j, respectively, the outsider's effort level and signal, when player i has the information and is allowed to relay the information, the decisions that can appear in the game are ϵ_A^R, σ_A^R, ϵ_L^R, σ_L^R, ϵ_L^A, and σ_L^A. For each of these variables, δ (ν) can be used to represent (non)disclosure.

In addition, the timing of disclosures can be substantially limited without loss of generality. When the leader has authority, the researcher's disclosures to the agent are effective only before his proposal. Meanwhile, disclosure of the researcher's information by the agent can only affect the outcome before the leader's policy decision. Thus, agent can make his disclosures simultaneously with his proposal. Furthermore, disclosures by the outsider to the leader, if allowed, will take place at the same time as the agent's proposal and disclosure decisions for notational convenience, since the agent's strategy will not affect that part of its communication strategy and vice versa. Similar timing restrictions can be applied when the agent has authority, except that the agent's proposal and leader's decision are substituted with the agent's decision.

With these restrictions on the nature and timing of information disclosure, the stages of the game can be succinctly stated as follows:

1) Nature chooses the state of the world $w \in \{0, 1\}$.

2) The researcher chooses the level of research effort e and receives a random signal s about the state of the world, whose accuracy increases with e.

3) The researcher decides whether to convey each of e and s to the agent.

4) The agent makes or proposes policy, $x^A \in \{0, 1\}$, and decides whether to relay each item of information he received from the researcher to the leader. Except under oversight, the researcher decides which of e and s to disclose to the leader.

5) Under administration or oversight, the leader makes the final policy decision, $x^L \in \{0, 1\}$.

2.3 Strategies and Beliefs

To notate pure strategies, it helps to distinguish between intended transmission and actual reception of information. The variables ϵ^i_j and σ^i_j defined above indicate whether player i would transmit information to player j given the opportunity. Reception of an item of information will be represented either by the true value in the case of transmission or by $\emptyset$ in the case of no transmission. Then, for the agent and leader, the sets $\mathring{E}_j \equiv \emptyset \cup \mathbb{R}_+$ and

$\mathring{S}_j \equiv \emptyset \cup \{0,1\}$ will represent the possibilities for what player j has learned about e and s, respectively, and $\mathring{e}_j$ and $\mathring{s}_j$, will respectively represent elements of these sets. Although the leader may receive information from either other player, the identity of the transmitter will turn out to be irrelevant, so additional notation for the sender can be omitted.

Now the parties' strategies can be expressed as follows: The researcher selects effort level e. After receiving the signal s, it decides what to communicate to the agent, and, when possible, to the leader. Its strategy is unaffected by the agent's, so its communications can be represented as ordered pairs $(\epsilon_i^R, \sigma_i^R) : E \times S \to \{\delta, \nu\}^2$, $i \in \{L, A\}$. Overall, the researcher's strategy can be denoted by $\Sigma^R \equiv (e; (\epsilon_A^R, \sigma_A^R); (\epsilon_L^R, \sigma_L^R))$. The agent's strategy consists of her proposal or decision and intentions to disclose and can be written as $\Sigma^A \equiv (x^A; (\epsilon_L^A, \sigma_L^A)) : \mathring{E}_A \times \mathring{S}_A \to \{0,1\} \times \{\delta, \nu\}^2$. Finally, the leader's strategy is just her policy choice when she has one, $\Sigma^L \equiv x^L : X^A \times \mathring{E}_L \times \mathring{S}_L \to \{0,1\}$. Notation for pure strategies is sufficient since mixed strategies do not play an important role except in borderline cases.

The fundamental set of beliefs center on the state of the world, which the players update as they receive information. The outsider's beliefs derive from its effort and the signal, which it always observes, so its posterior probabilities will not be notated. For the other active players $i \in \{L, A\}$, let β_1^i map the information s/he receives to a posterior probability (belief) that $w = 1$: i.e., $\beta_1^A : \mathring{E}_A \times \mathring{S}_A \to [0,1]$, and $\beta_1^L : X^A \times \mathring{E}_L \times \mathring{S}_L \to [0,1]$. Since full strategy-belief profiles are very extensive, only the most essential parts will be highlighted in the main text, and the propositions will describe only equilibrium path strategies.

2.4 Additional Terminology

In categorizing various scenarios, it will be useful to introduce some more terms. First, the limitations on the players' preferences described above imply that each strictly prefers one of the policies initially. It can be said that the player has a *bias* toward that policy, denoted by $\tilde{x}^i \equiv \arg\max_w q_w b^i_w$, and against the other policy, $1 - \tilde{x}^i$. Thus, the outsider is always biased toward policy 1, which is appropriate since interest groups are likely to consistently favor one side of an issue. For example, pharmaceutical companies generally want approval for their drugs, at least in the absence of contradictory information. In relation to the outsider, each of the other active players will be termed *advocative* if s/he also has a bias toward policy 1 and *adversarial* if s/he is biased toward policy 0. Players can differ not only in terms of the policy toward which they are biased, but also in the strength of their biases. The degree of a player's bias toward policy x can be measured as $B^i_x \equiv 2q^i_x b^i_x/(q_0 b^i_0 + q_1 b^i_1) - 1$, so that the quantity is negative if the player is biased against that policy.

Intuitively, a player will always want policy to follow a signal that follows that player's bias, while that player will want policy to match a contradictory signal only if it is supported by enough effort. A formal condition can be stated:

Lemma 1. *After research, a player prefers to have $1 - \tilde{x}^i$ enacted if and only if $s = 1 - \tilde{x}^i$ and $e \geq g^{-1}(q^i_{\tilde{x}^i} b^i_{\tilde{x}^i}/(q_0 b^i_0 + q_1 b^i_1)) \equiv e^i$. Otherwise, that player strictly prefers $\tilde{x}^i$.*

Proof. Proofs of all numbered results except Corollaries 4 and 13 are in Appendix A. ■

The quantity e^i can be called a player's *standard of proof*, which increases with the bias toward the policy toward the policy that player initially prefers. This quantity is the minimum effort level at which the leader and agent prefer to have the policy follow the signal rather than to always match his or her presumptive preference. Following Bayes' rule, the agent or leader's expected payoff when the policy follows the signal is

$$EU^i_f(e) \equiv g(e)(q_0 b^i_0 + q_1 b^i_1), i \in \{L, A\}, \tag{1}$$

which is increasing in e. Then the standard of proof e^i is the minimum level of effort that satisfies $EU^i_f(e) \geq q_{\tilde{x}^i} b^i_{\tilde{x}^i}$.

A few more terms for the researcher are worth defining. First is its *signal-constrained optimum*, the amount that it would devote to its research knowing that the signal would be followed. Bayes' Rule implies that this expected payoff is

$$EU^R_f(e) \equiv g(e)(q_0 b^R_0 + q_1 b^R_1) - c(e). \tag{2}$$

Then the signal-constrained optimum, denoted by $\hat{e}$, satisfies the first-order condition

$$g'(\hat{e})(q_0 b^R_0 + q_1 b^R_1) = c'(\hat{e}). \tag{3}$$

Related to this effort level is how the payoff in Equation (2) compares to $q_1 b^R_1$, its payoff if it selects policy 1, toward which it is biased, with no research effort. It is *(un)motivated* (i.e., to do research) if $EU^R_f(\hat{e}) > (<) q_1 b^R_1$.[1] Finally, the most effort that the researcher is willing to expend and have the signal be followed, rather than have policy x always be chosen, is its *discouragement point* for that policy. This point is defined as $\bar{e}_x \equiv \max\{e : EU^R_f(e) \geq q_x b^R_x\}$, with $\bar{e}_1$ existing only for a motivated researcher. The functional form assumptions on $g(\cdot)$ and $c(\cdot)$ imply that $\bar{e}_0 > \hat{e}$, and, when $\bar{e}_1$ exists, that $\bar{e}_1 \in (\hat{e}, \bar{e}_0)$.

3 Model Results

The solution concept for this game is perfect Bayesian equilibrium (PBE) in pure strategies. Understanding why the leader would want to delegate to an agent or cut off her communi-

[1] The borderline case of $EU^R_f(\hat{e}) = q_1 b^R_1$ does not add any insight and is omitted.

cations with the outsider requires an analysis of the equilibria under each of the three game forms and comparison of her payoff given various parameters for the three players. Since the researcher generates the information, a useful benchmark is how the leader would fare if it had authority to set the policy. Depending on its preferences and research technology, it would either maximize its payoff under Equation (2) or summarily choose policy 1:

Proposition 2. *If allowed to select the policy, an unmotivated researcher would set* $e = 0$ *and* $x = 1$*, while a motivated researcher would set* $e = \hat{e}$ *and* $x = s$*.*

This result, can be used to represents total lack of regulation. In the drug approval setting, it implies that a firm might market a drug without doing any research on it. Though it may be difficult to imagine a setting in which people would dare to sell drugs without doing any research, Congress estimated in 1906, the year when it first legislated federal controls on drugs, that there were 50,000 so-called "patent medicines" in the drug industry (Carpenter 2010, 77-78). This proposition indicates that there is much scope for improvement. For example, an adversarial leader facing an unmotivated researcher would do better at least by always selecting policy 0.

3.1 Equilibria under Administration

In the administration game form, the leader has final policymaking authority and can scrutinize whatever information that the researcher offers. It turns out that, in equilibrium, researcher has no problem disclosing its information to the leader. The next proposition states the equilibria under administration in general terms:

Proposition 3. *Under administration the unique PBE with respect to effort and policy choice is* $e^* = \max\{\hat{e}, e^L\}$ *and* $x^{L*} = s$ *when* $e^L < \bar{e}_{\tilde{x}^L}$ *and* $e^* = 0$ *and* $x^{L*} = \tilde{x}^L$ *when*

$e^L \not\leq \bar{e}_{\tilde{x}^L}$, *except that when* $e^L = \bar{e}_{\tilde{x}^L}$, *both equilibria can obtain. The researcher can induce the first equilibrium with* $\epsilon_L^R = \sigma_L^R = \delta$ *when* $s = 1 - \tilde{x}^L$ *and the second with* $\epsilon_L^R = \delta$.

Remark. The condition that $e^L \not\leq \bar{e}_{\tilde{x}^L}$ can result from $e^L > \bar{e}_{\tilde{x}^L}$ or an unmotivated outsider's lack of a value of $\bar{e}_{\tilde{x}^L}$ when $\tilde{x}^L = 1$. Also, though it may communicate with the agent, that player is unnecessary because it can directly convey information to the leader.

The specific equilibrium results depend on the leader's bias. For an advocative leader, whether the researcher is motivated also matters. With an unmotivated researcher, effort is zero and policy 1 always obtains, the same result as if the researcher were acting on its own. With a motivated researcher, how much effort she induces depends on her standard of proof. When $e^L \leq \hat{e}$, her standard of proof is not binding, and the outsider expends effort at its signal-constrained optimum for policy matching the signal, as if it had decision-making authority. When $e^L \in (\hat{e}, \bar{e}_1]$, her standard of proof is binding and induces the outsider to exert additional effort to meet the standard of proof, after which policy matches the signal. Finally, when $e^L > \bar{e}_1$, the researcher is unwilling to incur the cost needed to meet the standard of proof and instead sets effort at zero for policy 1. When the leader is adversarial, then only her standard of proof is relevant. The results are the same as those for an advocative leader facing a motivated agent, except that the upper bound for effort that she can extract is the discouragement point for policy 0.

Compared to allowing the researcher to decide policy, taking control of decision-making only helps the leader, primarily because she can always summarily select the policy toward which she is biased to earn a reservation payoff of $q_{\tilde{x}^L} b^L_{\tilde{x}^L}$. Both adversarial and advocative leaders can induce motivated researchers to expend additional research to meet their standard of proof. Adversarial leaders particularly benefit because they can also stimulate unmotivated researchers to exert effort in the first place, and because they can incentivize motivated researchers to increase their effort up to a higher limit, since $\bar{e}_0 > \bar{e}_1$.

However, given that the leader has authority to choose the policy, waiting for the outsider's research often does not benefit her compared to summarily selecting the policy for which she has a bias as if the researcher were not present. Whenever her standard of proof is binding and does not discourage effort altogether, her payoff is the same as if she committed to selecting $\tilde{x}^L$ from the beginning. Meanwhile, when a motivated researcher's effort is its signal-constrained optimum, the result is the same as if it were by itself. The only case in which the leader exceeds the reservation payoff of $q_{\tilde{x}^L} b^L_{\tilde{x}^L}$ from research that the outsider would not have conducted by itself is when she is adversarial, it is unmotivated, and its signal-constrained optimum exceeds her standard of proof.

These difficulties result from the fact that, when the standard of proof binds, the researcher can expend just enough research to make her weakly prefer to select policy according to the signal, in which case it effectively denies the leader any surplus over her reservation value. The effort and signal are observable, and the outsider is willing to disclose both of these items, so these limits on the leader's payoff do not result from any informational advantage that the researcher retains. Because she can receive both items of information, she cannot commit to summarily select the policy toward which she is biased for any effort level above her standard of proof. Thus, she would prefer to prevent low-quality information from reaching her, but she cannot when she has decision-making authority and the "freedom" to communicate with the third party.

3.2 Equilibria under Delegation

One potential solution to the problem that the leader faces in having to receive and respond to information from the outsider is to prevent herself from doing the latter by irrevocably granting an agent the authority to set policy. Under delegation the agent assumes the leader's role, receiving information and making a decision identically. Substituting the agent for the leader in Proposition 3 yields an analogous equilibrium result:

Corollary 4. *Under administration the unique PBE with respect to effort and policy choice is $e^* = \max\{\hat{e}, e^A\}$ and $x^{A*} = s$ when $e^A < \bar{e}_{\tilde{x}^A}$ and $e^* = 0$ and $x^{A*} = \tilde{x}^A$ when $e^A \nleq \bar{e}_{\tilde{x}^A}$, except that when $e^A = \bar{e}_{\tilde{x}^A}$, both equilibria can obtain. The researcher can induce the first equilibrium with $\epsilon^R_A = \sigma^R_A = \delta$ when $s = 1 - \tilde{x}^A$ and the second with $\epsilon^R_A = \delta$.*

Whether the leader's payoff is higher or lower delegating to an agent depends on their preferences and whether the agent is motivated. Ignoring borderline cases, one can formally state when delegation is better for her as follows:

Proposition 5. *The leader's payoff is higher from delegation than from administration: (a) when $\max\{e^L, \hat{e}\} < e^A \leq \bar{e}_{\tilde{x}^A}$ and (b) when $e^A \leq \hat{e}$ and $e^L < \hat{e}$, with $\tilde{x}^A = 0$, $\tilde{x}^L = 1$, and an unmotivated researcher.*

The generally necessary condition $\max\{e^L, \hat{e}\} < e^A < \bar{e}_{\tilde{x}^A}$ corresponds to two intuitive principles. First, the agent's standard of proof must not exceed the discouragement point corresponding to his bias. Otherwise, the outsider will do no research, and the agent will summarily select the policy for which he has a bias, which the leader cannot strictly prefer to making this kind of choice herself. Second, among the leader's and agent's standards of proof and the outsider's signal constrained optimum, the agent's standard of proof must be the highest. If the leader's standard of proof is the highest, then even if the agent induces research, the effort level will less than needed to satisfy her, in which case she would better off summarily selecting the policy according to her bias. If the signal-constrained optimum is the highest, either the agent or the leader will, for the most part, induce effort at $\hat{e}$, making administration and delegation equally good for the leader. In contrast, when the agent's standard of proof is the highest and does not exceed the relevant discouragement point, he can induce extra effort in a way that satisfies the leader and increases her expected payoff in the form of policy that is more likely to be correct.

For Proposition 5(b), an unmotivated researcher would induce policy 1 from an advocative leader by exerting and disclosing zero effort, but it must meet an adversarial agent's standard of proof. The policy can at best (for the outsider) match the signal, so it maximizes at its signal-constrained optimum, which meets both players' standards of proof.

Overall, when committing authority to an agent with the right preferences benefits the leader, it does so usually by forcing the researcher to satisfy the agent's standard of proof that exceeds hers. This mechanism differs from the logic that agencies gather information that political leaders cannot (Bendor and Meirowitz 2004). Here, delegation prevents the leader from acting on information that she might receive from the researcher. As in other agency models, whether the leader can commit to delegate is a significant issue; however, unlike in canonical models, this difficulty does not persist past the agent's policy choice, since the leader would not want to reverse the agent's decision ex post (cf. Callander 2008).

In situations outside those in Proposition 5, delegation does not improve the leader's payoff and will often reduce it. If she can choose whether to delegate, she can avoid cases in which she would do worse than under administration. However, if an agent must be chosen for many decisions, then she may have to trade off cases in which she gains against those in which she loses. Surrendering authority to an agent is a rather blunt way of avoiding the challenge of a researcher too willing to provide low-quality information. The remaining form of the game provides another way for the leader to increase her policy payoff over administration.

3.3 Equilibria under Oversight

Oversight is an intermediate form of control: she can still make the final decision, but she cannot obtain information directly from the researcher. Instead, it can only communicate with the agent, who then decides what, if anything, to convey to the leader along with his policy proposal. This game form helps distinguish the effect of withholding information from the effect of committing authority. Equilibria are no longer necessarily unique in this setting, even in terms of effort and policy choices. However, it is always possible to identify the PBE that yields the leader her highest payoff.

To begin with, whenever delegation yields the leader a higher payoff than administration, there exists a functionally equivalent oversight PBE in which the agent never discloses the researcher's information and the leader always ratifies the agent's proposal:

Proposition 6. *Under oversight, when* $e^A \leq \bar{e}_{\tilde{x}^A}$ *and* $e^L \leq \max\{\hat{e}, e^A\}$*, there exists a PBE that maximizes the leader's equilibrium payoff with* $e^* = \max\{\hat{e}, e^A\}$*,* $\epsilon_A^{R*} = \sigma_A^{R*} = \delta$ *when* $s = 1 - \tilde{x}^A$*,* $\epsilon_L^{A*} = \sigma_L^{A*} = \nu$*, and* $x^{A*} = x^{L*} = s$*.*

In an overlapping set but not identical of circumstances, there exists another PBE that maximizes the leader's payoff, one in which an adversarial agent proposes policy 1 when discloses the researcher's effort level and signal when the former meets his standard of proof and the latter points to policy 1, but discloses nothing and proposes policy 0 otherwise. As in the other equilibrium, the leader always accedes to the agent's proposal, although in fact a proposal is not necessary, since the agent's intentions can be inferred from his disclosure choices.

Proposition 7. *Suppose the game form is one of oversight,* $\tilde{x}^A = 0$*, and* $e^A \leq \bar{e}_0$*. In addition, if* $\tilde{x}^L = 1$ *and* $e^L \leq \max\{e^A, \hat{e}\}$*, or if* $\tilde{x}^L = 0$ *and* $e^L \leq \bar{e}_0$*, there exists a PBE* *that maximizes the leader's equilibrium payoff with* $e^* = \max\{\hat{e}, e^A, e^L\}$*,* $x^{A*} = x^{L*} = s$*,* $\epsilon_A^{R*} = \sigma_A^{R*} = \epsilon_L^{A*} = \sigma_L^{A*} = \delta$ *when* $s = 1$*, and* $\epsilon_L^{A*} = \sigma_L^{A*} = \nu$ *when* $s = 0$*.*

Compared to delegation, oversight does not increase the leader's payoff above that under administration in many more cases. However, oversight also does not result in a lower utility for the leader than administration in many situations in which delegation would. Intuitively, the leader preserves her payoff under administration with the decision-making authority that she retains under oversight, often by summarily selecting the policy toward which she is biased. The only scenario in which oversight underperforms administration is when an adversarial agent discourages research in the former game form, whereas an adversarial leader in the latter induces research that yields her a surplus above her reservation payoff from always selecting policy 0, $q_0 b_0^L$. In these scenarios, however, delegation yields an equally low payoff.

Overall, when all parameter values are considered, one can find that oversight achieves all the benefits of delegation compared to administration with few of delegation's costs.

Theorem 8. *Assume that, under oversight, a PBE that maximizes the leader's equilibrium payoff obtains. Then the three game forms can be ranked in terms of her utility as follows:*

(a) Whenever delegation outperforms administration, oversight does so equally.

(b) Administration outperforms delegation when (i) $\tilde{x}^A = \tilde{x}^L$ *and* $\max\{\hat{e}, e^A\} < \min\{e^L, \bar{e}_{\tilde{x}^L}\}$*, (ii)* $\tilde{x}^A = \tilde{x}^L$*, the outsider is motivated, and* $e^L < \hat{e} < \bar{e}_{\tilde{x}^L} < e^A$*, and (iii)* $\tilde{x}^A \neq \tilde{x}^L$ *and either* $e^A \not\leq \bar{e}_{\tilde{x}^A}$ *or* $\max\{e^A, \hat{e}\} < e^L$*. Oversight yields her as much as administration, but not more, except possibly when* $\max\{e^L, \hat{e}\} \leq e^A$ *and* $e^L \leq \bar{e}_{\tilde{x}^L}$ *in case (iii).*

(c) Administration equally outperforms delegation and oversight when $\tilde{x}^A = \tilde{x}^L = 0$*, the outsider is unmotivated, and* $e^L < \hat{e} < \bar{e}_0 < e^A$*.*

(d) Three game forms yield the same payoff in the remaining cases: (i) $\max\{e^A, e^L\} \leq \hat{e}$*, apart from when* $\tilde{x}^L = 0$ *and* $\tilde{x}^A = 1$ *with an unmotivated outsider, and (ii)* $\tilde{x}^L = \tilde{x}^A$*,* $\hat{e} < e^L$*, and* $e^A \nleq \bar{e}_{\tilde{x}^A}$*.*

3.4 Optimal Choice of Game Form and Agent

Oversight and delegation each have the potential to benefit the leader under in various settings. For institutional design purposes, however, the most useful equilibria for her in these two modes are those that can apply regardless of whether the leader is adversarial and regardless of whether the researcher is motivated or not. The motivation for this criterion is that the leader cannot control her preferences or those of the outsider, but she may be able to influence the agent preferences that apply through the choice of agent (see Bertelli and Feldmann 2007). Then the most readily helpful equilibria are those involving an adversarial agent in Propositions 5–7.

In general, an adversarial agent with a greater standard of proof is better up to a point, since he induces additional effort. However, a standard of proof that is too high can discourage the outsider from research altogether. This intuition underlies the next result:

Proposition 9. *Suppose* $\bar{e}_0$ *is fixed and the leader can select among a set of adversarial agents with* $e^A \leq \bar{e}_0$*. Also, suppose she does not know* $\tilde{x}^L$ *or* e^L *when she selects the game form and agent but will know after the agent's proposal. If oversight is available, she maximizes her utility with that game form and the agent with the highest* e^A*. If not, she maximizes her utility with delegation or administration and the same agent.*

Not counting the potential equilibria in Theorem 8(b)(iii), if the leader can select an agent with any preferences, her best agent is one who requires the maximum amount of effort that does not discourage research (cf. Bueno de Mesquita and Stephenson 2007, 614–15). When the leader cannot select an agent for each policy decision, calibrating his standard of proof to match the researcher's discouragement point for policy 0 is likely infeasible. Nonetheless, since the leader's payoff increases with effort when policy matches the signal, the usefulness of an agent who can enforce a higher standard of proof remains important. Since an adversarial agent with a higher standard of proof also has a greater bias toward policy 0 than the leader, his and the researcher's biases toward that policy will sometimes lie on opposite sides of hers.

It is worth noting that the "right" agent benefits the leader solely by virtue of his preferences, rather than because of his expertise. Other models that rely just on the agent's preferences have differing results about what kinds of agents are beneficial. First, Proposition 9 contrasts considerably with models in which the best agent preferences lie in between the leader's and the researcher's so that he can elicit more precise messages about its private information in a delegated cheap-talk setting (Dessein 2002, Gailmard and Patty 2013a). When the agent merely proposes rather than sets a policy, however, the leader benefits instead from a well-chosen agent with preferences relative to hers on the opposite side of the third party's (Ivanov 2010, Ambrus, Azevedo, and Kamada 2013). In these mediated cheap-talk models, the third party is willing to transmit more detailed messages because, in response to the agent's incentive to propose policies further away from the other two players' preferred ones, the leader will select policies closer to her and its preferred ones. Here, in contrast, the result lies in inducing additional effort from the third party by making it fearful of policy that is, in expectation, more adverse to its interests.

Other models highlighting the benefits of a more adversarial agent can be found in Bertelli and Feldmann (2007), in which his extreme preferences offset those of an interest group in policy bargaining; and Rogoff (1985), in which a conservative central banker with extra concern about reducing inflation beneficially does so given wage-setters' attempts to anticipate his response to economic shocks. Of these models, Rogoff's model most closely approximates the reverse rationale of having an agent to prevent the leader from receiving or using the same information,[2] since a conservative central banker does better than an equally expert leader with preferences matching social welfare. The current game points to a larger set of policymaking settings in which the reverse rationale may apply.

Overall, the model suggests that incorporating an adversarial agent with a high standard of proof and giving him the exclusive authority to make policy or the sole ability to communicate with an outside group helps the leader avoid the problem of an outsider's providing information of just barely sufficient to satisfy her. The next section considers ways in which the researcher might try to frustrate this institutional arrangement.

4 Extension of the Model to Regulatory Capture

The assumption, maintained until now, that the leader has preferences identical to the "public," the true principal, is relaxed with the possibility of regulatory capture. Although capture has various definitions (see Levine and Forrence 1990, Dal Bó 2006), the term here can be understood as steps by the researcher to influence the leader or agent such that the principal's payoff decreases. By Proposition 2, a motivated researcher would like to research at its signal-constrained optimum and have policy follow the signal, while an unmotivated researcher would like to avoid expending any effort and induce summary selection of policy 1. Thus, it has a reason to attempt regulatory capture as well as a maximum degree to which it is willing to do so.

[2] In particular, cheap-talk models involve only messages from the third party about its private information, because by assumption, it is not able to credibly disclose the information that is generated.

The outsider can attempt to influence either player. As Carpenter (2013) observes, one can distinguish statutory capture, which occurs apart from any agency action, from agency capture, in which an outside frustrates legislative intent through its influence on the agency. In this section, statutory capture will be represented by attempts to influence the leader, whereas agency capture will be modeled as steps to influence the agent. For each other active player, it has two techniques, τ, for capture: First, it can engage in *bias-shifting* (β), in which it causes a player's policy preferences, b^i_w, to change so that he or she has a different bias with respect to the two policies. Second, following Laffont and Tirole (1991), it can effectuate a *quasi-contract* (κ), in which a player is compensated for taking a different action than his or her policy preferences would dictate.

The outsider's cost for bias-shifting directed at player i can be denoted as a function $c^i_\beta(B^i_0 - \check{B}^i_0)$, where B^i_0 is that player's natural bias toward policy 0 and $\check{B}^i_0$ is that player's final bias when captured. Meanwhile, the cost of quasi-contractual compensation can be represented as $c^i_\kappa(V^i - \check{V}^i)$, where V^i is the player's policy payoff in an equilibrium without capture[3] and $\check{V}^i$ is that player's payoff from policy set according to the quasi-contract. It is convenient to further define $\Delta B^I_0 = B^i_0 - \check{B}^i_0$ and $\Delta V^i = V^i - \check{V}^i$. Since the goal is merely to understand how the different mechanisms operate for each player, rather than to define the researcher's optimal combination of capture strategies, it is sufficient to specify that, $\forall \tau \in \{\beta, \kappa\}, \forall i \in \{L, A\}$, c^i_τ is a strictly increasing function of its argument, to indicate roughly that more capture is more difficult for the researcher.

[3]In the case of oversight, the relevant equilibrium is the one yielding the principal the highest payoff.

4.1 Attempts at Statutory Capture

For statutory capture, the two methods of influencing the leader have different effects because she can select delegation or oversight. First, if the leader is using oversight and delegation in a particular case and it will yield more than her administration payoff, bias-shifting requires a fixed cost for her to be willing to return to administration. Proposition 9 implies that small values of ΔB_0^i do not help the outsider:

Proposition 10. *Suppose $e^L \leq e^A \leq \bar{e}_0$ and $\tilde{x}^A = 0$ and the principal can select among game forms. Bias-shifting of the leader does not affect the research effort or policy selection as long as $e^L \leq e^A$ continues to hold.*

Thus, if the leader can rule out quasi-contracts, she can mitigate capture with a strongly adversarial agent, even though it is subject to bias-shifting that might come from political pressure. To benefit from bias-shifting, the researcher would need to make ΔB_0^L large enough for leader to prefer administration and summarily selection of policy 1.

For a quasi-contract, small amounts of compensation to the leader, (i.e., values of ΔV^L near zero) can cause her to select an agent with less bias toward policy 0. As compensation to the principal increases, the researcher can induce her to take actions that are correspondingly more favorable to it.

Proposition 11. *Suppose that $e^L < \bar{e}_0$ and the leader can choose the game form and an adversarial agent with any $e^A \bar{e}_0$. For quasi-contracts with the leader, a non-empty interval $[0, \Delta_1 V^L)$ exists in which the researcher can only induce her to select an adversarial agent with a lower standard of proof. For some $\Delta_2 V^L \geq \Delta_1 V^L$, it can induce her to adopt her strategy under administration. If this policy outcome differs from what it would select acting alone, there exists $\Delta_3 V^L > \Delta_2 V^L$ such that the policy outcomes of Proposition 2 obtain.*

Therefore, a key contrast between statutory bias-shifting and quasi-contracts is that the former operates in an all-or-nothing fashion, while the latter can achieve more graduated results, starting from minimal costs. If the two methods are combined, then they may substitute for each other. For example, if an unmotivated researcher agrees with the leader in principle to have an agent with a lower standard of proof, then it needs a lower level of bias-shifting to induce the leader to always select policy 1.

4.2 Attempts at Agency Capture

Agency capture becomes relevant when the leader chooses delegation or oversight. Unlike for the leader, bias-shifting and quasi-contracts for the agent are essentially equivalent methods in the following sense:

Proposition 12. *Starting from any adversarial agent with* $e^A \in (\max\{\hat{e}, e^L\}, \bar{e}_0]$ *in oversight or delegation, the researcher can effect any standard of proof under capture,* $\check{e}^A < e^A$*, with* $\check{e}^A \in (\max\{\hat{e}, e^L\}, \bar{e}_0]$*, while remaining adversarial, with some level of bias-shifting* ΔB_0^A *or amount of compensation* ΔV^A*.*

The result that bias-shifting can have the same impact as a quasi-contract is consistent with the notion that so-called cultural capture, by which an interest group influences agency officials' preferences through human contact and can thereby sway regulation, as well as the argument that focusing on interest-based capture is incomplete (Kwak 2013). However, the exchangeability of cultural capture and interest-based capture does not extend to statutory capture because the leader has delegation and oversight. She can escape bias-shifting through the use of another player, but, by assumption, the agent cannot.[4]

[4]Even if an agency could employ yet another party to avoid directly facing the interest group, doing so would only remove the problem one step, as the group could seek to capture that party.

For completeness' sake, it is worth observing that the degree to which the outsider would want to influence the agent depends on what game form the leader selects and whether she can change the game form based on whether capture is occurring. If the leader delegates and cannot revoke the agent's authority when the outsider exerts its influence, then the outsider may want to influence the agent so that his standard of proof falls below the principal's. Otherwise, it can only recreate the policy outcome that would obtain under administration.

4.3 Inferences about Capture

With the mechanisms of capture clarified, it becomes possible to determine how the outsider's influence can be inferred from actions taken by the players. While it might be possible to observe capture directly, such as with a recording of a conversation about a quid pro quo, it is realistically likely that a player consciously subject to capture would act so that such evidence cannot be discovered. Thus, for the remainder of the discussion, actions from which the public can detect capture will be limited to what the leader can observe in the baseline model: her choice of game form and agent, her or his policy decision (or, under oversight, the agent's proposal), and any of the researcher's information that she receives.

In the context of this model, statutory capture is likely to be quite difficult to detect. Choosing an agent with a low standard of proof would show that she was subject to capture by quasi-contract. However, it may not be clear which available agent has the greatest bias toward policy 0 that benefits the principal. Furthermore, if there is ex ante uncertainty about the researcher's preferences or research technology, the leader would be right to select an agent with a somewhat lower standard of proof to prevent him from discouraging the outsider's research.

A clearer sign of statutory capture would the leader's decision not to use an agent at all, but to make the policy decision herself via administration. This evidence is also not unambiguous since she might not have any agent available who would benefit her under delegation or oversight. However, it can be argued that, in relative terms, opting for administration is stronger evidence than choosing an agent different from what the principal would prefer. In the case of Congress, this intuition contrasts with iron-triangle style arguments that agencies exist to benefit the interest groups that they regulate, and even that Congress creates this arrangement. Though delegation to or oversight of an agency may represent an attempt to avoid "making a hard decision," this avoidance can be socially beneficial when the problem is that outside groups will only submit information that barely satisfies legislators. Thus, employing an agent is not only not a indication of congressional capture, but it can also be a means to mitigate capture in the form of statutory bias-shifting.

Meanwhile, agency capture could be inferred if the agent transmits information about an effort level below his standard of proof, assuming that the latter is known. If the researcher has engaged in capture, she might try to withhold the outsider's effort level from the leader. However, the converse is not necessarily true, as Proposition 6 indicates that, given the right parameters, she (and thus the principal) can benefit when the agent adopts a policy of nondisclosure. Applied to Congress, this result responds to claims that it has neglected its oversight responsibility in a way that is different from arguments that oversight is actually robust (Aberbach 1990) or that it has fire alarms as an alternative (McCubbins and Schwartz 1984). Instead, lack of oversight facilitates better policy-making by inducing more policy research effort from outside research groups.

A similar argument can be made about the oversight equilibrium in Proposition 7, in which an adversarial agent only discloses the researcher's information when the effort is high enough and the signal contradicts her bias toward policy 0. Although the agent seems to be reporting only the researcher's successes in this equilibrium, the effort that he reveals exonerates him of capture. Moreover, the leader would not want the agent to report all research results, because then the researcher would effectively be able to transmit low-quality information to her, which would result in a lower payoff for the leader and the principal. Like the leader's decision to use an agent, these disclosure patterns are not only not necessarily signs of capture, but they also can facilitate her attempt to avert the effects of statutory bias-shifting.

4.4 Measures Against Capture

The model also has implications for what kinds of measures are likely to be effective in combating capture. In theory, three methods can be considered: (1) complete transparency of the researcher's information whenever the agent has them, (2) transparency of its information only when the agent proposes policy 1, and (3) efforts to keep the agent's bias for policy 0 relatively high. In practice, the first measure roughly corresponds to President Obama's Open Government Initiative (see Coglianese 2009), while the third corresponds to his attempts to tighten ethics rules for executive branch officials (see Thurber 2011).

The discussion about detecting agency capture makes clear that the first measure, transparency of both items of the research information, will generally be useless or even counterproductive. The nature of delegation and oversight implies the following result:

Corollary 13. *Disclosure of the researcher's effort level and signal to the leader does not affect policy outcomes under delegation but causes the administration policy outcomes to obtain under oversight.*

In particular, complete transparency means that the researcher can effectively communicate directly with the leader as under administration. Under oversight, this measure, if intended to mitigate capture, will ironically allow the researcher to achieve what it would want from capture without having to engage this type of activity. The Obama Administration's exhortation to agencies to release documents more proactively under the Freedom of Information Act (Coglianese 2009, 533) may work for prior policymaking decisions that were not transparent, but it may have unintended consequences for future decisions to the extent that it "successfully" induces more disclosure.

Corollary 13 also adds some nuance to one of the key results in Ting (2008), which states that having an employee report to the principal when the manager would reject a project regardless of its quality can only benefit the principal. That result roughly corresponds to the case in Theorem 8(c), in which the agent's standard of proof is so high that it discourages the outsider from researching at all. However, part (a) of Theorem 8 highlights cases in which a moderately high standard of proof for an adversarial agent can benefit the principal, but only when the agent can withhold either the researcher's information. Ting (2008) does not include an analogous result because it considers only two quality levels. The oversight equilibrium results in the present model suggest that, if multiple quality levels are possible in a whistleblowing setting and the manager accepts projects at fewer quality levels than the employee and the principal, then the principal's desire to have the employee report project quality might be less absolute. Discouraging whistleblowing could encourage the employee to exert more effort so that the quality is high enough for the manager to approve, whereas encouraging it might incentivize the employee to put in less effort and indicate a quality level that is satisfactory for the politician but not the manager, resulting in more approvals of lower quality projects.

The next option for combating is a conditional form of transparency, in which the agent reports the researcher's information only when he proposes policy 1. Since method would effectively detect capture, an important empirical question whether it can be implemented. One challenge that may arise is defining "policy 1," although this identification is simple for some categories of policymaking, like drug approvals. Also, in the face of uncertainty about the agent's preferences, it may be unclear upon observing a fairly low level of effort whether the agent has been captured or is merely acting according to a weaker bias toward policy 0 that he naturally has. In addition, since there are two oversight equilibria that yield the same policy outcomes under the conditions in Theorem 8(a), changing from the one in which the agent discloses nothing to one in which the agent transmits information with a proposal of policy 1 could be a challenge. To the extent that equilibrium selection represents culture (see Kreps 1990), it may not be easy for an agent accustomed to the former equilibrium to transition to the latter one.

The third possibility is preventing the adversarial agent from lowering his standard of proof. In the model, this entails increasing the cost of bias-shifting and quasi-contractual compensation. The Office of Government Ethics (OGE) proscribes for executive branch officials various kinds of behavior linked to influence by interest groups, such as gifts exceeding a nominal value and employment in a related industry after too short a period of time. These measures are designed in part to prevent officials from biasing their policymaking toward interest groups, including in an unconscious way. If it is difficult to stop capture at the policymaking stage, it is arguably helpful for OGE to prevent bureaucrats from becoming less adversarial in the first place through interactions outside of any decision-making processes.

The main challenge is in enforcement. Unambiguously illicit activity, like bribery, requires substantial resources to punish and deter, although the same might be said for information nondisclosure if an agency can claim that it lacks information. However, much of the behavior that ethics regulations target is legal except for government officials. Thus, if they are aware of what actions are improper, they are likely to abstain from them and truthfully certify that they have done so on reporting forms. Furthermore, some restricted activities, like post employment lobbying, cannot be hidden. If maintaining agency officials' preference is feasible, then the Obama Administration's ethics reforms measures for executive branch officials (see Thurber 2011) are likely to be more effective than unconditionally applied transparency measures, as the Obama Administration's seem to be (see Coglianese 2009).

5 Application to FDA Pharmaceutical Regulation

A general implication of the model is that, rather than administer a policy program herself, a principal can do better if she employs the right kind of agent either to make the final policy decisions or to be the sole collector of information. The reason is not so that the agent can obtain knowledge that the principal cannot, but something of the reverse: so that he can prevent the principal from receiving information that she is perfectly capable of understanding. A policy area that arguably implicates many of the features of this model is the FDA drug approval process. In terms of the game, either the agency or its employees serve as adversarial agents facing the outsiders, pharmaceutical companies. This application heavily on the account in Carpenter (2010), so unless otherwise noted, page numbers in this section are citations to this source.

5.1 The FDA as Adversarial Gatekeeper

When an agent benefits the principal, it is because he stimulates more research effort from the outsider. The outsider is not formally required to engage in any level of research, but it cannot have policy 1 enacted if it does not submit research that satisfies the him. Carpenter observes this kind of dynamic operating in drug regulation when he observes, "the Administration's gatekeeping power enacts a system of incentives that induces the production of far more information (and higher quality information) from drug companies and medical researchers than would otherwise have occurred " (751). In particular, the FDA's gatekeeping power "stems from its ability to veto product entry" (16).

In addition to showing the value of gatekeeping, the model also suggests that this power encourages the most research if the agent is highly adversarial toward the outsider. In the case of the FDA, its reviewers induce large amounts of research arguably because they would prefer that the drug not be marketed in the absence of sufficient evidence supporting the drug. This notion is consonant with the idea that "the agency would have to negate an appreciable fraction of new drug applications. If approval became so happily predictable as to become perceivably deserved, the incentives for drug companies to conduct exhaustive, careful, and clinical trials would vanish" (493). Although rejecting some applications, regardless of personal preferences, might be a viable strategy in a repeated game context, such a strategy is at least easier to pursue if FDA reviewers actually value safety over drug innovation a priori.

There is some evidence that these reviewers are adversarial. To begin with, the FDA was one of the two main forces that maneuvered for amendments to the 1906 Pure Food and Drugs Act (80), and the act that passed gave the agency its current gatekeeping authority. In general, the FDA seems to have been more consistently adversarial compared to the general public and the most vocal interest groups. It has had to withstand criticisms of a so-called "drug lag," according to which it was allegedly taking too long to approve new medicines (374); campaigns by patient advocacy groups to make new cancer and AIDS drugs available (410–11, 429), and calls to use external reviewers for new drug applications (NDAs) (458). The highly adverse media response to the agency's initial reaction of the drug Activase in 1987 (3–4) supports the general intuition that the people, directly or through their elected representatives, might more readily approve a drug if it could directly access the relevant information and make the final decision. Although there have also been congressional hearings questioning whether the FDA should have allowed particular drugs, they do not establish that Congress or the relevant committees are generally more adversarial than the agency's policymakers: first, there have been hearings expressing concern about slow approvals and lack of innovation (337), and second, as described in more detail below, even the first type of hearings may reflect institutional design concerns rather than committee members' underlying policy preferences.

It is harder to show directly that an adversarial stance is necessary for the FDA's gatekeeping authority to be effective, since there does not appear to be a period in which the FDA consistently approved drugs with too little evidence. However, there is evidence in other settings suggesting that, in general, gatekeeping power alone is insufficient to induce probing research. In related area of medical devices, Harris (2008) has reported in the *New York Times* that "disputes tend to pit agency managers, who often lean toward approving drugs or devices when the data are equivocal, against agency scientists, who want more certain trial results before allowing the products to be sold" (A15). A different agency, the now-defunct Minerals Management Service (MMS), had the authority to reject oil and gas lease applications based on safety an environment concerns, but it appears to have approved applications even when its scientists concluded that these were significant issues (see Urbina 2010, May 14). More generally, the MMS "faced criticism . . . for generally favoring the oil industry over public and environmental safety concerns" (Neill and Morris 2012, 636).

A more adversarial agency could conceivably have induced more research as to whether prospective lessees could adequately and cost-effectively address potential hazards.

5.2 Oversight and Delegation

The second element of the model that appears to operate in the FDA's pharmaceutical regulation is in the idea that a leader can benefit when the agent prevents her from receiving information or from acting on it. One can view the game form as one of oversight or delegation, depending on which actors are assigned the key roles in the game. There is modest support for the idea that the public benefits from congressional oversight if Congress is the leader and the FDA as a whole is the agent. Meanwhile, Carpenter's account provides rather strong evidence of delegation premised on the reverse rationale if an FDA manager is the leader and scientists lower in the hierarchy play the role of agent.

Though it is intuitive to view the game involving Congress and the agency as one of delegation, it is also possible to interpret it as one of oversight. Delegation assumes an act of commitment, and in theory, at least, Congress can reassert its authority through new legislation (see Callander 2008, 124). Members of Congress can also attempt to influence FDA informally. Carpenter reports "numerous cases in which legislator applied pressure behind the scenes and lobbied for the approval of a particular drug" (337). In the related area of the FDA's monitoring activities (inspections and analyses of product samples), Shipan (2004) finds that the agency can sometimes be responsive to congressional committees. In addition, the leader always always accedes to FDA's proposal in the equilibria in Propositions 6 and 7, so lack of oversight in the model's definition does not follow from rare decisions by members of Congress to reverse FDA drug approval decisions. Overall, it seems plausible that the oversight game could be in effect.

Perhaps the most difficult aspect of the reverse rationale to establish for Congress is that some of its members can sufficiently understand the clinical trials to determine whether a new drug is safe and effective. Admittedly, it would rare for a member of Congress to have the skill to generate the studies that come from clinical trials. However, inability to create information need not imply that they cannot comprehend information. Even if they cannot personally dissect a study, they may be able to rely on trusted staff or outside scientists for their opinions about the studies. Empirically, there is mild support for the proposition that members of the relevant committees would feel confident in drawing their own conclusions about a drug's safety and efficacy. First, a few oversight hearings in the past have focused on particular drugs on the market (338–39). Second, the agency's "technical reputation" has come under attack in the past by AIDS activists (456), and it has more recently "suffered as top scientists have fled the agency or have complained publicly about being overruled or ignored" (748). Thus, even if their level of scientific knowledge is not as high as those of FDA officials, the gap in expertise may be small enough for some congresspersons to believe that they can interpret experimental data, either directly or through surrogates they trust more than FDA reviewers.

If legislators are exercising oversight in pharmaceutical regulation and some of them feel they have sufficient expertise to understand the evidence supporting a drug application, then the information filtering is clearly occurring and supports their continuing oversight, even though the primarily motive for this filtering is not to withhold information from Congress. Currently, the FDA discloses documents related to a NDA if the agency approves the medicine, but it does not release any materials related to the application if it rejects the prospective drug (Lurie and Zieve 2006, 89). This pattern of disclosure decisions corresponds to the equilibrium in Proposition 7. McGarity and Shapiro (1980) indicates that a primary reason that the FDA has cited for withholding this information is that it constitutes trade secrets (868–69). This work argues that the information should be disclosed, with an embargo on its use to support future applications to "ensure adequate research incentives" (884). However, the model suggests that the current withholding incentivizes research in its own way. Specifically, if data from all NDAs were released, firms might be able to expend less effort in research and rely on political pressure to have their drugs approved, anyway.

As for the working relationship between agency scientists and managers, Carpenter's account provides strong support for the notion that delegation prevents less adversarial players from making decisions based on information that they are capable of understanding. Through rulemaking, the FDA formally delegated authority as far down as the "directors and deputy division directors of the various drug review devisions" (484). Informally, true authority may lie in entry-level medical officers (see 483). Although entry-level officers may have more specialized expertise (id.), it is less plausible that their immediate supervisors and some higher officers lack sufficient expertise for an informed review. Instead, commitment of authority to entry-level officers might be motivated by the belief that they are the most adversarial agents within the FDA. Support for this notion comes from an activist who asserted that this kind of delegation insulates them not only from sponsoring firms, but also their overseers (490). Analogously, formal delegation to relatively junior directors might also be rationalized by the idea that they are more adversarial than more] senior officers, even if they are not as adversarial as entry-level reviewers.

5.3 Mitigating Capture

One important principle from the results on capture is that the use of an agent can protect against this phenomenon when bias-shifting is the key method since small amounts of bias-shifting directed at the leader will likely yield nothing for the outsider. If quasi-contracts with the leader are impossible, pharmaceutical firms should direct their efforts at influence toward the agent. If the FDA as a whole is perceived as the agent, this idea implies that they should target the agency rather than Congress. Unfortunately, it is difficult to compare the degree to which these companies have tried to exert their influence on each institution. However, the fact that committees have held hearings questioning the approval of particular drugs may reflect cases in which the sponsoring firm tried to influence the review process at the FDA rather than Congress. Such hearings would be consistent with a recognition that the agency should be adversarial, even though the committee members themselves might have difficulty rejecting a product with the same information.

Within the FDA, the idea that entry-level medical reviewers are more adversarial than directors was discussed as if these people's preferences were not susceptible to influence. However, formal delegation to division directors and informal delegation to initial reviewers could reflect an awareness by officers that they are susceptible to influence and should thus grant authority to less senior employees. Then, up to a point, the review process remains intact even if these officers are somewhat swayed, provided that the medical officers are not.

Because agency capture is a risk, a second principle is that ethics rules to keep the agent adversarial are better than additional disclosures in mitigating capture of the agent. This distinction is relevant to current policy discussions at the FDA. Specifically, the FDA has considered disclosing more information from NDAs, including those that are ultimately rejected (Asamoah and Sharfstein 2010). Although the standard tradeoff is between current knowledge about potential treatments and future innovation, the model suggests that releasing information about failed NDAs might empower firms and patient advocates who disagree with the rejection. They might be able to appeal to legislators or more senior agency officials to intervene with arguments that supposed risks are not as great as reviewer concluded, especially compared to the benefits

Instead, the more important challenge is keeping medical officers or the agency as a whole more adversarial. For the agency, the delegations already mentioned can be portrayed as one method of maintaining a high standard of proof. For individual officers, ethics rules could be helpful. Five FDA employees were convicted for accepting bribes to approve generic drugs in 1989 (Gibbons 1991), so officials are clearly subject to capture. Though bribery has always been illegal, ethics rules could prevent officers from being influenced to that degree or to a lesser extent. Pharmaceuticals constitute an area in which various ethical issues arise, so the model suggests that the ethics of government officials' policymaking should be as thoroughly examined as the ethics of other actors' decisions.

6 Conclusion

Based on information that is observable (albeit not contractible) and whose quality depends on an outside party's effort, the model presented in this paper offers what appears to be a new logic for a leader's use of an agent to prevent her from obtaining information that she could understand. It differs not only from the idea that an agent exists to apply his expertise and gather information that she cannot comprehend, but also from the notion that he exists to elicit information from a regulated party by virtue of policy preferences that are closer to that party's preferences than those of the leader. This rationale clearly contrasts with the expertise purpose, and it also differs from the standard information elicitation reason since the leader will want an agent who is more opposed to the outsider's preferences than she.

The "reverse" nature of this rationale continues into the analysis of capture. This model presents a plausible situation in which the purpose of an agent is not to facilitate capture by allowing interest groups to obtain favorable policy away from public scrutiny, but to reduce the incidence of regulatory capture by forcing interest groups to face agencies rather than political leaders. If leaders are subject to pressure that shifts their bias but can avoid quasi-contracts, then can and prefer to pass authority, or at least information-gathering, to an agent who can credibly threaten unfavorable policy can elicit higher quality information because he is naturally set against the interested party. Though the agent himself can be captured, incomplete disclosures that might seem to evince capture not only do not necessarily indicate influence by interest groups but instead may be essential for the leader to benefit from oversight.

The FDA's drug approval process arguably provides a concrete example in which this reverse rationale operates. The FDA and its scientific reviewers can generally be expected to be more adversarial than the general public and thereby induce more research from drug sponsors than if, in theory, each drug approval were decided according to popular will. There are even some hints, in the agency's delegations to junior-level employees and in congressional investigations of approved drugs, that actors might be aware of this dynamic. Since FDA officials do not have a monopoly on the relevant scientific knowledge, it is plausible that the current system usefully denies other actors either certain types information or the ability to act on that information. Even if specialized knowledge is one rationale for FDA regulation, the model at least indicates the desirability of relying on a regulator simply because she is adversarial and requires a high standard of proof—independently of any expertise advantage she may have over other decision-makers.

The overall logic of the model is that, when an outsider directly faces a leader who can understand it, it can generate information that just barely satisfies her so that she will base her policy on that evidence. Unable to commit to demand higher quality evidence, she can benefit by relying on an agent who is not satisfied except by much higher-quality evidence. Thus, the agent functions as a natural commitment device. Though committing to delegation or oversight may be a challenge, the example of FDA drug approvals suggests that it is possible and can yield public benefits. Therefore, more exploration of the presence and potential usefulness of this alternative informational rationale for agencies in the administrative state is warranted.

A Proofs of Numbered Results

Proof of Lemma 1 Using Bayes' rule, if $s = \tilde{x}^i$, then $x = s$ is always preferred to $x = 1-s$:

$$\frac{g(e)q_{\tilde{x}^i}b^i_{\tilde{x}^i}}{g(e)q_{\tilde{x}^i} + (1-g(e))q_{1-\tilde{x}^i}} \geq \frac{(1-g(e))q_{1-\tilde{x}^i}b^i_{1-\tilde{x}^i}}{g(e)q_{\tilde{x}^i} + (1-g(e))q_{1-\tilde{x}^i}} \tag{A.1}$$

since $q_{\tilde{x}^i}b^i_{\tilde{x}^i} > q_{1-\tilde{x}^i}b^i_{1-\tilde{x}^i}$ by definition and $g(e) \geq \frac{1}{2}$. On the other hand, if $s = 1 - \tilde{x}^i$, Bayes' rule implies $x = s$ is preferred to $x = 1 - s$ when

$$\frac{g(e)q_{1-\tilde{x}^i}b^i_{1-\tilde{x}^i}}{g(e)q_{1-\tilde{x}^i} + (1-g(e))q_{\tilde{x}^i}} \geq \frac{(1-g(e))q_{\tilde{x}^i}b^i_{\tilde{x}^i}}{g(e)q_{1-\tilde{x}^i} + (1-g(e))q_{\tilde{x}^i}}. \tag{A.2}$$

Algebra yields $e \geq e^i$ as defined in the Lemma to satisfy the inequality in the Lemma. ■

The next two lemmas are each used to prove more than one of the numbered results in the text and build on Lemma 1:

Lemma A.1. *The strategy for a decision-maker $i \in \{L, A\}$ includes the following compo-*

nents:

$$x^i = \begin{cases} \tilde{x}^i & \text{if } \mathring{e}_i < e^i \text{ or } \mathring{s}_i = \tilde{x}^i \\ 1 - \tilde{x}^i & \text{if } \mathring{e}_i > e^i \text{ and } \mathring{s}_i = 1 - \tilde{x}^i. \end{cases} \tag{A.3}$$

Proof. This result follows from Lemma 1 and the decision-maker's power to act on what information s/he observes. ■

Lemma A.2. *If the oversight game form applies and $x^{L*} = s$, $e^* \leq \max\{e^L, \hat{e}, e^A\}$. If an equilibrium exists in which $e^* = \max\{e^L, \hat{e}, e^A\}$ and $x^{L*} = s$, L cannot receive a higher payoff in equilibrium.*

Proof. Suppose $e > \max\{e^L, \hat{e}, e^A\}$. Since $e > e^L$, Lemma A.1 implies that, if A has the effort level and the signal, he can induce L to select $x^L = s, \forall s$, setting $\epsilon_L^A = \sigma_L^A = \delta, \forall s$. Because $e > e^A$, A prefers $x^L = s$ and will prefer to disclose information as described such that L selects $x^L = s$, provided that he can do so. Then, for any $e > \max\{e^L, \hat{e}, e^A\}$, R can receive $EU_f^R(e)$ with $\epsilon_A^R = \sigma_A^R = \delta$, so that, through A's disclosures, L selects $x^L = s$.

In any proposed equilibrium with $x^L = s$ and effort at some $\dot{e} > \max\{e^L, \hat{e}, e^A\}$, R would receive $EU_f^R(\dot{e})$. However, the above paragraph implies that R can increase his utility by selecting some $\ddot{e} \in (\max\{e^L, \hat{e}, e^A\}, \dot{e})$ and setting $\epsilon_A^R = \sigma_A^R = \delta$. Then R would receive $EU_f^R(\ddot{e}) > EU_f^R(e^*)$. This inequality holds because the concavity of $g(\cdot)$ and convexity of $c(\cdot)$ imply that $EU_f^R(e)$ decreases with effort $e \geq \hat{e}$. Thus, R is not best-responding if $x^L = s$ and $e > \max\{e^L, \hat{e}, e^A\}$.

The first statement is thus established. The second statement follows from the first statement, which implies that any equilibrium with $e > \max\{e^L, \hat{e}, e^A\}$ would not have $x^{L*} = s$ and so would yield L a weakly lower payoff, and from the fact that any equilibrium with $e^* < \max\{e^L, \hat{e}, e^A\}$ also would yield a weakly lower payoff. ■

Proof of Proposition 2 For any e, Lemma 1 implies that R would either select $x = s$ or $x = 1$ after both signals. The former yields $EU_f^R(e)$, while the latter yields $q_1 b_1^R - c(e)$. The assumptions on $g(e)$ and $c(e)$ imply that, ex ante, the R would like to set either $e = 0$ and $x = 1, \forall s$, or $e = \hat{e}$ and $x = s$. The definitions of types of researchers imply that an unmotivated researcher would prefer the former and a motivated researcher the latter. Ex post, Lemma 1 implies that it would select policy consistently with its ex ante preferences. Specifically, an unmotivated researcher has $q_1 b_1^R > \frac{1}{2}(q_0 b_0^R + q_1 b_1^R)$, and a motivated researcher has $EU_f^R(\hat{e}) > q_1 b_1^R$, so that $g(e)(q_0 b_0^R + q_1 b_1^R) > q_1 b_1^R$. ■

Proof of Proposition 3 If $\tilde{x}^L = 1$ and $\bar{e}_1$ does not exist, then R's highest possible payoff comes uniquely from $e = 0$ and $x^L = 1$, and, by Lemma A.1, it can assure this outcomewith $\epsilon_L^R = \delta$ after $e = 0$. Otherwise, $\bar{e}_{\tilde{x}^L}$ exists, and further analysis is needed.

First, $e \in (0, e^L)$ cannot occur in equilibrium. By Lemma 1, such an equilibrium would require $x^L = \tilde{x}^L$ for both signals. Then R would receive $q_{\tilde{x}^L} b_{\tilde{x}^L}^R - c(e)$, and it would deviate by setting $e = 0$ and $\epsilon_L^R = \delta$ to induce $x^L = \tilde{x}^L$ by Lemma A.1. R would prefer to similarly deviate from any proposed equilibrium in which $e > \bar{e}_{\tilde{x}^L}$ because it would receive less than $q_{\tilde{x}^L} b_{\tilde{x}^L}^R$, regardless of the policy selected.

If $e^L > \bar{e}_{\tilde{x}^L}$, then $e = 0$ is the only possible equilibrium effort level. Suppose, instead, that $e^L < \bar{e}_{\tilde{x}^L}$. For any $e \in [e^L, \bar{e}_{\tilde{x}^L}]$, an equilibrium would entail $x^L = s$. L would deviate from $x^L = 1 - s$ and from $x^L = 1 - \tilde{x}^L, \forall s$, while R would prefer to deviate from any equilibrium in which $x^L = \tilde{x}^L$ by setting $e = 0$ and $\epsilon_L^R = \delta$. If $\hat{e} \in (e^L, \bar{e}_{\tilde{x}^L})$, R would prefer to deviate from any equilibrium in which $e \neq \hat{e}$ by selecting $e = \hat{e}$ and $\epsilon_L^R = \sigma_L^R = \delta$. (Note that $\hat{e} = \bar{e}_{\tilde{x}^L}$ is ruled out since $EU_f^R(\hat{e}) \neq q_1 b_1^R$ by assumption in Footnote 1.) If, instead, $e^L \in [\hat{e}, \bar{e}_{\tilde{x}^L})$, it would prefer to deviate from any equilibrium in which $e = 0$ by selecting some $e \in (e^L, \bar{e}_{\tilde{x}^L}))$

and $\epsilon_L^R = \sigma_L^R = \delta$, and from any equilibrium in which $e \in (e^L, \bar{e}_{\tilde{x}^L}))$ by setting a lower e in that interval and $\epsilon_L^R = \sigma_L^R = \delta$; however, with no minimum value in the interval, the only permissible equilibrium effort level is e^L. Finally, if $e^L = \bar{e}_{\tilde{x}^L}$, then the two possible equilibrium effort levels are 0 and e^L.

In all these cases, an equilibrium can be constructed in which L sets $x^L = 1 - \tilde{x}^L$ only when $\mathring{e} \geq e^L$ and $\mathring{s} = 1 - \tilde{x}^L$. R receives $q_{\tilde{x}^L} b^R_{\tilde{x}^L}$ for $e = 0$ and $EU_f^R(e)$ for $e \geq e^L$. If $e^L \geq \bar{e}_{\tilde{x}^L}$, $q_{\tilde{x}^L} b^R_{\tilde{x}^L} > EU_f^R(e), \forall e > e^L$, so it optimizes by setting $e = 0$ and $\epsilon_L^R = \sigma_L^R = \delta$. If $e^L \leq \bar{e}_{\tilde{x}^L}$, $EU_f^R(e) > q_{\tilde{x}^L} b^R_{\tilde{x}^L}, \forall e \geq e^L$, and $\arg\max e \in [e^L, \bar{e}_{\tilde{x}^L}] EU_f^R(e) = \max\{e^L, \hat{e}\}$. In this case, she can ensure $x = s$ by setting $\epsilon_L^R = \sigma_L^R = \delta$ after $e = \max\{e^L, \hat{e}\}$. Lemma A.1 implies no deviation by L. ■

Proof of Proposition 5 (a): Proposition 3 implies that either $e^* = 0$ and $x^{L*} = 1$ when $\tilde{x}^L = 1$, $\tilde{x}^A = 0$, and $e^L \not\leq \bar{e}_1$ or $e^* = \max\{e^L, \hat{e}\}$ and $x^{L*} = s$ otherwise under administration, for a payoff that can be expressed as $EU_f^L(e)$ for some $e \in \{e^L, \hat{e}\}$. Under delegation Corollary 4 implies $e^* = e^A$ and $x^{A*} = s$ for a payoff of $g(e^A)(q_0 b_0^L + q_1 b_1^L) > \max_{e \in \{e^L, \hat{e}\}} EU_f^L(e)$.

(b): Proposition 3 implies $e^* = 0$ and $x^{L*} = 1$ under administration, for a payoff of $q_1 b_1^L$, whereas Corollary 4 implies $e^* = \hat{e}$ and $x^{A*} = s$ under delegation, for a payoff of $g(\hat{e})(q_0 b_0^L + q_1 b_1^L) > q_1 b_1$ since $\hat{e} > e^L$. ■

Proof of Proposition 6 That $e^* \leq e^L$ when $x^{L*} = s$ implies that L does not receive less than her reservation payoff. Also, she cannot exceed her reservation payoff of $q_{\tilde{x}^L} b^L_{\tilde{x}^L}$ if the same policy occurs after each signal or if $x^L = 1 - s, \forall s$. These facts and Lemma A.2 imply that the given equilibrium maximizes her payoff if it exists. Assume that the conditions stated in the proposition hold. It is sufficient to specify that L's strategy includes the rules in Lemma A.1 and $x^L = x^A$ when $\mathring{e}_L = \mathring{s}_L = \emptyset$. A's strategy can be $\epsilon_L^A = \sigma_L^A = \nu$ for any information it receives from R, with $x^A = 1 - \tilde{x}^A$ if any only if $\mathring{e}_A \geq e^A$ and $\mathring{s}_A = 1 - \tilde{x}^A$.

Faced with L and A's strategies, R will receive $q_{\tilde{x}^A} b^R_{\tilde{x}^A} - c(e)$ unless $e \geq e^A \geq e^L$, and it displays both items of information when $s = 1 - \tilde{x}^A$. When $\bar{e}_{\tilde{x}^A} \geq e^A$, it prefers to research at $e \in [e^A, \bar{e}_{\tilde{x}^A}]$, since $EU^R_f(e) \geq q_{\tilde{x}^A} b^R_{\tilde{x}^A}$ for these levels of effort. Since $EU^R_f(e)$ decreases with effort from its maximum at $\hat{e}$ (due to concavity of $g(\cdot)$ and $c(\cdot)$), its best response entails (1) $e = \hat{e}$ if $\hat{e} \geq e^A$ and $e = e^A$ otherwise to exert the least effort needed so that $x^L = s$; and (2), when $s = 1 - \tilde{x}^A$ to set $\epsilon^R_A = \sigma^R_A = \delta$. Given R's effort, A maximizes his utility if he induces $x^L = s$ from L. Since he does so by setting $x^A = s$ and $\epsilon^A_L = \sigma^A_L = \nu$, his strategy is a best response. Finally, since $e \geq e^L$ in equilibrium, L is best-responding: behind either proposal, with no other information, is a signal matching the proposal, supported by enough effort to persuade a leader inclined toward the opposite policy. ■

Proof of Proposition 7 As in Proposition 6, the facts that $e^* \leq e^L$ when $x^{L*} = s$, so that L does not receive less than her reservation payoff; that she cannot exceed her reservation payoff of $q_{\tilde{x}^L} b^L_{\tilde{x}^L}$ if the same policy occurs after each signal or if $x^L = 1 - s, \forall s$, combine with Lemma A.2 to imply that the given equilibrium maximizes her payoff if it exists. Assume that the conditions given in the proposition hold. L's strategy can be partially filled in with $x^L = 0$ when $x^A = 0$ and $\mathring{e}_L = \mathring{s}_L = \emptyset$, and $x^L = 1 - \tilde{x}^L$ when $\mathring{e}_L = e^L$ and $\mathring{s}_L = 1 - x^L$ (and is not unique beyond this specification). A's strategy can be $x^A = 1$ and $\epsilon^A_L = \sigma^A_L = \delta$ when $\mathring{e}_A \geq e^A$ and $\mathring{s}_A = 1$, and $x^A = 0$ and $\epsilon^A_L = \sigma^A_L = \nu$ otherwise.

Faced with these two players' strategies, R will receive $q_0 b^R_0$ unless $e \geq \max\{e^A, e^L\}$ and displays both items of information when $s = 1$. Because $\max\{e^A, e^L\} < \bar{e}_0$, it prefers to research at $e \in [\max\{e^A, e^L\}, \bar{e}_0]$, since $EU^R_f(e) \geq q_0 b^R_0$ for these levels of effort. Since $EU^R_f(e)$ decreases with effort from its maximum at $\hat{e}$ (due to concavity of $g(\cdot)$ and $c(\cdot)$), its best response entails (1) $e = \hat{e}$ if $\hat{e} \geq \max\{e^A, e^L\}$ and $e = \max\{e^A, e^L\}$ otherwise to exert the least effort needed so that $x^L = s$; and (2), when $s = 1$ to set $\epsilon_A^R = \sigma_A^R = \delta$. Given R's effort, A maximizes his utility if he induces $x^L = s$ from L. Since he does so by setting $x^A = 1$ and $\epsilon_L^A = \sigma_L^A = \delta$ when he observes the effort and $\mathring{s}_A = 1$ and by choosing $x^A = 1$ and $\epsilon_L^A = \sigma_L^A = \nu$ otherwise, his strategy is a best response. Finally, since $e \geq e^L$ in equilibrium, L is best-responding: by Lemma A.1 when she observes the effort level and signal and because $\mathring{e}_L = \mathring{s}_L = \emptyset$ implies a signal of 0 with $e = \max\{\hat{e}, e^A, e^L\}$, so that either type of leader prefers policy 0. ■

Proof of Theorem 8 (a): The conditions in Proposition 5 are a subset of the conditions under which the equilibrium in Proposition 6 exists: when $\max\{e^L, \hat{e}\} < e^A \leq \bar{e}_{\tilde{x}^A}$, or when $e^A \leq \hat{e}$ and $e^L < \hat{e}$ with $\tilde{x}^A = 0$, $\tilde{x}^L = 1$ and an unmotivated researcher $e^A \leq \bar{e}_{\tilde{x}^A}$ and $e^L \leq \max\{\hat{e}, e^A\}$. (In the latter case $\hat{e} < \bar{e}_0$ establishes that $e^A \leq \bar{e}_{\tilde{x}^A}$.) From Corollary 4 and Proposition 6, $e^* = \max\{\hat{e}, e^A\}$ and $x^* = s$, which implies that delegation and oversight yield the same payoff. Since Proposition 5 is refers to the conditions under which delegation outperforms administration, it follows that oversight outperforms administration by the same amount under these conditions. Also, since $e^* = \max\{e^L, \hat{e}, e^A\}$ in these cases, Lemma A.2 implies L cannot do any better.

(b)(i): Here, $e^* = e^L$ and $x^{L*} = s$ when $e^L \leq \bar{e}_{\tilde{x}^L}$ or $e^* = 0$ and $x^{L*} = \tilde{x}^L, \forall s$, under administration for her default payoff by Proposition 3, whereas $e^* \max\{\hat{e}, e^A\} < e^L$ and $x^* = s$ under delegation for a payoff less than her default payoff by Corollary 4. Under oversight, however, L can achieve her default payoff by setting $x^L = \tilde{x}$ unless $\mathring{e}_L \geq x^L$ and $\mathring{s}_L = 1 - \tilde{x}^L$ as she would under administration. For A it is sufficient to specify that $\epsilon_L^A = \sigma_L^A = \delta$ when $\mathring{e}_A \geq e^L$ and $\mathring{s}_A = 1 - \tilde{x}^L$. If $e < e^L$, there is nothing that R can disclose to A and have him relay to L so that she would prefer $x^L = 1-\tilde{x}^L$. Then $x^L = \tilde{x}^L, \forall s$ and R will receive $q_{\tilde{x}^L} b^R_{\tilde{x}^L} - c(e)$, which is maximized at $e = 0$. However, if $e \geq e^L$, he can disclose both items, which A will relay at least when $s = 1-\tilde{x}^L$, for $EU^R_f(e)$. Since $e^L > \hat{e}$, R prefers the lowest in this range, e^L. If $e^L \leq (>)\bar{e}_{\tilde{x}^L}$ then $g(e^L)(q_0 b^R_0 + q_1 b^R_1) - c(e^L) \geq (<) q_{\tilde{x}^L} b^R_{\tilde{x}^L}$. When $e = e^L$ (0), $x = s$ ($\tilde{x}^L, \forall s$). These are the same effort and policy choices as under administration, so L's payoff is the same if the equilibrium exists. A prefers to follow his strategy for $e \geq e^L$: when he observes the signal, his strategy leads L to select the same policy he would after each signal. If he does not observe the signal, any disclosure yields $x^L = 1 - \tilde{x}^L$. P will not defect, since on the equilibrium path she is selecting her preferred policy based on R and A's strategies and on Lemma 1. Also, Lemma A.2 implies that L a higher payoff since $e^L = \max\{e^L, \hat{e}, e^A\}$.

(b)(ii): Proposition 3 implies that $e = \hat{e}$ and $x = s$, for a payoff of $g(\hat{e})(q_0 b^L_0 + q_1 b^L_1)$, but Corollary 4 implies $e = 0$ and $x = \tilde{x}^L, \forall s$, for a payoff of $q_{\tilde{x}^L} b^L_{\tilde{x}^L} < g(\hat{e})(q_0 b^L_0 + q_1 b^L_1)$. However, L can recover her administration payoff under oversight by setting $x^L = 1 - \tilde{x}^L$ when $\mathring{e}_L = \mathring{s}_L = \emptyset$. Then R, which is motivated and has $\hat{e} \in (e^L, e^A)$, maximizes its payoff by setting $e = \hat{e}$, disclosing at least s when $s = \tilde{x}^L$, and setting $\epsilon^R_A = \sigma^R_A = \nu$ when $s = 1 - \tilde{x}^L$. Then $\sigma^A_L = \delta$ when $\mathring{s}_A = \tilde{x}^L$ to avoid having $x^L = 1 - \tilde{x}^L$, while A cannot disclose anything when $s = 1 - \tilde{x}^L$. Finally, L is best-responding: by Lemma A.1, $x^L = \tilde{x}^L$ should follow $\mathring{s}_L = \tilde{x}^L$, and when $\mathring{s}_L = \emptyset$, $s = 1 - \tilde{x}^L$, she prefers $x^L = 1 - \tilde{x}^L$ since $e = \hat{e} > e^L$. L cannot achieve a higher payoff with $x^L = x^A$ or $x^L = \tilde{x}^L$ when $\mathring{e}_L = \mathring{s}_L = \emptyset$. Lemma A.1 implies A's best response would entail $\epsilon^A_L = \sigma^A_L = \nu$ whenever $\mathring{e}_A < e^A$ and setting $x^A = \tilde{x}^L$ as needed to induce $x^L = \tilde{x}^L$. Since $e^A > \bar{e}_{\tilde{x}^L}$, R would not be best-responding if $e \geq e^A$ and would prefer to set $e = 0$ and accept $x^L = \tilde{x}^L$, which it can ensure with $e = 0$ and $\epsilon^R_A = \delta$. Meanwhile, if $\mathring{e}_A = \emptyset$, A would not be best-responding if $e < e^A$ and A did not act as needed to ensure

$x^L = \tilde{x}^L$. Since $x^L = \tilde{x}^L$ for $e < e^A$ for both signals in any equilibrium L receives her default payoff or less.

(b)(iii): Corollary 4 implies that, under delegation, L receives less than her default payoff under administration: When $e^A \not\leq \bar{e}_{\tilde{x}^A}$, $x^* = \tilde{x}^A, \forall s$, for $q_{\tilde{x}^A} b^L_{\tilde{x}^A} < q_{\tilde{x}^L} b^L_{\tilde{x}^L}$. When $e^A \leq \bar{e}_{\tilde{x}^A}$ but $\max\{\hat{e}, e^A\} < e^L$, $e^* = \max\{\hat{e}, e^A\}$ for $\max^L_{e \in \{\hat{e}, e^A\} EU^L_f(e) < q_{\tilde{x}^L} b_{\tilde{x}^L}}$. However, oversight allows L to recover her administration payoff. L's strategy be $x^L = \tilde{x}^A$ if and only if $\mathring{e}_L \geq e^L$ and $\mathring{s}_L = \tilde{x}^A$. Then R maximizes his payoff as follows: If $e^L \leq \bar{e}_{\tilde{x}^L}$, it sets $e = \max\{\hat{e}, e^L\}$ and $\epsilon^R_A = \sigma^R_A = \delta$. When $s = \tilde{x}^A$, A best-responds with $\epsilon^A_L = \sigma^A_L = \delta$ so that L will select $x^L = \tilde{x}^A$. If $e^L \not\leq \bar{e}_{\tilde{x}^L}$, R maximizes by setting $e = 0$, which leads to $x = \tilde{x}^L$ regardless of the disclosures. L is best responding since she selects the right policy after each equilibrium disclosure. With the same efforts and policy choices as in Proposition 3, L can obtain the same payoff under oversight as under administration.

L cannot receive a higher payoff with oversight when $\max\{\hat{e}, e^A\} < e^L$ by Lemma A.2. Suppose $\max\{\hat{e}, e^A\} \geq e^L$, but $e^A \not\leq \bar{e}_{\tilde{x}^A}$. First, whenever $\max\{e^A, e^L\} \leq \hat{e}$, and $e^A \not\leq \bar{e}_{\tilde{x}^A}$, $\tilde{x}^A = 1$ and R is unmotivated. Also, $\tilde{x}^L = 0$, in which case $e^* = \hat{e} = \max\{e^L, \hat{e}, e^A\}$ under administration. Then Lemma A.2 implies that L cannot do better.

Now suppose $\max\{e^L, \hat{e}\} \leq e^A$. For $\mathring{e}_L = \mathring{s}_L = \emptyset$, L must set $x_L = \tilde{x}^L$. Otherwise, Lemma A.1 implies A's best response would entail $\epsilon^A_L = \sigma^A_L = \nu$ whenever $\mathring{e}_A < e^A$ and setting $x^A = \tilde{x}^L$ as needed to induce $x^L = \tilde{x}^L$. If $\tilde{x}^L = 1$, $e^A > \bar{e}_0$, and R is not best-responding if $e \geq e^A$ since it would prefer to set $e = 0$. Meanwhile, if $\mathring{e}_A = \emptyset$, A would not be best-responding if $e < e^A$ and A did not act as needed to ensure $x^L = 0$. Thus, if $x_L = 0$ or $x^L = x^A$, having A and R best-respond entails $e = 0$ and $x^L = 0$, in which case L would defect with $x^L = 1$. If $\tilde{x}^L = 0$ and $x_L = 0$ or $x^L = x^A$ when $\mathring{e}_L = \mathring{s}_L = \emptyset$, then R can guarantee $q_1 b^R_1$ by setting $e = 0$ and $\epsilon^R_A = \delta$, since A will disclose and propose as needed to secure $x^L = 1$. If R is unmotivated, he would just set $e = 0$, and L would defect by setting $x^L = 0$. Even if R is motivated, it would defect if $e^A > \bar{e}_1$ since it could set $e = 0$, and A would defect if $e \leq \bar{e}_1$ and $x^L = s$ by always disclosing and proposing such that $x^L = 1$. That leaves $e \leq \bar{e}_1$ and policy not always matching the signal, so his best payoff is from $e = 0$, which would lead L to defect with $x^L = 0$. Thus, $x_L = \tilde{x}^L$ for $\mathring{e}_L = \mathring{s}_L = \emptyset$. Also, L cannot benefit if $x^L = \tilde{x}^A$ after $\mathring{e}_L = \emptyset$ and $\mathring{s}_L = \tilde{x}^A$. Then R would maximize with $e = \hat{e}$, $\epsilon_A^R = \nu$, $\sigma_A^R(\hat{e}, \tilde{x}^A) = \delta$, and $\sigma_A^R(\hat{e}, \tilde{x}^L) = \nu$.

Suppose $\max\{e^L, \hat{e}\} \leq e^A$ and $e^L \not\leq \bar{e}_{\tilde{x}^L}$. If $\tilde{x}^L = 0$, then R would never set $e \geq e^L > \bar{e}_0$ since it would prefer $e = 0$ followed by any policy. Then any equilibrium with $e < e^L$ cannot yield more than $q_0 b_0^L$, her payoff under administration. If $\tilde{x}^L = 1$, then when $\mathring{e}_L = \mathring{s}_L = \emptyset$, $x^L = 1$. R can guarantee $q_1 b_1^R$ with $e = 0$ and $\epsilon_L^A = \sigma_L^A = \nu$. There is no equilibrium in which R's utility is higher and in which L would be best-responding since $e^L \not\leq \bar{e}_1$. So it must be that $e^* = 0$ and $x^{L*} = 1, \forall s$, which is the same result as under administration.

If $\max\{e^L, \hat{e}\} \leq e^A$, $e^L \leq \bar{e}_{\tilde{x}^L}$ and $\bar{e}_{\tilde{x}}^A \not\leq e^A$, then L may be able to exceed her payoff from administration and delegation with oversight. This will not happen when $\tilde{x}^A = 0$, $\tilde{x}^L = 1$, and R is unmotivated, in which case R maximizes his payoff with $e = 0$ and $\epsilon_L^A = \sigma_L^A = \nu$. Also, if $\tilde{x}^A = 1$, $\tilde{x}^L = 0$, and $b_0^R < 0$, in which case L must set $x^L = 0$ when $\mathring{e}_L > e^L$ and $\mathring{s}_L = \emptyset$. Otherwise, R always prefers $x^L = 1$ ex post and would set e or slightly above e^L and $\epsilon_A^R = \delta$ and $\sigma_A^R = \nu$, and A would set $\epsilon_L^A = \delta$ and $x^A = 1$ as necessary so that $x^L = 1$. With $x^L = 0$ when $\mathring{e}_L > e^L$ and $\mathring{s}_L = \emptyset$, R can only aim for $x^L = s$, in which case it selects $e = \max\{\hat{e}, e^L\}$, which would occur under administration or delegation.

(c) Under these conditions, Proposition 3 implies a payoff of $EU_f^L(\hat{e}) > q_0 b_0^L$ under administration, while Corollary 4 implies a payoff of $q_0 b_0^L$ under delegation. Oversight cannot yield L more than delegation. L's strategy after $\mathring{e}_L = \mathring{s}_L = \emptyset$ cannot be $x^L = 1$, or else R

would maximize its utility by setting $e = 0$ and withholding both items of information. If instead, the strategy for $\mathring{e}_L = \mathring{s}_L = \emptyset$ is $x^L = 0$ or $x^L = x^A$, A can ensure that $x^L = 0$ for any $\mathring{e}_A < e^A$. Since $e^A > \bar{e}_{\tilde{x}^L}$, R would not be best-responding if $e \geq e^A$ and would prefer to set $e = 0$ and accept $x^L = \tilde{x}^L$, which it can ensure with $e = 0$ and $\epsilon_A^R = \delta$. Meanwhile, if $\mathring{e}_A = \emptyset$, A would not be best-responding if $e < e^A$ and A did not act as needed to ensure $x^L = \tilde{x}^L$. Since $x^L = \tilde{x}^L$ for $e < e^A$ for both signals in any equilibrium L receives no more than her default payoff, which is less than $EU_f^L(\hat{e})$ under administration.

(d)(i): Among cases in this set, Proposition 3 and Corollary 4 imply that $e^* = \hat{e}$ under both administration and delegation. Then Lemma A.2 implies that there cannot be any equlibria with $e^* > \hat{e}$ with $x^{L*} = s$, and any other equilibria would yield L no more than $EU_f^L(\hat{e})$. The exception is when $\tilde{x}^L = \tilde{x}^A = 1$, and R is unmotivated, in which case $e^* = 0$ in both modes. Then oversight adds nothing since R can set $e = 0$ and $\epsilon_A^R = \delta$, in which case A would set $\epsilon_L^A = \delta$ as needed to ensure $x^L = 1$.

(d)(ii): Proposition 3 implies $e^* = e^L$ or $e^* = 0$, yielding $q_{\tilde{x}^L} b_{\tilde{x}^L}$, and Corollary 4 entails $e^* = 0$ for the same payoff. L's payoff is the same, but not more, under oversight. If $\tilde{x}^L = 0$, then $\mathring{e}_L = \mathring{s}_L = \emptyset$ should not lead to $x^L = 1$. Otherwise, R can guarantee $q_1 b_1^R$ through $x^L = 1$ with $e = 0$ and $\epsilon_L^A = \sigma_L^A = \nu$, R would not be best-responding if $e \geq e^A > \bar{e}_0$, and A would not be best-responding if R received more than $q_1 b_1^R$. Thus, if $x^L = 1$ when $\tilde{x}^L = 0$ and $\mathring{e}_L = \mathring{s}_L = \emptyset$, $e = 0$ and $x^L = 1$, causing L to defect with $x^L = 0$. Then if $x^L = 0$ or $x^L = x^A$ when $\mathring{e}_L = \mathring{s}_L = \emptyset$, A will set $\epsilon_L^A = \sigma_L^A = \emptyset$ and set $x^A = 0$ as needed to ensure $x^L = 0$ when $\mathring{e} < e^A$. Again, R would not be best-responding if $e \geq e^A > \bar{e}_0$. Then in any possible equilibrium A must induce $x^L = 0$ if $\mathring{e} = \emptyset$. The result is $e = 0$ and $x^L = 0$, which yields the same payoff as under administration or delegation. If $\tilde{x}^L = 1$ and R is unmotivated, $e = 0$ with $\epsilon_A^R = \epsilon_L^A = \delta$ to ensure that $x^L = 1$. Even if R is motivated,

$e^A > \bar{e}_1$, a best response by implies that there is no equilibrium with $e > \bar{e}_1$ (or else R is not best-responding) or with $e \leq \bar{e}_1$ and $x^L = s$ (or else A is not best-responding). Then $e = 0$ and $x^L = 1$, which yields L the same payoff as under administration or delegation. ■

Proof of Proposition 9 If oversight is available, Theorem 8 implies that oversight dominates delegation and that oversight dominates administration when an agent with $\tilde{x}^L = 0$ and $e^A \leq \bar{e}_0$ is involved. Most of the equilibria in which L benefits from oversight involve $\max\{e^L, \hat{e}\} = e^A = e^* \leq \bar{e}_0$ and $x^{L*} = s$. Among such equilibria, she prefers the greatest value of e^A because her payoff, represented by Equation (1), increases with e. The exception is those alluded to Proposition 5(b), but even then, L would prefer a higher value of e^A to achieve an equilibrium in the previous category. If oversight is not available, Theorem 8 neither delegation nor administration dominates among agents with $\tilde{x}^L = 0$ and $e^A \leq \bar{e}_0$. However, she maximizes by choosing an agent with the greatest value of e^A. Under administration, the e^A is not relevant. Under delegation, however, Corollary 4 implies that her utility weakly increases with e^A provided that $e^A \leq \bar{e}_0$. ■

Proof of Proposition 10 Theorem 8 implies that, when $e^L \leq e^A < \bar{e}_0$ and $\tilde{x}^A = 0$, only cases (a) and (d)(i) apply, in which case L prefers oversight or delegation to administration. As long as $e^L \leq e^A$ after stautory bias-shifting, she will still prefer oversight or delegation. Meanwhile, e^A and $\hat{e}$ not affected by statutory bias-shifting, so the equilibria effort and policy selections in Corollary 4 and Propositions 6 and 7 are unaffected. ■

Proof of Proposition 11 L's payoff exceeds her default only when $e > e^L$ and $x^* = s$. Her payoff would be represented by Equation (1). Then her maximum possible payoff is $EU_f^L(\bar{e}_0)$ given her choice of agents. Proposition 5 implies that L would continue to at least prefer oversight or delegation as long as $\max\{e^L, \hat{e}\} \leq e^A$. Thus, if L only receives $EU_f^L(\bar{e}_0) - \max_{e\in\{e^L,\hat{e}\}} EU_f^L(e) \equiv \Delta_1 V^L$, she would still be willing to use delegation or oversight with an agent who has $e^A = \max\{e^L, \hat{e}\}$.

Proposition 3 implies that $\max_{e\in\{e^L,\hat{e}\}} EU_f^L(e)$ is her payoff under administration if $e^L \leq \bar{e}_{\tilde{x}^L}$. Since $e^L < \bar{e}_0$, if $\tilde{x}^L = 0$, then $\Delta_2 V^L = \Delta_1 V^L$. The same is true if $\tilde{x}^L = 1$ and R is motivated. If $e^L \leq \bar{e}_1$, $e^* = \max\{e^L, \hat{e}\}$ under any game form. If $e^L > \bar{e}_1 > \hat{e}$, then her payoff under administration is $q_1 b_1^L = EU_f^L(e^L)$. If, however, $\tilde{x}^L = 1$ and R is unmotivated and $\hat{e} > e^L$, then her payoff under delegation or oversight is $EU_f^L(\hat{e})$, whereas her administration payoff is $q_1 b_1^L < EU_f^L(\hat{e})$. Here, $\Delta_2 V^L = EU_f^L(\bar{e}_0) - q_1 b_1^L < \Delta_1 V^L$ is necessary for L to select administration.

If R is motivated, R prefers $\hat{e}$ but does not receive it if $e^L > \hat{e}$. Then L needs $\Delta_3 V^L = EU_f^L(\bar{e}_0) - EU_f^L(\hat{e}) > EU_f^L(\bar{e}_0) - q_1 b_1^L = \Delta_2 V^L$ to accept R's referred outcome. If R is unmotivated, R prefers $e = 0$ and $x^L = 1$ but does not receive it if $\tilde{x}^L = 0$. Then $\Delta_3 V^L = EU_f^L(\bar{e}_0) - q_1 b_1^L > \max\{EU_f^L(\hat{e}), q_0 b_0^L\}$. ■

Proof of Proposition 12 Following Corollary 4 and Proposition 6 or 7, A's utility under delegation or oversight in the two respective cases is $EU_f^A(e^A)$ and $EU_f^A(\check{e}^A)$, for $\Delta V^A = EU_f^A(e^A) - EU_f^A(\check{e}^A)$. With the definition of B_0^A and A's standard of proof, the two respective biases are $B_0^A = 2g(e^A) - 1$ and $\check{B}_0^A = 2g(\check{e}^A) - 1$, for $\Delta B_0^A = 2(g(e^A) - g(\check{e}^A))$. ■

References

Aberbach, Joel D. 1990. *Keeping a Watchful Eye: The Politics of Congressional Oversight.* Washington, DC: Brookings Institution.

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

Ambrus, Attila, Eduardo M. Azevedo, and Yuichiro Kamada. 2013. "Hierarchical Cheap Talk." *Theoretical Economics* 8 (January): 233–61.

Asamoah, Afia K., and Joshua M. Sharfstein. 2010. "Transparency at the Food and Drug Administration." *New England Journal of Medicine* 362 (June 24): 2341–43.

Beermann, Jack M. 2006. "Congressional Administration." *San Diego Law Review* 43 (February-March): 61–158.

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

Bertelli, Anthony, and Sven E. Feldmann. 2007. "Strategic Appointments." *Journal of Public Administration Research and Theory* 17 (January): 19–38.

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

Callander, Steven. 2008. "A Theory of Policy Expertise." *Quarterly Journal of Political Science* 3 (July): 123–40.

Carpenter, Daniel. 2010. *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton: Princeton University Press.

Carpenter, Daniel. 2013. "Detecting and Measuring Capture." In *Preventing Regulatory Capture: Special Interest Influence, and How to Limit It*, ed. Daniel Carpenter and David Moss. New York: Cambridge University Press. Forthcoming.

Carpenter, Daniel P. 2004. "Protection without Capture: Product Approval by a Politically Responsive, Learning regulator." *American Political Science Review* 98 (November): 613–31.

Coglianese, Cary. 2009. "The Transparency President? The Obama Administration and Open Government." *Governance* 22 (October): 529–44.

Coglianese, Cary, Richard Zeckhauser, and Edward Parson. 2004. "Seeking Truth for Power: Informational Strategy and Regulatory Policymaking." *Minnesota Law Review* 89 (December): 277–341.

Dal Bó, Ernesto. 2006. "Regulatory Capture: A Review." *Oxford Review of Economic Policy* 22 (Summer): 203–25.

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

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

Gailmard, Sean, and John W. Patty. 2012. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15: 353–77.

Gailmard, Sean, and John W. Patty. 2013a. *Learning While Governing: Information, Accountability, and Executive Branch Institutions*. Chicago: University of Chicago Press.

Gailmard, Sean, and John W. Patty. 2013b. "Stovepiping." *Journal of Theoretical Politics*. Forthcoming.

Gibbons, Ann. 1991. "Can David Kessler revive the FDA?" *Science* 252 (April 12): 200–03.

Graham, John D., Paul R. Noe, and Elizabeth L. Branch. 2006. "Managing the Regulatory State: The Experience of the Bush Administration." *Fordham Urban Law Journal* 33 (May): 953–1002.

Harris, Gardiner. 2008. "F.D.A. Scientists Accuse Agency Officials of Misconduct." *New York Times* (November 18): A15.

Huber, John D., and Charles R. Shipan. 2002. *Deliberate Discretion? The Institutional Foundations of Bureaucratic Autonomy*. New York: Cambridge University Press.

Ivanov, Maxim. 2010. "Communication via a Strategic Mediator." *Journal of Economic Theory* 145 (March): 869–84.

Kreps, David M. 1990. "Corporate Culture and Economic Theory." In *Perspectives on Positive Political Economy*, ed. James E. Alt and Kenneth A. Shepsle. Cambridge, UK: Cambridge University Press.

Kwak, James. 2013. "Cultural Capture and the Financial Crisis." In *Preventing Regulatory Capture: Special Interest Influence, and How to Limit It*, ed. Daniel Carpenter and David Moss. New York: Cambridge University Press. Forthcoming.

Laffont, Jean-Jacques, and Jean Tirole. 1991. "The Politics of Government Decision-making: A Theory of Regulatory Capture." *Quarterly Journal of Economics* 106 (November): 1089–1127.

Landis, James M.C. 1938. *The Administrative Process*. New Haven, CT: Yale University Press.

Levine, Michael E., and Jennifer L. Forrence. 1990. "Regulatory Capture, Public Interest, and the Public Agenda: Toward a Synthesis." *Journal of Law, Economics, and Organization* 6 (April): 167–98.

Lurie, Peter, and Allison Zieve. 2006. "Sometimes the Silence Can Be Like the Thunder: Access to Pharmaceutical Data at the FDA." *Law and Contemporary Problems* 69 (Summer): 85–97.

McCarty, Nolan. 2013. "Complexity, Capacity, and Capture." In *Preventing Regulatory Capture: Special Interest Influence, and How to Limit It*, ed. Daniel Carpenter and David Moss. New York: Cambridge University Press. Forthcoming.

McCubbins, Mathew D., and Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols versus Fire Alarms." *American Journal of Political Science* 28 (February): 165–79.

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

McGarity, Thomas O., and Sidney A. Shapiro. 1980. "The Trade Secret Status of Health and Safety Testing Information: Reforming Agency Disclosure Policies." *Harvard Law Review* (March): 837–88.

Neill, Katharine A., and John C. Morris. 2012. "A Tangled Web of Principals and Agents: Examining the Deepwater Horizon Oil Spill through a Principal–Agent Lens." *Politics and Policy* 40 (August): 629–56.

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

Rogoff, Kenneth. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (November): 1169–89.

Shipan, Charles R. 2004. "Regulatory Regimes, Agency Actions, and the Conditional Nature of Congressional Influence." *American Political Science Review* 98 (August): 467–80.

Thurber, James A. 2011. "The Contemporary Presidency: Changing the Way Washington Works? Assessing President Obama's Battle with Lobbyists." *Presidential Studies Quarterly* 41 (June): 358–74.

Ting, Michael M. 2008. "Whistleblowing." *American Political Science Review* 102 (May): 249–67.

Tirole, Jean. 1986. "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations." *Journal of Law, Economics, and Organizations* 2 (Autumn): 181–214.

Urbina, Ian. 2010, May 14. "U.S. Said to Allow Drilling Without Needed Permits." *New York Times*: A1.

Weber, Max. 1978. *Economy and Society*. Trans. Ephraim Fischoff et al. Berkeley, CA: University of California Press.