



Learning from Oversight: Fire Alarms and Police Patrols Reconstructed

Arthur Lupia; Mathew D. McCubbins

Journal of Law, Economics, & Organization, Vol. 10, No. 1. (Apr., 1994), pp. 96-125.

Stable URL:

<http://links.jstor.org/sici?sici=8756-6222%28199404%2910%3A1%3C96%3ALFOFAA%3E2.0.CO%3B2-T>

Journal of Law, Economics, & Organization is currently published by Oxford University Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/oup.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Learning From Oversight: Fire Alarms and Police Patrols Reconstructed

Arthur Lupia and Mathew D. McCubbins
University of California, San Diego

While the delegation of policy-making authority from legislators to bureaucrats is ubiquitous in modern democracies, there is considerable disagreement about the consequences of this type of delegation. Some scholars point to the fact that bureaucrats tend to have policy-relevant expertise, assume that bureaucrats will use their expertise to systematically mislead legislators, and conclude that delegation and abdication are equivalent. Other scholars point to the extensive use of legislative oversight, assume that oversight is sufficient to abate the problems associated with bureaucratic expertise, and conclude that delegation produces more effective governance. We depart from previous scholarship by developing models of delegation and oversight that allow us to derive, rather than assume, conditions under which legislators can adapt successfully to bureaucratic expertise. With these conditions in hand, we identify conditions under which delegation to the bureaucracy produces more effective governance and conditions under which delegation and abdication are equivalent.

1. Introduction

The delegation of policy-making authority is ubiquitous in modern democracies. In this article, we focus on a legislature's decision to delegate such authority to the bureaucracy. The potential advantage of this type of delegation is that bureaucrats may have expertise that legislators desire. As a result, legislators may be able to rely on the bureaucracy to formulate policies that they themselves would have formulated if they had spent the time and resources necessary to acquire the bureaucracy's level of expertise. The potential drawback of delegation is that a bureaucrat who possesses both expertise and policy-making authority can also take actions that make legislators worse off than if they had never delegated. In this case, the act of delegation is equivalent to abdication, wherein the will of the people, as expressed by the people's elected representatives, neither constrains nor motivates the formulation of public policy.

We thank Will Heller, Jonathan Katz, Keith Krehbiel, Roger Noll, Brian Sala, and Michael Thies for helpful comments. Professor McCubbins acknowledges the support of the National Science Foundation grant no. SES 9022882 and the support of the Ford Foundation.

This article is about the consequences of delegation. In choosing this topic, we join a vigorous debate. On one side of this debate are scholars who argue that the presence of bureaucratic expertise is sufficient to make delegation and abdication equivalent (Lowi, 1979; Huntington, 1965; Niskanen, 1971; and Weber, in Gerth and Mills, 1946). On the other side of this debate are scholars who argue that legislative oversight, budgeting, and administrative procedure are sufficient to mitigate the potentially deleterious consequences of delegation (see, e.g., Harris, 1964; Fiorina, 1977; Weingast and Moran, 1983; McCubbins, Noll, and Weingast, 1987; and Kiewiet and McCubbins, 1991).

While both sides of this debate agree that legislators' ability to adapt to bureaucratic expertise is the key to understanding delegation's consequences, neither set of scholars has demonstrated the sufficiency they claim. For instance, scholars who argue that delegation is manageable assume, but do not demonstrate, that legislators learn from oversight and administrative procedure. By contrast, scholars who argue that delegation is equivalent to abdication typically reject the possibility that legislators can learn about bureaucratic expertise.

We depart from previous scholarship by deriving, as opposed to assuming, conditions under which legislators can adapt successfully to bureaucratic expertise. With these conditions in hand, we then render a more reasoned judgment about the consequences of delegation. That is, we identify conditions under which delegation to the bureaucracy will be beneficial for legislators and conditions under which delegation and abdication are equivalent.

We conduct our analysis by focusing our attention on a primary means by which legislators overcome bureaucratic expertise: oversight. The literature on legislative oversight has come to focus on two broad categories: police-patrol oversight and fire-alarm oversight. [For a survey, see Ogul and Rockman (1990). The original definitions and arguments are due to McCubbins and Schwartz (1984).] Police-patrol oversight is the centralized and direct approach to uncovering hidden knowledge and is what most people think about when they discuss the oversight function of legislatures. An example of police-patrol oversight is a legislator who personally conducts an audit of agency activity. Fire-alarm oversight, on the other hand, is relatively passive, indirect, and decentralized. Legislators who conduct fire-alarm oversight establish "a system of rules, procedures, and informal practices that enable [interested third parties] to examine administrative decisions . . . [and] to seek remedies from agencies, courts, and [the legislature] itself" (McCubbins and Schwartz, 1984:166).

A prerequisite to understanding the consequences of delegation, then, is to determine the conditions under which legislators can learn from police-patrol or fire-alarm oversight. By definition, police-patrol oversight is likely to be an effective way for legislators to track bureaucratic actions. However, it is also likely to be very costly in terms of the time and resources needed to conduct it. Because of its high costs, several scholars have argued that police-patrol oversight cannot be conducted in sufficient amounts to prevent delegation

from lapsing into abdication (Lowi, 1979; Dodd and Schott, 1979; Huntington, 1965).

Fire-alarm oversight, by contrast, has cost advantages for legislators relative to police-patrol oversight. Under a fire-alarm system, interested third parties, such as constituents or interest groups, bear the lion's share of the costs of learning about bureaucratic activities. It is this efficiency that McCubbins and Schwartz argue forms the basis of a legislative preference for fire-alarm oversight. The validity of the McCubbins-Schwartz and other institutionalist arguments, however, depends upon the crucial assumption that legislators can learn from fire-alarm activity. Unfortunately, there is neither justification nor evidence for this assumption. People who are given the opportunity to play the role of fire alarm are also given an opportunity to benefit from their ability to deceive legislators. In short, false alarms are always possible.

If police-patrol oversight is costly and fire alarms are disingenuous, then legislators may be unable to adapt to the potentially deleterious consequences of bureaucratic expertise, and it may well be that delegation and abdication are equivalent. We argue, in contrast, that even the simultaneous existence of bureaucratic expertise, the high cost of police-patrol oversight, and the possibility of perfidious fire alarms is not sufficient to equate delegation with abdication. We support this argument by identifying three general characteristics of oversight that allow legislators to learn about bureaucratic expertise. We call each characteristic a condition for learning.

The first condition for learning is a legislator's ability to observe costly action by a bureaucratic agent who possesses expertise. The logic here is basically the same as the idea behind the old adage "actions speak louder than words." In general, if we observe an informed individual taking actions that are costly (e.g., the action involves effort or precludes other actions being taken), then we can learn something about how the individual's preferences or beliefs about the world compare to our own (since the individual would not have paid the costs had there not been an expected net benefit from the actions taken).

The second condition for learning is the existence of a cost associated with making particular statements. One example of this type of cost is a penalty for lying, where it costs the information provider more to lie than to tell the truth. The penalty for lying can be thought of as the potential loss in a valued reputation for honesty. A penalty for lying works by affecting an information provider's benefit-cost calculation. The information provider benefits from lying when doing so increases the likelihood that we (as information receivers) take the action that is preferred by the information provider. It will be worthwhile for an information provider to lie only if the expected benefit of lying outweighs the expected cost. If the penalty is large enough to make certain statements unprofitable as lies, then we can infer that such statements, when made, must be true. Thus, a penalty for lying allows us to learn from information providers, even when they have an incentive to lie.

The third condition for learning is a similarity of preferences over outcomes. If outcomes that are good for one person are also good for another, and if bad outcomes for one are also bad for another, then neither person will have an incentive to lead the other to take an action that produces an outcome that is bad for either. The dynamic effect of preference similarity can be stated as follows: the more likely it is that the information provider shares our preferences over outcomes, the greater is the likelihood that the information provider will truthfully reveal what he knows and the greater, therefore, is the weight that we should assign to the information provider's claim being true.

When these conditions for learning exist, and are sufficiently informative, then legislators will be able to avoid the potentially deleterious consequences of delegation. Such avoidance is possible because legislators are better able to distinguish between beneficial and detrimental bureaucratic actions and because bureaucrats, who anticipate the legislators' enhanced abilities, are more likely to believe that they will be subject to legislative discipline for taking actions that are contrary to legislative interests. In other words, when legislators can create the conditions for learning, they can create incentives for bureaucrats to act in accordance with legislative interests. In the absence of these conditions for learning, the prospect of such legislative discipline disappears and delegation has the same consequences as outright abdication.

The remainder of the article proceeds as follows. In Section 2, we explain the basic relationship between delegation and its consequences. In Section 3, we describe the structure and findings of a model of legislative decision-making that includes delegation and the possibility of police-patrol oversight. In Section 4, we present a model that focuses on the relationship between fire-alarm oversight and the consequences of delegation. In Section 5, we conclude with a review of the ways in which legislators can learn from oversight and the extent to which their ability to learn affects the consequences of delegation. The two Appendices contain the technical foundations of our models of oversight and the derivation of our results.

2. The Potential Consequences of Delegation

What are the consequences of legislative delegation to the bureaucracy? In modern parlance, delegation is seen as a game between a principal—in our case, Congress and its members—and an agent—a bureaucratic agency. Delegation occurs when the principal assigns to the agent the responsibility to take certain types of actions.

There are two key issues involved in determining the consequences of delegation. The first is *agenda control*. Does the agent make recommendations to the principal, who then decides on a course of action, or does the agent present the principal with a *fait accompli*? As Congress has always retained the right to overturn agency decisions, we believe that the agenda control present in congressional delegations to the bureaucracy resembles the former type of agenda control rather than the latter. We thus model delegation as a sequential game that begins when an agent offers a take-it-or-leave-it

policy proposal to the principal and ends when the principal either accepts it (perhaps by doing nothing) or rejects it in favor of the status quo policy (i.e., the policy on that issue is left as it was before the agent's offer).

The second important determinant of delegation's consequences is the *extent of agency expertise* (Lupia and McCubbins, 1993; and McCubbins and Noble, 1993). To see the effect of expertise, consider two nonoverlapping cases. In the first case, the agent has no expertise about the policy choice (i.e., the principal and the agent have the same information). Under these circumstances, the legislative principal is as able as the bureaucratic agent to distinguish bureaucratic proposals that are beneficial (for the principal) from bureaucratic proposals that are deleterious. The consequence of delegation in this case is that the agent will make only those proposals that it believes will make both the agent and the principal better off than they would have been under the status quo.

In the second case, the agent has expertise (i.e., the principal is relatively uninformed about the consequences of accepting the agent's proposal). If the principal believes that the agent is relatively likely to make an offer that is better for the principal than the status quo, then the principal should accept the offer. Otherwise, the principal should stick with the status quo. The consequences of delegation for the principal, then, are beneficial only when the agent prefers policies similar to those preferred by the principal. If, by contrast, the agent prefers policies different from those preferred by the principal, but, ironically, the principal believes otherwise, the principal may accept a proposal that should be rejected. The consequence of delegation in the case of agency expertise, then, may be that the outcome is worse for the principal than the existing policy.

We believe that bureaucrats generally possess knowledge about their policy choices that legislators do not. It is thus important to identify conditions under which legislators can circumvent the detrimental consequences of bureaucratic expertise. It is widely recognized that oversight is the cornerstone of legislative efforts to efface bureaucratic expertise. We examine, in turn, the utility of the two types of oversight—police patrols and fire alarms—discussed in the literature.

3. A Model of Delegation and Police-Patrol Oversight

We model delegation with police-patrol oversight as a multistage, single-shot game between two players: a bureaucratic *agent* and a legislative *principal*.¹ The agent's task is to propose changes in policy. The legislative principal's task is to choose between the agent's policy proposal (or offer), denoted *o*, or the existing policy of the government, called the status quo and denoted *sq*.

1. For simplicity, we treat the legislature as an individual actor. In effect, we assume that individual legislator preferences and the existing legislative institution have already interacted to produce a single legislative preference ordering over the possible alternatives. Our "legislative interest," then, is a generalization of what is often called the "median legislator's (or median committee member's) preferences."

Both o and sq are represented as points on the unit interval $[0,1]$. Also represented on this interval is an *ideal point* for each player. We assume that both players have single-peaked preferences, which means that neither player strictly prefers an outcome that is relatively far from its ideal point to an outcome that is relatively close to it.² Thus, each player's objective is to obtain the policy, o or sq , that is closest to its own ideal point. Unless stated otherwise, we assume that the value of each parameter in the game is common knowledge.

The exception to the common-knowledge assumption is that the location of the agent's ideal point may be known only by the agent (it is the agent's private information). Since the agent's objective is to make a proposal that yields an outcome that is as close as possible to its own ideal point, the assumption that the agent's ideal point is private information implies that the location of o , the agent's policy proposal, may also be private information. Thus, while the agent knows whether o or sq is closer to the principal's ideal point, the principal may be uncertain. We represent police-patrol oversight as a decision by the principal to take a costly action that informs it about the location of o .³

We model the principal-agent relationship as a series of four events, depicted in Figure 1. First, the values of several parameters are determined exogenous to the play of the game. These values are: the location of the status quo (sq), the location of the principal's and agent's ideal points, the magnitude of all costs, and the prior information about the location of the agent's ideal point. Second, the agent decides whether or not to issue a proposal. If the agent makes no proposal, the game ends with each player's utility determined by the distance between their own ideal point and the status quo. If the agent makes a proposal, it pays a cost, c_a , makes an offer, o , and forces the principal to choose between o and sq . Third, the principal decides whether or not to engage in police-patrol oversight. Police-patrol oversight is costly; the principal pays (cost of monitoring) $c_m \geq 0$, but in return learns the location of the agent's offer, o .⁴ The principal then ends the game by either accepting the agent's policy offer, o , or rejecting it in favor of the status quo, sq .⁵

2. Specifically, the principal's utility function is $-|X - P|$ and the agent's utility function is $-|X - A|$, where P is the principal's ideal point, A is the agent's ideal point, and $X \in \{o, sq\}$. Our results are dependent on neither unidimensionality nor the utility functions stated; these particular assumptions are used to explain the logic of the model in relatively simple terms. It is easy to verify that our results are robust to either the assumption that o and sq are points in n -dimensional space (n finite) or the assumption that utility functions are continuous over $[0,1]$ and single-peaked.

3. Though McCubbins and Schwartz identified several characteristics of police-patrol oversight, we concentrate here on just one: that police-patrol oversight is an active effort by legislators to obtain sufficient information about the bureaucracy's actions so as to remedy bureaucratic missteps.

4. Although we assume that the principal knows this cost, our findings easily generalize to the case where paying c_m buys the principal knowledge of o with a probability about which the players have common priors.

5. In many legislative environments there is an asymmetry in the costs of accepting or

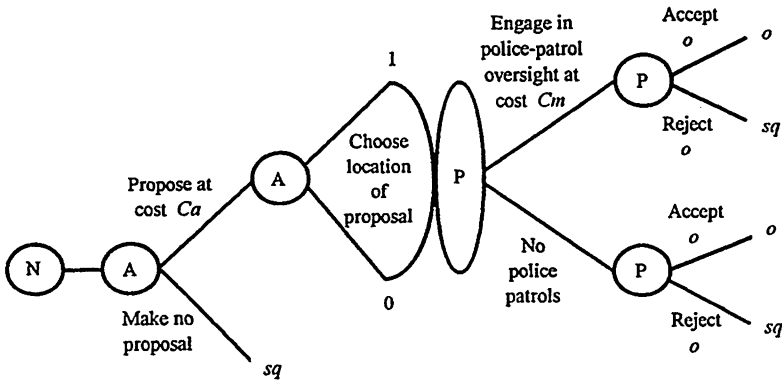


Figure 1. A model of police-patrol oversight. N refers to a move by Nature.

3.1 Learning from Police-Patrol Oversight

Under what conditions are the consequences of delegation beneficial for a principal who can conduct police-patrol oversight? If the agent expects the principal to engage in police-patrol oversight, then the agent knows that the principal will have the ability to distinguish proposals that are better for her than sq from those that are worse. The agent should expect the principal to reject any proposal that is worse for her than sq . Therefore, the only way the agent can recover its proposal costs is to make a proposal that is better for the principal than sq . As a result, employing police patrols ensures that the consequences of delegation are positive for the principal.

In contrast, when the agent does not expect the principal to undertake police-patrol oversight, the consequences of delegation for the principal can be detrimental. In this case, the agent makes an offer that is worse for the principal than is the status quo when (i) the best policy for the agent to purpose is worse for the principal than the status quo, and (ii) the principal's relative ignorance leads her to believe that o is closer to her ideal point than sq (see Proposition 2).

So, what determines whether or not the principal conducts police-patrol oversight? Simply put, the principal must expect that the benefits will outweigh the costs.⁶ The benefits from police-patrol oversight are highest when

rejecting an agent's policy offers. For simplicity, we examine the case where these costs are equal. Extending our model to the more general case of asymmetric costs is trivial; it merely requires the addition of one constant term to the principal's utility function (described in Appendix A). The consequences of adding this constant term to our description of the effect of each type of oversight are straightforward as well: the costlier action will be chosen less often, all else constant.

6. Some readers may wonder at this point about the possibility of the following infinite regress: believing that the agent will choose proposal strategy X , the principal will choose not to pay c_m ; but knowing this, the agent will choose proposal strategy X' ; but knowing this, the principal will choose to monitor, in which case the agent will once again be better off choosing X .

the principal is uncertain about which of sq or o will provide a larger payoff, but believes that choosing the alternative that is farther from her ideal point will be very costly to her. By contrast, if the principal believes that mistakes are not costly, or that costly mistakes are avoidable, due to sufficient information, then the benefits of police patrol-oversight will be small. All else constant, the principal is unlikely to find police-patrol oversight to be worthwhile in these cases.

3.2 Learning from Costly Effort

In addition to learning from police-patrol oversight, the principal in our model can, upon observing that the agent has taken costly action, form more accurate beliefs about the location of o , and, as a result, the consequences of accepting the proffered policy (see Lemma 1 and Proposition 1). In general, one person can learn about a second person's hidden knowledge by observing the choices that the second makes when some of the second person's actions are costly. What can be inferred by the first person is that the choice made by the second must be sufficiently valuable to the second person so that the expected gain from this choice outweighs the costs.

While the logic that drives the effect of costly entry also underlies Spence's (1973) finding about the ability of an employer to distinguish unskilled job applicants from skilled job applicants, placing costly action in the context of delegation provides its own unique inference. This dynamic is depicted in Figure 2. At the top of Figure 2 we present the principal's prior beliefs about the location of o as a uniform distribution. (While our results hold for any prior distribution, we illustrate the uniform case because it is the easiest to draw.) At the bottom of the figure we depict the principal's posterior beliefs after it learns that the agent has paid $c_a > 0$ to propose o . The difference between the principal's prior and posterior beliefs is that, having observed a proposal under proposal cost c_a , the principal can infer that o must be sufficiently distinct from sq to make the agent's payment of the proposal cost worthwhile. The principal's observation of the agent's costly action allows her to conclude that o cannot be located between the points a and b , where a is $sq - c_a$ and b is $sq + c_a$.⁷ In effect, the principal's ability to observe costly effort

To see why regress does not occur in this model, notice, first, that the agent cannot commit to any proposal strategy besides the one that places the policy offer at the point that is closest to the agent's ideal point, subject to the constraint that the principal is unable to detect that it is farther from her own ideal point than sq (i.e., if the agent could make the principal believe that she was choosing any other strategy, she would always have an incentive to defect back to the stated strategy). Second, that the principal has no private information implies that the value to the principal of choosing o is based solely on what is common knowledge to the other player. Third, that the principal moves last implies that she cannot commit to any particular monitoring strategy besides the one that maximizes her own expected utility. Thus, the agent knows at the time she chooses her strategy whether or not the principal will find it worthwhile to employ police patrols. This knowledge prevents the infinite regress.

7. For simplicity, we assume that the magnitude of the agent's costs are common knowledge. Our conclusions about the effect of costly action on the principal's ability to learn are robust to cases where the principal holds commonly known beliefs about the magnitude of agent costs.




Figure 2. Learning from agent's costly action.

by the agent produces relatively accurate posterior beliefs and gives the principal a greater ability to distinguish beneficial proposals from harmful ones.

It follows that when a proposal is made, the larger the proposal cost, the larger the difference between σ and sq (the larger the distance ab) must be. While the existence of observable, costly actions does not necessarily allow the principal to learn all of the bureaucrat's expertise, it does allow the principal to approximate with relative accuracy the minimum possible difference between σ and sq .

The principal's ability to observe costly action has an important impact on the consequences of delegation. When the presence of costly effort increases the principal's knowledge about the minimum distance between σ and sq , it also enables the principal to make more accurate distinctions about which alternative it should choose. When the agent knows that the principal is capable of making accurate distinctions, it also knows that it can recover its proposal costs only by making proposals that are better for the principal than sq . Therefore, the presence of observable and costly action can improve the consequences of delegation for the principal.

4. A Model of Fire-Alarm Oversight

While police-patrol oversight of agency activity is presumed to be effective, it is difficult to imagine that legislators actually have the resources (including time) that would be required to overcome their lack of expertise for more than but a few policy areas. Indeed, it is the costliness of police-patrol oversight that has led many scholars to be skeptical of the ability of legislatures to control the choices of the bureaucracies they create.

McCubbins and Schwartz, in arguing that members of Congress will prefer a fire-alarm oversight system, assumed implicitly that the statements made by

fire alarms to members of Congress always would be informative. This assumption is problematic, however. Simply put, constituents and others, in pulling their fire alarms, may attempt to mislead legislators. To make matters worse, it may be impossible for legislators to differentiate lying fire alarms from truthful fire alarms. Thus, if police patrols are too expensive for widespread use, then to understand the consequences of delegation, we must be able to answer the question “Can legislators learn from fire alarms who have the potential to mislead them?”

We model fire-alarm oversight as a situation where legislators receive signals from informed parties. These informed parties can be either the bureaucrats themselves (informed second parties) or constituents who have an interest in and information about bureaucratic activities (informed third parties). We isolate the effects of expertise and oversight on the consequences of delegation by modeling a situation where a single bureaucratic agent, who may or may not have the same policy preferences as the legislative principal, can use its previously delegated authority to propose an alternative to an existing policy. For expositional simplicity, our fire-alarm game begins after the agent has made a proposal (i.e., this game is between the information-providing fire alarm and the legislative principal). A single fire alarm then chooses whether or not to signal the principal about the consequences of the agent’s policy choice. The principal then decides to accept or reject the proposal. If the principal obtains a sufficient amount of information from oversight activities, we will show that it can influence the consequences of delegation (i.e., the policy choice). Otherwise, delegation is equivalent to abdication.

4.1 A Description of the Model

We again model a multistage, single-shot game between two players: a *fire alarm* and a *principal*. The principal’s task is to render a judgment on whether a previously offered bureaucratic proposal, again denoted o , or the existing policy of the government, again denoted sq , should prevail. As was true in the police-patrol model, these two policies and player ideal points are represented as points on the unit interval $[0,1]$. Also as before, each player’s objective is to obtain the policy, o or sq , that is closest to its ideal point.⁸ Unless stated otherwise, we assume that the value of each parameter in the game is common knowledge.

The exceptions to the common-knowledge assumption here are similar to those presented in the police-patrol model. The locations of the proposal o and the fire alarm’s ideal point may be known only by (i.e., are the private information of) the fire alarm. This is equivalent to saying that the fire alarm knows something about the consequences for the principal of the acceptance

8. We assume that the principal’s utility function is $-|X - P|$ and the fire alarm’s utility function is $-|X - F|$, where P is the principal’s ideal point, F is the fire alarm’s ideal point, and $X \in \{o, sq\}$. For the same reasons as stated in note 2, our results are not dependent on either unidimensionality or the utility functions stated.

of o that the principal is uncertain about (i.e., whether o or sq is closer to the principal's ideal point). Similarly, the fire alarm also knows something that the principal does not know about the fire alarm's own incentives.

The remaining important characteristic of our model is how we structure the principal/fire-alarm interaction. There are two actions we care about: what the fire alarm says and what the principal does in response to what the fire alarm says. To isolate the factors that lead these actors to take particular actions, we model their interaction as a series of three events, depicted in Figure 3. First, the values of several parameters are determined exogenous to the play of the game. These values are: the location of the status quo (sq) and the bureaucratic policy proposal (o), the location of the principal's and the fire alarm's ideal points, the magnitude of all costs, and player prior beliefs. Second, the fire alarm decides to send one of two messages to the principal. The fire alarm can tell the principal either that " o is closer to your ideal point than is sq " (i.e., o is "better") or that " o is at least as far from your ideal point as is sq " (i.e., o is "worse").⁹ In making this statement, the fire alarm has the option of telling the truth or lying. If the fire alarm lies, both players know that it is penalized an amount $t \geq 0$.¹⁰ Third, after observing the signal sent by the fire alarm, the principal either accepts o or rejects it in favor of sq .

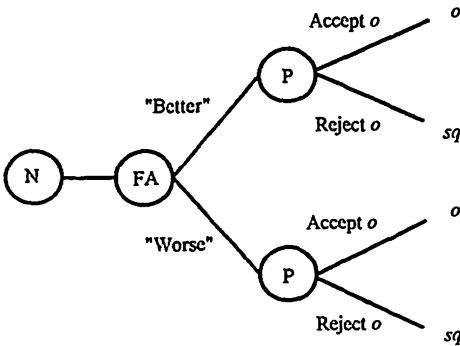
4.2 The Conditions for Learning from Fire-Alarm Oversight

We use the model to derive two conditions for learning from fire-alarm oversight: (i) the existence of penalties for lying on the fire alarm; and (ii) the perceived degree of similarity between fire-alarm and principal preferences. With the exception of a few extreme cases, neither of the two conditions we identify is either necessary or sufficient for learning. However, it is a necessary condition that at least one of the two conditions hold if the principal is to learn from fire-alarm oversight. In addition, straightforward comparative statics demonstrate that the amount that the principal can learn from fire-alarm oversight is nondecreasing in the size of the penalty for lying or the degree of preference similarity.

The first condition for learning is the existence of a penalty for lying (see Lemma 4 and Proposition 3 in Appendix B). This condition is a straightforward application of the economic concept of opportunity costs to the context

9. While the fire alarm is restricted to the statement "better" or the statement "worse," this conception of fire-alarm/principal communication allows us to draw general inferences about the effect of oversight in a parsimonious manner. As we will demonstrate, the context in which the binary "better/worse" message is delivered determines the extent to which the principal can use these messages to make relatively accurate "how much better than" and "how much worse than" inferences.

10. For expositional simplicity, we examine the case where the penalty is applied with certainty when the fire alarm's message is untruthful. Let p^* be the exogenously determined probability that the truthfulness of the fire alarm's signal will be verified. Since the case we discuss is equivalent to the case where the penalty t/p^* is applied with probability p^* , our simplification is made without a loss of generality. We think of the penalty for lying as the expected damage to an information-provider's reputation if he or she lies.



Note: FA knows the location of o , P may not.

Figure 3. A model of fire-alarm oversight. N refers to a move by Nature.

of communication. That is, when the marginal cost of lying is positive, lying will be a worthwhile activity only when the expected benefit of lying outweighs the expected cost.

Figure 4 shows how a penalty for lying affects the principal's beliefs about the location of o (again, for ease of illustration, we draw a uniform distribution to describe the principal's beliefs). Part (a) depicts the principal's prior beliefs; parts (b) and (c) depict the principal's posterior beliefs after it learns that the fire alarm, who faces penalty for lying $t > 0$, has sent the messages "better" and "worse," respectively. The difference between the principal's prior beliefs and the principal's posterior beliefs after observing the message "better" is that the principal can infer that o cannot actually be both close to sq and a little worse for the principal than sq (i.e., o cannot be between points d and e on Figure 4). The principal can make such an inference because it is common knowledge that this set of potential locations of o does not offer the fire alarm sufficient benefit to make lying worthwhile. Similarly, if the fire alarm sends the message "worse," then the principal can correctly infer that o cannot actually be both close to sq and a little better for the principal than sq (i.e., o cannot be between points f and g on Figure 4.) In short, if the fire alarm sends a particular message in the presence of a penalty for lying, then the principal can correctly infer that either (i) the message must be true or (ii) the message is false and the fire alarm believed that it was worthwhile to pay the penalty. Thus, what the principal learns in the presence of a penalty for lying is not necessarily that the fire alarm has told the truth, but rather that particular locations of o must be impossible given the fire alarm's actions. With this knowledge, the principal can generally use the content of the fire alarm's message to make more accurate inferences about the consequences of the bureaucratic proposal (the location of o).

The second condition for learning is the degree of similarity between the

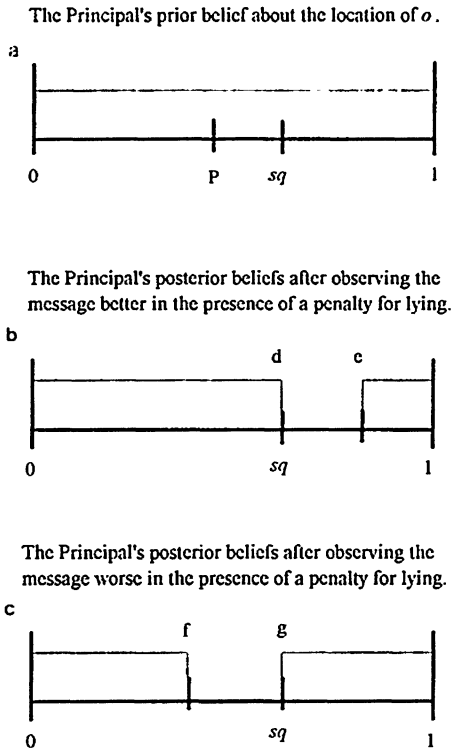


Figure 4. Learning from penalties for lying.

principal's and fire alarm's preferences over outcomes (see Lemma 5 and Corollary 1 in Appendix B). Underlying this insight is the principle that people can learn from others whose preferences are known to be similar to their own, since people with similar interests have little incentive to mislead each other. The condition we identify here follows that originally derived in Crawford and Sobel (1982) and applied insightfully in the context of lobbying by Austen-Smith (1993a, 1993b) and in the context of committee/floor relationships by Gilligan and Krehbiel (1987, 1989).¹¹ In short, a fire alarm who

11. Differences between the Crawford-Sobel model and our model allow us to draw unique inferences that are particularly useful for the problem at hand. To see this, we briefly review the differences between our model and Crawford-Sobel. The two modeling approaches differ in three important ways, the first two of which make our model more general than Crawford-Sobel. First, unlike the information receiver in the Crawford-Sobel model, our principal is uncertain about the location of the fire alarm's (i.e., the sender's) ideal point. Second, unlike Crawford and Sobel's sender, our fire alarm can send untruthful information. The third difference makes Crawford and Sobel's model more general than ours. Their information provider chooses its message from an infinite vocabulary, while our fire alarm can say only "better" or "worse." Since the human vocabulary is neither infinite nor binary, both assumptions about the size of the message space are

shares the principal's preferences over outcomes has relatively little to lose by revealing information about the location of o to the principal. Therefore, the more likely it is that the fire alarm and the principal prefer the same outcome, the greater is the weight that the principal assigns to the fire alarm's claim being true.

To see this, refer to Figure 5 and first consider the extreme case where the principal is certain that it and the fire alarm share the same preferences over outcomes (case b). Because the principal knows that the fire alarm wants to have the principal accept only those offers that are better than the status quo and reject all those that are worse, the fire alarm's message can be treated as though it were true. In contrast, when the principal is certain that it and the fire alarm have vastly different preferences over outcomes (case d), the fire alarm's message can be treated as though it were uninformative (regardless of whether it is actually true or false.) When the principal is uncertain about the similarity of its and the fire alarm's preferences over outcomes (case c), the principal's posterior beliefs about the location of o depend on the likelihood that it and the fire alarm share the same preferences over outcomes. When the fire alarm signals "better," the principal's posterior beliefs place more weight on states of the world where "better" is true and less weight on states of the world where "better" is false than is true when this likelihood is low.


In summary, if either the expected penalty for lying or the likelihood that the principal's and fire alarm's preferences are similar is sufficiently large, then the principal can use fire-alarm oversight to learn about the consequences of accepting bureaucratic proposal o (see Corollary 1 in Appendix B). As the strength of these conditions increases—in other words, as the penalty for lying increases and/or as the perceived similarity of preference between the principal and the fire alarm increases—so does the ability of the principal to learn from fire alarms (see Proposition 4 and Corollary 1 in Appendix B). Under such conditions, legislators will likely prefer fire-alarm oversight to police-patrol oversight, as McCubbins and Schwartz argued.

The ability to learn from fire alarms affects the consequences of delegation. The presence of informative fire alarms allows the principal to more accurately determine which alternative, o or sq , is better for it. When the principal can distinguish between o and sq , it can reject proposals that are worse than the status quo. If the proposal is being made by an agent who pays a cost to bring the proposal, as in our model of police patrols, then when fire alarms are sufficiently informative, the principal can discourage the agent from making proposals that are not better for her than sq .¹² The presence of informative


extreme. The benefit of our approach is that we can identify the simultaneous effect of contextual variables like preference proximity, penalties for lying, and costly action, which are very difficult to identify when the message space has an infinite number of elements.

12. For simplicity, we do not present a fire-alarm model that includes the agent as a strategic player. We can, however, conjoin the fire-alarm and police-patrol models presented in the Appendixes. Once the models are conjoined, we are able to demonstrate this result.


The Principal's prior belief about the location of θ .



The Principal's posterior beliefs about the location of θ after observing the message better from a fire alarm who is known to have the same ideal point.



The Principal's posterior beliefs about the location of θ after observing the message better from a fire alarm who is very likely to have the same ideal point.



The Principal's posterior beliefs about the location of θ after observing the message better from a fire alarm who is known to have the opposite preferences over outcomes as does the Principal.




Figure 5. Learning from similarity of preferences (holding penalty for lying = 0).

fire-alarm oversight can thus improve the consequences of delegation for the principal.

5. Conclusion

It has often been alleged that Congress's broad delegations of policy-making authority to the bureaucracy in this century have led to a sweeping abdication of congressional power. The widely recognized cause of this potential abdication is hidden knowledge: that is, bureaucrats have information that members of Congress do not have, information that is important for making or evaluating policy. Legislative oversight of the bureaucracy is the most obvious means of abating the problem of hidden knowledge.

Oversight takes two general forms: direct, central surveillance of bureaucratic activity, (i.e., police patrols); or indirect, decentralized monitoring of bureaucratic policy-making by third parties (i.e., fire alarms) (see McCubbins and Schwartz, 1984). Fire-alarm oversight, it is argued, is less costly than police-patrol oversight for legislators, because it relies on the monitoring activities of constituents and interest groups, rather than on the direct actions of legislators. However, fire alarms are not necessarily as effective as police patrols. Indeed, if legislators are unable to make accurate inferences about their agents' hidden knowledge from what fire alarms say, legislators who rely on such oversight will find they are unable to manage bureaucratic policy-making. Hence, not only must fire alarms be less expensive than police patrols for legislators to choose them, but it must also be the case that legislators can learn something from the statements made by fire alarms.¹³

We identify two conditions under which learning from fire alarms can occur: (a) penalties for lying and (b) a perceived similarity of preferences between fire alarms and legislators. If both of these conditions are weak or absent, and if police-patrol oversight is too expensive to be practical, then delegation will be equivalent to abdication. Otherwise, legislators can learn from fire alarms and manage the activities of their bureaucratic agents.¹⁴

The main problem legislators face, then, is in establishing the conditions under which learning through oversight can occur. However, legislators are not equally well suited for affecting the costs and benefits of each type of oversight. For instance, all that can be done to affect the costs and benefits of

13. In defining the conditions under which oversight can be effective in managing bureaucratic policy choices so that they are beneficial to legislative principals (and presumably, then to voters), we also have demonstrated that legislators need not become fully informed about a policy choice in order to manage it effectively. Rather, legislators need only learn enough to know whether the bureaucratic proposal is better or worse for them than the reversion. In sum, bureaucratic expertise and some legislative ignorance are not sufficient to equate delegation and abdication.

14. These general findings imply that the conditions under which we expect the principal to rely on fire-alarm oversight, police-patrol oversight, or neither (in the decision to accept or reject o) depend not only on the relative costs of fire-alarm and police-patrol oversight, as has been previously argued, but also on what can be learned from each type of oversight and the effect of oversight-induced learning on behavior and beliefs. It follows that the principal will depend exclusively on fire-alarm oversight when it is expected that credible fire alarms will participate and when the cost of police-patrol oversight is high relative to the expected gain to the principal from direct monitoring. This happens when one or both of the conditions for the legislative principal to learn from fire alarms is sufficiently strong to allow the principal to learn enough to make the delegation beneficial. The principal will depend exclusively on police-patrol oversight when none of the conditions for learning from the fire alarm are present or are too weak to be sufficiently informative, and when the cost of police-patrol oversight is low relative to the expected gain to the principal from this activity. The principal will use neither type of oversight when the conditions for learning are absent and the cost of police-patrol oversight is high relative to the expected gain to the principal from this activity. In this case, the principal decides whether to accept or reject o based only on its prior beliefs and observations of the agent's action. In this case, if the principal has sufficient prior information about the agent's policy choice, then it can effectively manage that choice so that the outcome is beneficial; if it does not have sufficient prior information, then delegation will equal abdication.

police-patrol oversight is to simplify the delegation problem by choosing more reliable agents, which itself requires better informed legislators or delegating only on those issues for which legislators have easy access to information. The usually high opportunity costs of a legislator's time would seem to render each remedy inviable.

By contrast, previous research by McCubbins et al. (1987), Mazmanian and Nienaber (1979), and Ferejohn (1987) suggest that the conditions for learning from fire alarms can be generated by the structure and process of bureaucratic decision-making and the process of oversight itself.¹⁵ For example, legislators can increase the likelihood that bureaucratic activity produces results that are favorable to certain constituents by making the cost of participation in the fire-alarm process relatively low for them. Consider the case of individual consumers who rarely have the resources to participate in agency decision-making without outside assistance. Several laws written in the 1970s, such as the Magnuson-Moss Act of 1974, created public participation programs that paid consumers and other underrepresented groups for participating in regulatory proceedings. These programs increased consumer representation in regulatory proceedings and gave consumers new opportunities to play the role of fire alarms.

We have identified conditions under which learning from oversight is possible. It follows that when legislators can design institutions that produce these conditions, delegation is likely to be an effective way for legislators to affect the determination of complex policy matters. In contrast, when legislators are unable to design an institution that produces any of the conditions under which learning from oversight is possible, the delegation of authority to bureaucratic agencies is likely to be equivalent to abdication.

Note: The two appendixes that follow describe the formal logic underlying how we think about police-patrol and fire-alarm oversight. In Appendix A, we develop a model of police-patrol oversight; in Appendix B, a model of fire-alarm oversight.

A. A Model of Police-Patrol Oversight

A.1 Basic Premises

Two players, called the *principal* and the *agent*, play a single-shot game. Unless otherwise stated, all aspects of this game are common knowledge. The purpose of the game is to choose one of two points, o , which is endogenously determined, and sq , which is exogenously determined, on the line segment $[0,$

15. For instance, legislators may be able to enhance the credibility of the signals they receive from a single fire alarm by making it relatively cheap and easy for competing fire alarms both to gain access to bureaucratic policy-making and to then send signals to legislators. The competing fire alarms could serve as a check on the truthfulness of the information provided by other fire alarms. Under certain conditions, which we explore elsewhere (Lupia and McCubbins, 1993), this system of competing fire alarms makes the signals sent by at least one of the fire alarms credible. This credibility is obtained without an exogenous penalty for lying on all of the fire alarms.

1]. This choice determines a payoff in utils for each player, and each player's objective is to maximize their own utility. The principal's utility function is $-|X - P|$ and the agent's utility function is $-|X - A|$, where P is the principal's ideal point, A is the agent's ideal point, and $X \in \{o, sq\}$.

While the exact location of P is commonly known, the location of A is the result of a single draw from the distribution α . Distribution α has positive mass at no more than three known points on $[0,1]$. A_n is a point that is on the same side of sq as P and is closer to P than is sq . A_f is a point that is on the same side of sq as P and is farther from P than is sq . A_o is a point that is on the opposite side of sq from P . For $i \in \{n, f, o\}$, let q_i be the prior probability that $A = A_i$. Only the agent knows the result of this draw.

The agent makes the game's first move when he decides whether or not to make a proposal o : $(o(sq, A, \alpha, c_a, P, c_m) \in [0,1])$. If the agent makes no proposal, the game ends with $o = sq$ as the point that determines player payoffs. If the agent proposes, it pays proposal cost c_a , selects the location of o , and forces the principal to choose between o and sq . Before she chooses sq or o , the principal can choose to pay c_m for an opportunity to learn the location of o ($\text{MON}(sq, \alpha, c_a, P, c_m) \in \{0,1\}$). The principal then ends the game by deciding whether or not to approve o : $(\text{APP}(sq, \alpha, \text{MON}, c_a, P, c_m) \in \{0,1\})$. To simplify the exposition, we assume that (i) if the expected benefit of taking a costly action (proposing, monitoring) is not strictly positive, then this action is not taken; and (ii) if sq and o provide the principal with the same expected utility, then the principal chooses sq .

The equilibrium concept we use is the sequential equilibrium concept of Kreps and Wilson (1982). In words, a sequential equilibrium for this game consists of:

(i) $\text{APP}^* = 1$ iff the expected utility for the principal from the proposal is greater than the utility from the status quo;

(ii) $\text{MON}^* = 1$ iff the expected utility for the principal from the monitoring is greater than the utility from not monitoring, given that optimal decision rule APP^* will be used to choose between o and sq .

(iii) $o^* \in \text{argmax}_{o \in [0,1]} \{-|o - P| - c_a, -|sq - P|\}$ or the agent maximizes her utility subject to principal decision rule $\{\text{MON}, \text{APP}\}$.

(iv) the principal's beliefs about the location of o are consistent.

Technical statements of each condition follow.

A.2 Equilibrium Behavior

We first examine the relationship between the agent's proposal cost and the principal's beliefs about the location of o . In short, proposal costs can dissuade an agent from proposing. This dissuasion occurs with certainty when the agent's ideal point is close enough to the status quo that the benefit from making a winning proposal could not be greater than the cost of entry. Let ε be the maximum distance for which this statement is true. The proofs of the following lemma and proposition follow straightforwardly from the assumption of expected utility maximization, the definition of the agent's utility function, and the definition of ε .

Lemma 1. If the agent's ideal point is within distance ε of sq he should make no proposal.

Proposition 1. Learning from costly entry: If the principal observes that an offer was made in the presence of entry cost c_a , then she can infer that $o \notin [sq - \varepsilon, sq + \varepsilon]$.

We now identify undominated strategies for the agent. For $i \in \{n, f, o\}$, let points y_i for which $-|y_i - P| > -|sq - P|$ and $-|sq - A_i| < -|y_i - A_i| - c_a$ be members of the set Y_i . Y_i represents the set of points that both an agent of type i and the principal prefer to sq .

Lemma 2. It is an undominated partial strategy for agent types A_n and A_o to propose their ideal points if they make a proposal. This description is also true for agent type A_f when Y_f is empty.

Proof. Notice first that conditional on the principal approving o , an agent maximizes his utility by proposing that o be located at his ideal point. P common knowledge implies that, in equilibrium, the agent will know precisely how the principal's strategy is conditioned on the proposal he makes. In the case where the agent knows that the principal will not monitor, he also knows that the principal's approval strategy is not conditioned on the actual location of o . In this case, the agent cannot be rewarded for choosing o at a point other than his ideal point. This is true even if a nonmonitoring principal believes that any agent that makes a proposal will propose his ideal point. Therefore, proposing that o be located at the agent's ideal point is an undominated partial strategy. In the case where the agent knows that the principal will monitor, agent type A_n knows that the principal will accept proposal $o = A_n$ and can do no better than proposing his ideal point. In contrast, agent type A_o (and A_f if Y_f is empty) knows that there exists no point on $[0, 1]$ that both he and the principal prefer to sq . Therefore, "do not participate" will be a dominant strategy for these agent types in this situation. Q.E.D.

Lemma 3. When Y_f is nonempty, and agent type A_f is certain that the principal will monitor, an undominated partial strategy for the agent is to propose the element of Y_f that is the same distance from P as is sq .

The proof of this lemma follows straightforwardly from Romer and Rosenthal (1978). In short, agent type A_f knows that his ideal point will be rejected by a monitoring principal. If Y_f is nonempty, then agent type A_f chooses the point that is both closest to his ideal point and (in the limit) at least as close to P as is sq . From the definition of A_f , it follows that this point will be the same distance from P as is sq .

To simplify the statement of the game's equilibrium, we provide some additional pieces of notation. To represent the effect of costly entry on the principal's beliefs about the location of o , let p_i equal q_i iff $-|A_i - sq| > \varepsilon$ and

zero otherwise. Let a denote the condition where monitoring provides greater expected utility than does choosing sq without monitoring. From Lemmas 2 and 3, it follows that if Y_f is empty, then a is true iff $|sq - A_n| > c_m$ and if Y_f is nonempty, then a is true iff $(\frac{p_n}{p_n+p_f} \times |sq - A_n|) > c_m$. Let b denote the condition where monitoring provides greater expected utility than does not monitoring and choosing o . From Lemmas 2 and 3, it follows that if Y_f is empty, then b is true iff

$$\left(\frac{p_j}{p_n + p_f + p_o} \times |A_f - P| - |A_n - P| \right) + \left(\frac{p_o}{p_n + p_f + p_o} \times |A_o - P| - |A_n - P| \right) > c_m,$$

and if Y_f is nonempty, then b is true iff

$$\left(\left(\frac{p_n}{p_n + p_f + p_o} - \frac{p_n}{p_n + p_f} \right) \times |A_n - P| \right) + \left(\frac{p_f}{p_n + p_f + p_o} \times |A_f - P| \right) + \left(\frac{p_o}{p_n + p_f + p_o} \times |A_o - P| \right) - \left(\frac{p_f}{p_n + p_f} \times |sq - P| \right) > c_m.$$

Proposition 2. The equilibrium of the police-patrol game is as follows:

(i) If not monitoring and choosing o provides greater expected utility than does monitoring (i.e., if b is false), if agent type i ($i \in \{n, f, o\}$) is drawn and $|A_i - sq| > c_a$, and

$$\left(\frac{p_n}{p_n + p_f + p_o} \times -|A_n - P| \right) + \left(\frac{p_f}{p_n + p_f + p_o} \times -|A_f - P| \right) + \left(\frac{p_o}{p_n + p_f + p_o} \times -|A_o - P| \right) > -|sq - P|,$$

then $o = A_i$ and the principal chooses o without monitoring.

(ii) If monitoring provides greater expected utility than does not monitoring (i.e., if a and b are true), then

(a) if A_n is chosen and $|A_n - sq| > c_a$, then $o = A_n$ and the principal accepts o .

(b) if A_f is chosen, $P < sq$ and $|2 \times |sq - P|| > c_a$, then $o = sq - (2 \times |sq - p|)$ and the principal accepts o .

(c) if A_f is chosen, $P > sq$ and $|2 \times |sq - P|| > c_a$, then $o = sq + (2 \times |sq - p|)$ and the principal accepts o .

(iii) Otherwise, no proposal is made and sq is the outcome.

Proof. In case (i), the principal finds it worthwhile to choose o without monitoring even given the expectation that all agent types whose ideal points

are not within ε of sq will propose their ideal points. The best response of any such agent in this case is to propose his ideal point if it is not within ε of sq . Expected utility maximization by the principal and Lemma 2 imply that the choice of any alternate strategy by either player is not a best response given available information. In case (ii), the principal finds it worthwhile to monitor. Agent behavior follows from Lemmas 2 and 3. The set of other cases includes either “the principal maximizes interim expected utility by choosing sq without monitoring” or “the principal would maximize interim expected utility by monitoring, but the agent, who is cognizant of what the principal will do, cannot make a proposal that makes both he and the principal better off than sq .” In each of these cases, the agent’s best response is to make no proposal. Q.E.D.

Appendix B. A Model of Fire-Alarm Oversight

B.1 Basic Premises

Two players, called the *principal* and the *fire alarm*, play a single-shot game. Unless otherwise stated, all aspects of this game are common knowledge. The purpose of the game is to choose one of two exogenously determined points, o and sq , on the line segment $[0,1]$. This choice determines a payoff in utils for each player and each player’s objective is to maximize their own utility. The principal’s utility function is $-|X - P|$ and the fire alarm’s utility function is $-|X - F|$, where P is the principal’s ideal point, F is the fire alarm’s ideal point, and $X \in \{o, sq\}$. For expositional simplicity, we discuss the case where $P < sq$. Our results are without a loss of generality to the case $P > sq$, which is equivalent, or to the case $P = sq$, which is trivial.

As before, the location of P is common knowledge, however the locations of o and F may be the private information of the fire alarm. We assume that the location of o is the result of a single draw from the distribution O and that location of F is the result of a single draw from the distribution Γ . O has density O' , Γ has density Γ' , and each has support on known, but undenoted, subsets of $[0,1]$. In effect, we assume that O and Γ are common knowledge and that only the fire alarm observes the result of the draw from each distribution. If either distribution has mass at more than one point, the fire alarm has private information. For expositional simplicity, we examine the case where $O'(sq) = 0$.

The fire alarm makes the game’s first move when he decides to send one of two messages, $M(\Gamma, sq, o, t) \in \{B, W\}$. B (better) means that o is closer to P than is sq . W (worse) means that o is farther from P than is sq . The fire alarm is not restricted to the transmission of a truthful message, but may have to pay an additional penalty for lying, t , if it chooses to dissemble. The principal must then either approve or disapprove $o(\text{APP}(\Gamma, sq, O, M, t) \in \{0, 1\})$. After making this choice, the game ends and the principal and the fire alarm receive a payoff in utils.

This article’s equilibrium concept is a variant of the sequential equilibrium concept of Kreps and Wilson. A sequential equilibrium consists of strategies that players believe to be the best responses to the chosen strategies of others,

given prior beliefs that are consistent and their use of an updating procedure that is based on Bayes' Rule. (Consistency implies that player beliefs assign positive probability to the true state of the world.)

The variation we introduce is that we assume that the principal utilizes an exogenously determined algorithm to determine whether or not to condition her beliefs on her knowledge of the fire alarm's strategy. We introduce this concept to simplify the formal statement of the model and the exposition that follows. The algorithm suggests that a principal with limited cognitive resources will opt to consider the fire alarm's statement if she expects, without explicitly considering all possible outcomes of the game, that doing so will increase the probability that she makes the same decision she would have made had she known the location of o . An algorithm with these characteristics is proposed in Section B.4. The remainder of our analysis focuses on the case where the algorithm directs the principal to use (or stated another way, the principal uses) the content of the fire alarm's message, the existence and magnitude of the penalty for lying, the probability that the fire alarm and the principal have the same preference ordering over $\{o, sq\}$, and Bayes' Rule to update her prior beliefs about the location of o . The validity of our results relies on the validity of this equilibrium concept, as we do not examine the consequences of play that strays from the equilibrium path. To simplify the exposition, we assume that if sq and o provide the principal with the same expected utility, then the principal chooses sq .

B.2 Interim Steps

Let τ be the smallest distance from the point sq for which the fire alarm could find the payment of a penalty for lying t to be worthwhile. Since sq , t , and the shape of the fire alarm's utility function are common knowledge, so is τ .

Lemma 4. In the presence of penalty for lying t , truth telling is an undominated partial strategy for the fire alarm when $o \in [sq - \tau, sq + \tau]$.

Proof. When $o \in (sq, sq + \tau]$, a fire alarm that signals B is required to pay the penalty for lying. Similarly, when $o \in [sq - \tau, sq)$, a fire alarm that sends the signal W is required to pay the penalty for lying. In each case, the definition of τ implies that the maximum possible benefit to the fire alarm of affecting the outcome cannot possibly be higher than the penalty for lying. Therefore, truth telling is an undominated partial strategy in the cases described. Q.E.D.

The following proposition follows straightforwardly from Lemma 4 and is offered without proof.

Proposition 3. Learning from a penalty for lying: The density of O at o (or a closed interval of small and positive length with endpoints that are equidistant from o) in the principal's posterior beliefs minus the density of O at that point (or interval) in the principal's prior beliefs is nondecreasing in t .

For example, when the principal observes B in the presence of penalty for lying t , she learns that o cannot be located in the interval $(sq, sq + \tau)$. Similarly, upon observing W , the principal learns that o cannot be located in the interval $(sq, sq + \tau)$. In each case, the principal learns that particular locations for o are impossible, which may be knowledge that she would not have possessed in the absence of a penalty for lying.

We now move to the relationship between the fire alarm's incentives for truth telling and the similarity of fire-alarm and principal preferences. We begin by making a preliminary claim whose proof follows the same logic as that found in Crawford and Sobel (1982).

Lemma 5. When it is common knowledge that $-|o - F| > -|sq - F|$ and $-|o - P| > -|sq - P|$, then the fire alarm should send B and the principal should treat the message as though it were true. Similarly, when it is common knowledge that $-|o - F| \leq -|sq - F|$ and $-|o - P| \leq -|sq - P|$, then the fire alarm should send W and the principal should treat the message as though it were true. When it is common knowledge that $t = 0$ and either $-|o - F| \leq -|sq - F|$ and $-|o - P| > -|sq - P|$ or $-|o - F| > -|sq - F|$ and $-|o - P| \leq -|sq - P|$, then the principal should disregard the content of the fire alarm's message.

We now apply the logic of Lemmas 4 and 5 to describe the simultaneous impact of penalties for lying and similarity of preferences on the principal's beliefs about the location of o . First, recall that for the case $IP < sq$, $(sq - (2 \times (sq - P)), sq)$ is the range of points on $[0,1]$ that provides the principal with higher utility than sq . Now, let s_B be the probability that o is better for both the principal and the fire alarm:

$$s_B = [O(sq) - O(sq - (2 \times (sq - P)))] \times \text{prob}[o: -|F - o| > -|F - sq| \text{ if } o \in (sq - (2 \times (sq - P)), sq)]$$

Let s_W be the probability that o is worse for both the principal and the fire alarm:

$$s_W = [1 - O(sq) + O(sq - (2 \times (sq - P)))] \times \text{prob}[o: -|F - o| \leq -|F - sq| \text{ if } o \notin (sq - (2 \times (sq - P)), sq)]$$

Since P is common knowledge, it is also common knowledge that $s_B + s_W \leq 1$.

Now, let d_B be the common prior probability that the principal and the fire alarm have different preferences over the set $\{o, sq\}$ when $o \in [sq - \tau, sq)$. Let d_W have an equivalent definition for the case $o \in [sq, sq + \tau)$. Probabilities d_B and d_W are the probabilities that the penalty for lying is large enough to persuade a fire alarm, who would otherwise find dissembling worthwhile, to send a truthful message.

$$\begin{aligned}
d_B \in [0,1] &= (O(sq) - (1 - O(sq - \tau)) \times \\
&\quad [\text{prob}(o: -|F - o| > -|F - sq|, -|o - P| \leq -|sq - P| \\
&\quad \text{if } o \in [sq - \tau, sq)) \\
&\quad + \text{prob}(o: -|F - o| \leq -|F - sq|, -|o - P| > -|sq - P| \\
&\quad \text{if } o \in [sq - \tau, sq))]
\end{aligned}$$

$$\begin{aligned}
d_W \in [0,1] &= (O(sq + \tau) - (1 - O(sq)) \times \\
&\quad [\text{prob}(o: -|F - o| > -|F - sq|, -|o - P| \leq -|sq - P| \\
&\quad \text{if } o \in [sq, sq + \tau)) \\
&\quad + \text{prob}(o: -|F - o| \leq -|F - sq|, -|o - P| > -|sq - P| \\
&\quad \text{if } o \in [sq, sq + \tau))]
\end{aligned}$$

It follows that $1 - s_B - s_W - d_B - d_W$ is the common prior probability that the fire alarm is a type that could find it profitable to lie.

We now combine insights from Lemmas 4 and 5 to propose updating schemes that are consistent with Bayes's Rule and that lead to beliefs that are consistent in the sense of Kreps and Wilson. While these are not the only updating schemes available, they are the minimal inference that the principal can draw if she decides to utilize information about penalties for lying and preference similarity. It follows that when $\tau < 2 \times (sq - P)$, the principal's posterior beliefs are related to her prior beliefs in the following manner:

$$\begin{aligned}
O'(o|B) &= \left(\frac{s_B}{1 - (s_W + d_W)} \times \frac{O'}{O(sq - \tau) - O(sq - 2 \times (sq - P))} \right) \\
&\quad + \left(\frac{1 - s_W - d_W - s_H - d_H}{1 - (s_W + d_W)} \times \frac{O'}{1 - O(sq + \tau) + O(sq - \tau)} \right) \\
&\quad o \in [sq - 2 \times (sq - P), sq - \tau)
\end{aligned}$$

(B could have been sent by a fire alarm who either has the same preferences over outcomes or is attempting to mislead.)

$$\begin{aligned}
O'(o|B) &= \left(\frac{d_B + s_B}{1 - (s_W + d_W)} \times \frac{O'}{O(sq) - (1 - O(sq - \tau))} \right) \\
&\quad o \in [sq - \tau, sq)
\end{aligned}$$

(B could have been sent by a fire alarm who either has the same preferences over outcomes or faces a large penalty for lying.)

$$O'(o|B) = 0 \quad o \in (sq, sq + \tau]$$

(If o were in this range, B would not be sent.)

$$\begin{aligned}
O'(o|B) &= \left(\frac{1 - s_B - d_B - s_W - d_W}{1 - (s_W + d_W)} \times \frac{O'}{1 - O(sq + \tau) + O(sq - \tau)} \right) \\
&\quad o \in [0, sq - 2 \times (sq - P)) \cup (sq + \tau, 1]
\end{aligned}$$

(B would be sent regardless of the truth.)

$$O'(o|W) = \left(\frac{s_W}{1 - (s_B + d_B)} \times \frac{O'}{1 - O(sq + \tau) + O(sq - 2 \times (sq - P))} \right) \\ + \left(\frac{1 - s_B - d_B - s_W - d_W}{1 - (s_B + d_B)} \times \frac{O'}{1 - O(sq + \tau) + O(sq - \tau)} \right) \\ o \in [0, sq - 2 \times (sq - P)) \cup (sq + \tau, 1]$$

(W could have been sent by a fire alarm who either has the same preferences over outcomes or is attempting to mislead.)

$$O'(o|W) = \left(\frac{d_W + s_W}{1 - (s_B + d_B)} \times \frac{O'}{O(sq + \tau) - (1 - O(sq))} \right) \\ o \in (sq, sq + \tau]$$

(W could have been sent by a fire alarm who either has the same preferences over outcomes or faces a large penalty for lying.)

$$O'(o|W) = 0 \quad o \in (sq - \tau, sq]$$

(If o were in this range, W would not be sent.)

$$O'(o|W) = \left(\frac{1 - s_B - d_B - s_W - d_W}{1 - (s_B + d_B)} \times \frac{O'}{1 - O(sq + \tau) + O(sq - \tau)} \right) \\ o \in (sq - 2 \times (sq - P), sq - \tau)$$

(W would be sent regardless of the truth.)

Similarly, when $\tau \geq 2 \times (sq - P)$, the principal's posterior beliefs are related to her prior beliefs in the following manner:

$$O'(o|B) = \left(\frac{s_B + d_B}{1 - (s_W + d_W)} \times \frac{O'}{O(sq) - O(sq - 2 \times (sq - P))} \right) \\ o \in [sq - 2 \times (sq - P), sq)$$

$$O'(o|B) = 0 \quad o \in [sq - \tau, sq - 2 \times (sq - P)] \cup (sq, sq + \tau]$$

$$O'(o|B) = \left(\frac{1 - s_B - d_B - s_W - d_W}{1 - (s_W + d_W)} \times \frac{O'}{1 - O(sq + \tau) - O(sq + \tau)} \right) \\ o \in [0, sq - \tau) \cup (sq + \tau, 1]$$

$$O'(o|W) = \left(\frac{d_W + s_W}{1 - (s_B + d_B)} \right)$$

$$\times \frac{O'}{O(sq + \tau) - (1 - O(sq)) + O(sq - 2 \times (sq - P)) - O(sq - \tau)} \\ o \in [sq - \tau, sq - 2 \times (sq - P)] \cup (sq, sq + \tau]$$

$$O'(o|W) = 0 \quad o \in [sq - 2 \times (sq - P), sq]$$

$$O'(o|W) = \left(\frac{1 - s_B - d_B - d_W}{1 - (s_B + d_B)} \times \frac{O'}{1 - O(sq + \tau) + O(sq - \tau)} \right) \\ o \in [0, sq - \tau] \cup (sq + \tau, 1]$$

It is easy to verify that this updating scheme renders the content of the fire alarm's message uninformative when $s_B = s_W = t = 0$ and perfectly credible when either $s_B = s_W = 1$ or when t is sufficiently high.

To simplify the statement and proof of this game's equilibrium, we make one additional assumption. We assume that if s_B , s_W , and τ are sufficiently low, the principal can commit to ignoring the content of the fire alarm's message. We make this assumption in response to the following scenario:

Variables s_B , s_W , and τ are sufficiently low that, when using the appropriate updating scheme, the principal's strategy is not conditioned on the content of the fire alarm's message. In effect, the principal ignores the fire alarm's advice. If the fire alarm anticipates that he will be ignored, then in the presence of a positive penalty for lying, he has no incentive to lie. If the fire alarm has no incentive to lie, then the principal should treat the content of his message as though it were true. However, if the principal were to change updating schemes and not ignore the fire alarm, then the fire alarm's incentive to lie would reappear. Since the values of s_B , s_W , and τ have not changed, then at this point, the first statement in this scenario would again be true.

Some scholars have dealt with this type of scenario by implementing the *revelation principal* (named by Myerson (1979)). Since we are attempting to identify conditions under which the principal can learn, employing the revelation principal would increase the range of conditions under which learning was possible but would also reduce the power of our argument. Our decision to assume that the principal ignores the fire alarm in this situation reflects our desire to shift the burden of proof onto the power of the conditions for learning in the remaining cases.

B.3 Equilibrium Behavior

We now state the equilibrium conditions for the case $t < 2 \times (sq - P)$. The statement for the other case follows straightforwardly. Notice that each condition for which o is chosen is first stated in words and then stated formally.

Proposition 4. If the algorithm directs the principal to consider information about the penalty for lying and preference similarity, and if the principal uses the Bayesian updating schemes defined above, then o is the equilibrium outcome of the game if and only if one of the following statements is true:

(i) o is better than sq for both players and the fire alarm is sufficiently credible (i.e., some elements of the set (s_B, s_W, d_B, d_W, t) are large enough to cause prior and posterior beliefs to diverge by such a degree that the principal's strategy depends on the content of the fire alarm's message) that he can persuade the principal to choose o .

$$o \in (sq - 2 \times (sq - P), sq - \tau), \quad -\int |o - P| dO'(o|B) > -|sq - P|,$$

and $-\int |o - F| > -|sq - F|.$

(ii) o is better than sq for the principal, o is not necessarily better than sq for the fire alarm, the expected penalty for lying faced by the fire alarm is larger than the maximum possible benefit from lying, and the fire alarm is sufficiently credible that he can persuade the principal to choose o .

$$o \in [sq - \tau, sq) \quad \text{and} \quad -\int |o - P| dO'(o|B) > -|sq - P|.$$

(iii) o is worse than sq for the principal and is better than sq for the fire alarm, and the fire alarm is sufficiently credible that he can persuade the principal to choose o .

$$o \in [0, sq - 2 \times (sq - P)] \cup (sq + \tau, 1], \quad -\int |o - P| dO'(o|B) > -|sq - P|,$$

and $-\int |o - F| > -|sq - F|.$

(iv) The fire alarm is not sufficiently credible to affect the principal's behavior and regardless of the fire alarm's action, the principal will accept o .

$$\min(-\int |o - P| dO'(o|B), -\int |o - P| dO'(o|W)) > -|sq - P|.$$

Proof. As a preliminary statement, notice first that the assumption of expected utility maximization and Lemmas 1 and 2 imply that the principal chooses o iff either the fire alarm signals B and $-\int |o - P| dO'(o|B) > -|sq - P|$ or the fire alarm signals W and $-\int |o - P| dO'(o|W) > -|sq - P|$. If case (i) is true, then signaling B is a unique best response for the fire alarm given his beliefs, and "accept if B is signaled" is a unique best response for the principal given her beliefs. If case (ii) is true, then lying is a dominated strategy for the fire alarm, therefore the same best responses as stated in case (i) apply here. If case (iii) is true, then the fire alarm can and would find it worthwhile to persuade the principal to choose her least preferred outcome, therefore the case (i) best responses also apply here. If case (iv) is true, then the fire alarm cannot affect the principal's choice of strategy. Thus, the principal responds to her prior beliefs by choosing o . Q.E.D.

Corollary 1. If prior beliefs are consistent, then the likelihood that the principal chooses the element of $\{o, sq\}$ that it would have chosen had it known the location of o is nondecreasing in $s_B + s_W$ and t .

Proof. It is sufficient to show that as either $s_B + s_W$ or t increases, the likelihood that o is chosen if $o \in (sq - 2 \times (sq - P), sq)$ is nondecreasing. It follows from Lemma 2, and is easy to verify using the updating tables, that the density of O at o (or a closed interval of small and positive length with endpoints that are equidistant from o and within the bounds $(sq - 2 \times (sq - P), sq)$) in the principal's posterior beliefs minus the density of O at that point (or interval) in the principal's prior beliefs is nondecreasing in s_B and s_W . Proposition 1 provided a similar statement for the effect of a penalty for lying. If the true location of o is closer to P than is sq , then it follows from Lemma 1, Lemma 2, and Bayes's Rule that the probability that the message B is sent and the magnitude of the probability mass placed on o (or a finite interval that has o as its center and boundaries within $(sq - 2 \times (sq - P), sq)$) are both nondecreasing in s_B , s_W , or t . As the probability mass on this interval increases, so does the likelihood that $-\int |o - P| dO'(o|B) > -|sq - P|$, and so does the likelihood that o is chosen. Q.E.D.

B.4 An Algorithm That Determines the Principal's Willingness to Hear a Fire Alarm

One of the problems faced in modeling signaling games is that the probability that the message receiver reacts to a message in a particular way is dependent on the actions of the message sender, which themselves are dependent on the probability that the message receiver reacts to a message in a particular way. This type of problem often requires modelers to make special assumptions in order to obtain useful results. Examples of these assumptions are the consistency requirements implicit in the definition of the sequential equilibrium (Kreps and Wilson) and in several refinements of the concept (see Banks, 1991, for a review). Our response to this problem is to invoke an algorithm that we believe is a good representation of how people deal with this type of situation. The algorithm suggests that a principal with limited cognitive resources will opt to consider the fire alarm's statement if she expects, without explicitly considering all possible outcomes of the game, that doing so will increase the probability that she makes the same decision she would have made had she known the location of o . This algorithm's invocation allows for the relatively simple formal statement of the model.

The algorithm's first inputs are s_B and s_W . The algorithm's next input is the principal's prior beliefs about the extent to which the fire alarm could benefit from making an untruthful statement. Let $q_{\text{lie}}(sq, o, F, t) = 1$ if $|(sq - F) - |o - F|| > t$ and 0 otherwise. The variable q_{lie} tells whether a fire alarm with ideal point f who observes o could find it profitable to make an untruthful statement. All else constant, the likelihood that $q_{\text{lie}} = 1$ is increasing in $|(sq - F) - |o - F||$, which is the maximum potential benefit from lying, and is decreasing in the magnitude of the expected penalty for lying.

If the principal knew F and o , she would know the value of q_{lie} . However, her information about F and o are limited to her knowledge of the distributions Γ and O . Let $Q_{\text{lie}}(\Gamma, sq, O, t) \in [0, 1]$ be the principal's prior belief about

the probability that the fire alarm could find it profitable to make an untruthful statement, where

$$Q_{\text{lie}}(\Gamma, sq, O, t) = \int \int q_{\text{lie}}(sq, o, F, t) dO' d\Gamma'.$$

Let $h(P, s_B, s_W, Q_{\text{lie}})$ be an exogenously determined correspondence that is everywhere nondecreasing in s_B and s_W and everywhere nonincreasing in Q_{lie} . Correspondence h denotes the principal's (common knowledge) expectation about the relationship between the content of the statement and the actual location of o . Let \underline{h} be an exogenously determined constant. We say that a principal of type p chooses to condition her inferences on the fire alarm's actions if and only if

$$h(P, s_B, s_W, Q_{\text{lie}}) > \underline{h}.$$

We have chosen to examine the case where this threshold is surpassed. Alternatively, the rule of thumb might dictate that the principal either discount or ignore information about the fire alarm. Fortunately, the case where the principal chooses to ignore this information is equivalent to the case where the fire alarm's entry costs are prohibitively high. In effect, we examine that case as well. The case where the principal discounts information in a systematic manner can be made equivalent to an analysis of the present model by exchanging current prior beliefs about the fire alarm's ideal point with relatively diffuse priors or by decreasing the value of t . Since the rule of thumb is solely a function of the common knowledge, we assume that the principal's inference technique is also common knowledge.

References

- Aberbach, Joel. 1990. *Keeping a Watchful Eye*. Washington, D.C.: Brookings.
- Austen-Smith, David. 1993a. "Interested Experts and Policy Advice: Multiple Referrals under Open Rule," *5 Games and Economic Behavior* 3–43.
- . 1993b. "Information and Influence: Lobbying for Agendas and Votes," *37 American Political Science Review* 799–833.
- Banks, Jeffrey S. 1991. *Signaling Games in Political Science*. New York: Harwood Academic Publishers.
- Crawford, Vincent, and Joel Sobel. 1982. "Strategic Information Transmission," *50 Econometrica* 1431–51.
- Dodd, Lawrence, and Richard Schott. 1979. *Congress and the Administrative State*. New York: Wiley.
- Ferejohn, John A. 1987. "The Structure of Agency Decision Processes," in M.D. McCubbins and T. Sullivan, eds., *Congress: Structure and Policy*. Cambridge: Cambridge University Press.
- Fiorina, Morris P. 1977. *Congress: Keystone of the Washington Establishment*. New Haven, Conn.: Yale University Press.
- Gerth, Hans H., and C. Wright Mills. 1946. *From Max Weber: Essays in Sociology*. New York: Oxford University Press.
- Gilligan, Thomas W., and Keith Krehbiel. 1987. "Collective Decision Making and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures," *3 Journal of Law, Economics & Organization* 145–93.
- Gilligan, Thomas W., and Keith Krehbiel. 1989. "Organization of Informative Committees by a Rational Legislature," *34 American Journal of Political Science* 531–64.

- Harris, Joseph. 1964. *Congressional Control of Administration*. Washington, D.C.: Brookings.
- Huntington, Samuel P. 1965. "Congressional Responses to the Twentieth-Century," in D. B. Truman, ed., *The Congress and America's Future*. New York: Prentice-Hall.
- Kiewiet, D. Roderick, and Mathew D. McCubbins. 1991. *The Logic of Delegation: Congressional Parties and the Appropriations Process*. Chicago: University of Chicago Press.
- Kreps, David M., and Robert Wilson. 1982. "Sequential Equilibria." *Econometrica* 50 863–894.
- Lowi, Theodore J. 1979. *The End of Liberalism: The Second Republic of the United States*. 2nd ed. New York: W. W. Norton.
- Lupia, Arthur, and Mathew D. McCubbins. 1993. "Designing Bureaucratic Accountability," forthcoming in *Law and Contemporary Problems*.
- Mazmanian, Daniel, and Jeanne Neinaber. 1979. *Can Organizations Change?* Washington, D.C.: Brookings.
- McCubbins, Mathew D., and Gregory W. Noble. 1993. In P. Cowhey and M. D. McCubbins, eds., *Structure and Policy in Japan and the United States*. Cambridge: Cambridge University Press.
- , Roger G. Noll, and Barry R. Weingast. 1987. "Administrative Procedures as an Instrument of Political Control," 3 *Journal of Law, Economics, & Organization* 243–77.
- , and Thomas Schwartz. 1984. "Police Patrols vs. Fire Alarms," 28 *American Journal of Political Science* 165–79.
- Myerson, Roger. 1979. "Incentive Compatibility and the Bargaining Problem." *Econometrica* 47: 61–73.
- Niskanen, William A. 1971. *Bureaucracy and Representative Government*. Chicago: Aldine-Atherton.
- Ogul, Morris S., and Bert A. Rockman. 1990. "Overseeing Oversight: New Departures and an Old Problem," 15 *Legislative Studies Quarterly* 5–24.
- Romer, Thomas, and Howard Rosenthal. 1978. "Political Resource Allocation, Controlled Agendas, and the Status Quo." *Public Choice* 33, 27–44.
- Spence, Michael. 1973. "Job Market Signaling," 87 *Quarterly Journal of Economics* 355–74.
- Weingast, Barry R., and Mark Moran. 1983. "Bureaucratic Discretion, or Congressional Control? Regulatory Policy Making by the Federal Trade Commission," 91 *Journal of Political Economy* 775–800.