

Stovepiping

Journal of Theoretical Politics
25(3) 388–411
©The Author(s) 2013
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI:10.1177/0951629813477277
jtp.sagepub.com



Sean Gailmard

Charles & Louise Travers Department of Political Science, University of California, Berkeley, CA, USA

John W Patty

Department of Political Science, Center for New Institutional Social Sciences, Washington University in Saint Louis, St Louis, MO, USA

Abstract

In hierarchical organizations lower-level agents can often censor the information that a higher-level principal has available to make a decision. We present a model of this interaction in which the principal can also access an independent source of unfiltered but lower-quality information besides that provided by the agent. This provision of outside information can be thought of as ‘stovepiping,’ the transmission of unfiltered information from analysts directly to decision-makers. Stovepiping can, in equilibrium, result in the agent passing along more information to the principal, precisely because the outsider’s information is of lower quality than the agent’s. But it can also lead the agent to ‘destroy’ information so that there is no basis for any policy change. Accordingly, stovepiping has countervailing effects on the principal’s utility. We discuss the comparative statics of equilibrium levels of information transmission with respect to the preferences of the principal, the outsider, and the probability that the outsider has access to the information held by the agent.

Keywords

Advice; information; transparency; hierarchies

... what the Bush people did was ‘dismantle the existing filtering process that for fifty years had been preventing the policymakers from getting bad information. They created stovepipes to get the information they wanted directly to the top leadership. Their position is that the professional bureaucracy is deliberately and maliciously keeping information from them. . . They always had information to back up their public claims, but it was often very bad information. . . ’. (Hersh, 2004, pp. 223–224, quoting Kenneth Pollack, a former National Security Council expert on Iraq)

Corresponding author:

John Patty, Department of Political Science, Center for New Institutional Social Sciences, Washington University in Saint Louis, 1 Brookings Drive, Campus Box 1063, St Louis, MO 63130, USA.

Email: jpatty@wustl.edu

Hersh's description of the decision making within the Bush Administration in the months leading up to the invasion of Iraq in 2003 is clearly disturbing in *ex post* terms: the certainty with which Donald Rumsfeld and his colleagues apparently believed that Saddam Hussein's regime possessed weapons of mass destruction contrasts starkly with the lack of evidence discovered after the military intervention. But this point merely highlights a potential argument in favor of political decision makers¹ possessing control over the processes by which they are presented with decision-relevant information. In institutional terms, such control was exercised by Rumsfeld himself through his redirection of the Northern Gulf Affairs Office. Retitled the Office of Special Plans (OSP) in 2002, this office reviewed information that had been collected through both public and confidential channels, although it collected little to no new intelligence on its own, as it was constituted as a policy group.²

The OSP was responsible for forwarding intelligence reports to the Secretary of Defense that were pertinent in its members' opinions to the planning for how best to deal with Iraq, particularly the possibility of direct military intervention. In this sense it illustrates *stovepiping*: an information transmission process in organizations in which decision makers obtain direct access to (a subset of) the same data that their bureaucratic subordinates possess, but without the data filtering and processing of data that subordinate analysts typically perform. Dreyfuss and Vest (2004) provides a direct illustration of how the OSP operated as a *stovepipe*:

raw intelligence is usually subject to a thorough vetting process, tracked, verified, and checked by intelligence professionals. But not at OSP—the material that it produced found its way directly into speeches by Bush, Cheney, and other officials. (Dreyfuss and Vest, 2004)

In the case of the debate leading up to the invasion of Iraq, it was clear that a stovepipe did not lack for smoke. As an unnamed former CIA official is quoted by Hersh (2003), '[w]e collect so much stuff that you can find anything you want.' After the Department of Defense's Inspector General had proffered their postmortem report about the handling of intelligence in the lead-up to the invasion of Iraq, Senator Jay Rockefeller (D, WV), the then Chairman of the US Senate Select Committee on Intelligence, described the activities of the OUSD(P) in the following very clear terms:

[i]ndividuals in that office produced and disseminated intelligence products outside of the regular intelligence channels. These intelligence products were inconsistent with the consensus judgments of the Intelligence Community. This office did this without coordinating with the Intelligence Community and as a result policy-makers received distorted intelligence. (Senator Jay Rockefeller, press release, 8 February 2007)

Of course, while captivating, the events of 2002–2003 are not the first time that the importance of information in determining public policy has been observed. Indeed, it has long been acknowledged that a critical source of bureaucratic power is a monopoly over the provision of decision-relevant information to a principal (Weber, 1946; Niskanen, 1971). For instance, in regulatory agencies, line bureaucrats acquire information from notice and comment proceedings, investigations of complaints, petitions for action, offeror provisions, incidental learning through the enforcement of existing regulations, analytical studies of regulated industries, discovery during judicial review and enforcement of agency actions, reviews by the Office of Management and Budget (OMB), audits by the Government Accountability Office (GAO), and congressional hearings. In policy-making agencies, staff bureaucrats acquire information about policy options and their

consequences from field contacts with affected parties or experts in the relevant scientific community, analysis of previously implemented policies or policies in other jurisdictions, and the received wisdom in the agency itself (even ‘groupthink,’ in pathological cases). These bits and pieces add up to a mountain of data that the agency staff filters, refines, and processes into information. This information, in turn, is transmitted to higher levels in the agency to make a decision: pursuit of an enforcement action, design of a new regulation, creation of a guidance document, or simply choice of a new policy or strategy.

Monopoly control over the gathering and processing of this data allows bureaucratic agents to exercise agenda setting power (Bendor et al., 1985). Bureaucrats can censor some of this information in the process of transmitting to higher-ups if they consider it low quality, not worthy of action, or at odds with options they wish the agency to pursue (whether for reasons of ideology, career building, status quo bias, avoiding new work, etc.). In this way, bureaucrats can use monopoly status over information provision to define the set of options available for action by higher-ups, and tilt that set in a direction they prefer.

Stovepiping is one way to challenge this control over information by analysts and line agents. The stovepipe effectively reaches from the unfiltered data directly and uninterrupted to the final decision maker(s). Thus, stovepiping inserts an agent with authority to pass along data ‘around’ another agent who gathers and processes it. The stovepiped data, while not refined and aggregated by the line agents’ expertise, is uncensored.³ The stovepipe in our model has two defining characteristics: first, the data it provides are necessarily not processed by the line agency. An underlying assumption of our analysis below is that this reduces the quality of the information, or put differently, that the line agency has useful analytic capacity.⁴ Second, the stovepipe shares some interest with the principal that the line agency does not share. This shared interest arises in stovepiping because the stovepipe wants the principal to take action in cases where the line agency does not; that is the reason for (to mix metaphors a bit) the stovepipe to do an ‘end run’ around the line agency in the first place.

From the standpoint of hierarchical control and accountability in bureaucracy, data filtering and censoring are problematic. But from the standpoint of ensuring that policy decisions are grounded in high-quality information, stovepiping is problematic as well. The literature on information provision in hierarchical organizations does not offer conceptual guidance in evaluating this dilemma, so in this article we provide a formal model to do so.⁵ The issue is obviously politically important, as the OSP vignette shows, although our purpose is not to offer an ‘explanation’ of any actor’s behavior in this case.⁶

Thus, our plan in this article is two-fold. First, we construct a simple model of line agency information transmission to the principal. The model considers the effect of stovepiping on both the quality of decision making (which the model makes precise) and the incentives of both the line agency and the outsider to submit information to the principal. Second, we address the effects of stovepiping on the utility of the principal. We demonstrate that these alternative information channels have two effects on the provision of decision-relevant information to the principal: a direct effect whereby the stovepipe reveals information the line agency does not reveal, and an indirect effect whereby the presence of a stovepipe leads the line agency to censor its reports to the principal less aggressively. At the same time, these alternative information channels are not always

beneficial for the principal: stovepiping can reduce the line agency's incentives to process information in the first place, and can give line agency incentives to destroy evidence before the stovepipe (or agency) can act on it.

I. A strategic model of stovepiping

In this section, we develop a strategic model of information transmission within a simple policy-making hierarchy. While this model is game theoretic, we omit some aspects of a full description of the game for purposes of clarifying the presentation of the interesting results.

Intuitively, the model captures a process in which a line agent A is formally charged with assembling information about the value of altering policy from a status quo, and can decide at its discretion to forward its information to the principal P who will make the final decision about whether to make policy, denoted by $a_P = 1$, or not, denoted by $a_P = 0$. We assume that A is more conservative than P in the sense that there are some reports on which P would want to act but are deemed unworthy of action by A . This is not particularly restrictive given P 's power to choose a_P : if A were less conservative, i.e. there are reports on which A would act but P would find unpersuasive as the basis of making new policy, then P could simply ignore them.⁷

With positive probability, an Outsider O also observes the same information observed by the agent about the value of making new policy (i.e. choosing $a_P = 1$). If O observes this information, it can choose to forward it to P regardless of A 's action. The Outsider is, for the time being, assumed to be less conservative (i.e. more likely to favor making new policy, $a_P = 1$) than both P and A in the sense that if either P or A believes that a report warrants making new policy, then O necessarily finds this information persuasive as well. In addition, A has substantive capacity to implement policy that brings it systematically closer to P 's desired outcome, but only if A involves itself by recommending a policy change; otherwise, P must pursue policy change without A 's substantive capacity. One way to think of this is that policy can result in a Type I error, taking action when it is inappropriate, and actors are ordered in terms of the cost they attach to Type I errors.⁸ While A attaches the greatest cost to Type I error, all actors also know that, given that policy change will take place, the probability of Type I error is minimized when A applies its substantive capacity to implementing the new policy.

For instance, A could represent a field agent in the FBI who observes untoward activity among suspected terrorist organizations, but is concerned that intervening against the suspects prematurely could undermine a more intricate case he or she is working on. On the other hand, P is less concerned about line agents' painstaking case development, and more concerned to 'do something' in the public eye about terrorism and ensure that no attacks happen on her watch; besides, if new action is taken inappropriately, P can always blame it on overzealous field agents. The outside informant O in this example could represent agents from another unit (e.g. another division of the FBI, the CIA, or local law enforcement) who believe the agents A are excessively lax or timid. Numerous other examples are also possible, e.g. with A representing the National Security Agency, P the president, and O a military commander in the field.

Alternatively, A could represent drug reviewers at the Food and Drug Administration who observe the quality of scientific protocols used to evaluate the efficacy and safety of

new drugs proposed by industry, while P represents a hierarchical principal in the FDA with formal authority to approve new drugs for marketing. Due to their acute concern for the scientific reputation of their agency (Carpenter, 2010), drug reviewers may be inclined against approving new drugs for marketing if they are supported by lax protocols. On the other hand, P may be more attuned to political pressure to break the ‘drug lag’ and bring new drugs to market. The outsider O could then represent other informed groups such as pharmaceutical research associations, patient advocacy groups, or government health researchers (e.g. the National Cancer Institute) who are more interested than FDA drug reviewers in seeing new drugs approved for marketing. While these outside groups have information about drug safety and efficacy to present to the principal, P still benefits from A ’s involvement; e.g. A may have specific expertise in crafting a drug label to convey intricate dosage information and side effects to physicians.

Structure of play. The line agent A observes the report r with probability $\pi_A \in (0, 1]$ and the Outsider O observes the report r with probability $\pi_O \in (0, 1]$. Whether A or O observes the report is independent of whether the other player observes the report, and neither player observes whether the other player observes the report. The probabilities π_A and π_O are common knowledge. Upon observing r , player $i \in \{A, O\}$ chooses whether to forward the report to the Principal: $a_i = 1$ if the report is forwarded by i and $a_i = 0$ otherwise. After A and O have made their decisions the Principal decides whether to promulgate a new policy ($a_P = 1$) or not ($a_P = 0$).

When $\pi_A < 1$, standard bureaucratic redundancy arguments⁹ imply that the Outsider’s presence is beneficial for the principal, P . Specifically, the Outsider’s presence raises the probability that at least one agent observes the report. As will become clear below, however, this is not the primary value of O to P ; rather O is valuable because of the changes it induces in A ’s reporting behavior. Note that we formally model A ’s and O ’s decisions to submit the report to P as simultaneous.

It may seem that stovepiping would be more naturally captured in a sequential game in which O observes A ’s report a_A , and then chooses its report a_O . However, it can be shown that the sequential move version when O follows A will generate the same subgame perfect Nash equilibria. In particular, O would report exactly the same set of reports in both versions of the game and, more importantly, O ’s decision is independent of A ’s action (and, hence, strategy) in *both* game forms.¹⁰

Information. The report r consists of one number, $q_r > 1$, which represents the quality of policy implementation based upon the report.¹¹ Report quality q_r is observed by the Principal if and only if the report is forwarded by either the Agency or the Outsider, and the Agency and/or Outsider can forward the report only if they observe it. Information is ‘hard’ in the sense that if either A or O forwards the report, they must do so truthfully; they cannot lie about its content. This is a straightforward way to represent either relational/reputational costs, or formal legal sanctions (as might apply to A), associated with lying about q_r ; essentially we are assuming that these costs are high enough to deter lying even if it is caught with very small probability. While A and O must tell the truth, they need not tell the ‘whole truth’: they may send no signal at all. This captures a situation in which information submitted to the agency comes with a paper trail; agents have some

discretion to decide what information (if any) to pass up the hierarchy, but passing up false information is formally prohibited.

Note that our model is not a ‘cheap-talk’ signaling model, in that the report r can be revealed credibly. A better description of the model is as one of ‘informational gate-keeping’ with multiple gatekeepers. The state of nature is r revealed with some positive probability to the agents who then independently decide whether to also reveal it to the principal. The agents have divergent preferences, in this case over the minimal quality of information required for the agent in question to prefer that the Principal make new policy, and divergent abilities. As we discuss below, the conservative informational gatekeeper (in this case the Agency, A) is also more effective at utilizing the information contained in any given report. Accordingly, one can conceive of the Agency as inherently possessing ‘more information’ than the Outsider.

The ‘policy-setting’ process. If $a_P = 0$, an exogenous status quo with known utilities remains in effect. If $a_P = 1$, P makes new policy with *precision* (or *quality*) equal to q_r . ‘Precision’ is assumed here to be a commonly valued characteristic of policy-making: high precision is preferred, *ceteris paribus*, by all three agents.¹² The precision of the report r is drawn from a commonly known distribution with cumulative distribution function (CDF) F . We assume that the observation of the report by the Agent and the Outsider (occurring with probabilities of π_A and π_O , respectively) are each independent of q_r .¹³

Payoffs. The Principal’s payoff function, u_P , is based on a_O , a_P , and q_r as follows:

$$u_P(a_O, a_P, q_r) = \begin{cases} b_P - \frac{1}{q_r} & \text{if } a_P = 1 \text{ and } a_A = 1, \\ b_P - \frac{\tau}{q_r} & \text{if } a_P = 1, a_O = 1, \text{ and } a_A = 0 \\ -x_P & \text{if } a_P = 0. \end{cases}$$

Thus, the first line covers cases where P makes new policy with A ’s input, the second line where P makes new policy with O ’s input but not A ’s, and the third where P keeps the status quo. The parameter $b_P \geq 0$ measures the Principal’s desire to promulgate new policies, *ceteris paribus*. For instance, the principal may desire to ‘leave her mark’ on policy or show some degree of activity to yet a higher-level (unmodeled) principal. The parameter $x_P \in (0, 1)$ represents the Principal’s payoff loss from the status quo policy. In our framework, the cost of promulgating a new policy ($a_P = 1$) stems from uncertainty about policy outcomes (which is inversely proportional to q_r). This could be described as the potential for blame for policy outcomes that would be (perhaps unfairly) attributed to the new regulation. Alternatively it could simply be standard spatial distance between a desired target outcome and the realized outcome.¹⁴ The parameter $\tau > 1$ is an exogenous and commonly known value that measures the increased uncertainty (i.e. lack of precision) resulting from policymaking *without* the guidance of the Agency. We assume that q_r and the observations of the report by the Agency and Outsider are all independently distributed.¹⁵

The Outsider’s payoff function, u_O , is similarly based on a_O , a_P , and q_r :

$$u_O(a_O, a_P, q_r) = \begin{cases} \sigma b_P - \frac{1}{q_r} & \text{if } a_P = 1 \text{ and } a_A = 1, \\ \sigma b_P - \frac{\tau}{q_r} & \text{if } a_P = 1, a_O = 1, \text{ and } a_A = 0 \\ -x_O & \text{if } a_P = 0. \end{cases}$$

The parameter σ measures the alignment between the Outsider’s and Principal’s direct preferences for ‘new policy making’ and $x_O \in (0, 1)$ represents the Outsider’s payoff loss from the status quo policy. The Outsider and the Principal effectively have identical preferences over policy outcomes and, if $\sigma \geq 1$, the Outsider’s preference for new policymaking are at least as strong as the Principal’s. Conversely, when $\sigma \in (0, 1)$, the Outsider has a *ceteris paribus* preference for new policy making, but this preference is weaker than the Principal’s.

Finally, the Agency’s payoff function, u_A , is based on a_O , a_P , and q_r as follows:

$$u_A(a_P, q_r) = \begin{cases} -\frac{1}{q_r} & \text{if } a_P = 1 \text{ and } a_A = 1, \\ -\frac{\tau}{q_r} & \text{if } a_P = 1, a_O = 1, \text{ and } a_A = 0 \\ -x_A & \text{if } a_P = 0. \end{cases}$$

where $x_A \in (0, 1)$ represents the Agency’s payoff loss from the status quo policy.

Several features of the payoffs are worth mentioning due to the role they play in our results. First, all players share preferences over implementation quality: higher values of q_r uniformly make the players better off conditional on the Principal making policy ($a_P = 1$). This holds regardless of whether the Agency, Outsider, or both forward the report to the Principal. This assumption is critical for our analysis, and while credible in many applications, it probably does not hold universally. For example, if q_r represents a sort of reduced form variance over policy outcomes, an actor might prefer greater variance if it makes it more likely for policy to exceed some desired threshold. That is, some actors might have conventional preferences for lower variance, while others might prefer greater variance because it raises the chance of policy exceeding a critical threshold, thus increasing the option value of enacting a policy change. In that case, players would not have common preferences over q_r , and the subsequent analysis would generally not hold.

Second, note that the model appears to represent divergence in the players’ preferences in two ways: they may have different preferences over whether to make new policy at all (i.e. $\sigma b_P > 0$) and/or they may receive different relative payoffs from the status quo (for example, $b_P = 0$ but $x_A \neq 0$). This difference is one in appearance only, as the action space is particularly coarse. In instrumental terms, policy is either changed or it is not. Accordingly, there is a duality built into the model’s parameterization: increasing x_P is equivalent to increasing b_P .¹⁶ However, the representation allows us to frame the comparative statics in direct terms so that the presentation of the results is clearer.

Third, no player cares about the quality of the report, q_r , unless the Principal chooses to make new policy ($a_P = 1$). This assumption creates an incentive for all players to defer to the status quo policy unless the report has high enough quality. Thus, although the players agree that precision in policy implementation (i.e. a high-quality report) is desirable, there is important heterogeneity in their motivations through the fact that they

will disagree over the cost of Type I errors or, in other words, what level of quality is sufficiently high to justify changing the status quo policy.

Fourth, and finally, we assume throughout that for each player $i \in \{A, O, P\}$, there is a sufficiently high level of quality such that player i would be made strictly better off by new policy being made (i.e. such that $a_P = 1$ is strictly preferred by player i to $a_P = 0$). Thus, Type II errors (failing to change policy when it is beneficial relative to the status quo) are a potential concern to all players, as are Type I errors (changing policy when it is not beneficial). Furthermore, we assume that there are sufficiently low-quality reports that the Agency would prefer not to make new policy.

2. Analysis

With the payoff functions in hand, it is simple to define the following sequentially rational policy making strategy for the Principal. Along with the content of the report r , q_r , there are essentially three cases to consider (one of which is invariant to the content of the report because the Principal does not observe it). These are described below.

- *Policy making with ‘everyone in the room.’* Given a report r with quality q forwarded by the Agency ($a_A = 1$), the Principal’s optimal policy choice is

$$a_P^*(q|a_A = 1, a_O = 0) = a_P^*(q|a_A = 1, a_O = 1) = \begin{cases} 1 & \text{if } b_P + x_P \geq 1/q, \\ 0 & \text{otherwise.} \end{cases}$$

- *Policy making with the agency absent.* Given a report r with quality q not forwarded by the Agency ($a_A = 0$) but forwarded by the Outsider ($a_O = 1$), the Principal’s optimal policy choice is

$$a_P^*(q|a_A = 0, a_O = 1) = \begin{cases} 1 & \text{if } b_P + x_P \geq \tau/q, \\ 0 & \text{otherwise.} \end{cases}$$

- *Policy making with no report.* Given a report r with quality q not forwarded by either the Agency or by the Outsider ($a_A = a_O = 0$), the Principal’s optimal policy choice is

$$a_P^*(q|a_A = 0, a_O = 0) = \begin{cases} 1 & \text{if } b_P + x_P > 1 \\ 0 & \text{otherwise.} \end{cases}$$

Before continuing, it is useful to note the nature of the Principal’s possible responses to a report: he either ‘pulls the trigger’ and makes new policy or he effectively ignores the report. This stark approach is useful because it allows us to focus on the interesting aspects of the informational structure.¹⁷ Specifically, the analysis here is of information handling *within* the bureaucracy.

Combined with the binary nature of policy making and the fact that reports can be well ordered in terms of their quality, $q(r)$, this initial analysis of the Principal’s incentives allows us to describe any player’s *ex ante* preferences about policy making in relation to any other’s. Accordingly, and throughout, we use the terms ‘conservative’ and ‘aggressive’ to describe a player’s policy making preferences in relation to the principal’s preferences.

Definition 1. A player i is conservative (relative to P) if the set of reports that i prefers $a_P = 1$ to $a_P = 0$ is a strict subset of the set of reports such that $a_P^*(q|a_A = 1, a_O = 1) = 1$. Conversely, a player i is aggressive (relative to P) if the set of reports that i prefers $a_P = 1$ to $a_P = 0$ is a strict superset of the set of reports such that $a_P^*(q|a_A = 1, a_O = 1) = 1$.

It is straightforward to verify the following facts:

1. the Agency, A , is conservative;
2. the Outsider, O , is conservative if $\sigma < 1$; and
3. the Outsider, O , is aggressive if $\sigma > 1$.

The responsive principal assumption. The following assumption is a useful baseline to rule out uninteresting cases.

Assumption 1. The Principal prefers leaving policy unchanged in the absence of a report: $b_P + x_P \leq 1$.

Without this assumption, there is never a reason (presuming sequential rationality by the Principal) for either the Agent or the Outsider to withhold a report. Even more specifically, when this assumption is violated the only reason for P to listen to a stovepipe is essentially to insure against the possibility that the Agency does not receive a report (which occurs with probability $1 - \pi_A$). In such cases, the stovepipe benefits both the Principal and the Agency in a standard redundancy fashion. While this is of course an empirically relevant consideration in real-world organizations, note that its effects are not the result of any interesting strategic interaction between any of the players.¹⁸

Note that the Outsider’s expected payoff is independent of his or her choice to report (a_O) if the Agency reports. Thus, the Outsider’s expected payoff from $a_O = 1$, conditional on $a_A = 0, q_r$, and the Principal’s optimal strategy as described above, is¹⁹

$$V_O(a_O = 1; q) = \begin{cases} \sigma b_P - \tau/q & \text{if } b_P + x_P \geq \tau/q, \\ -x_O & \text{otherwise.} \end{cases} \tag{1}$$

It is without loss of generality to presume that the Outsider reports only if the Principal would act upon the report (this is always known by the Outsider conditional on observing r , since τ, b_P , and x_P are all common knowledge and the Outsider knows q_r). This strategy is written formally as

$$a_O^*(q) = \begin{cases} 1 & \text{if } \min[b_P + x_P, \sigma b_P + x_O] \geq \tau/q \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

In equilibrium, then, the Agency’s strategic incentives are derived using (2), which leads to the following expected payoff function for the Agency:²⁰

$$V_A(a_A|q, \pi_O) = \begin{cases} -1/q & \text{if } b_P + x_P \geq 1/q \text{ and } a_A = 1 \\ -\frac{\pi_O \tau}{q} - (1 - \pi_O)x_A & \text{if } \min[b_P + x_P, \sigma b_P + x_O] \geq \tau/q \text{ and } a_A = 0 \\ -x_A & \text{otherwise.} \end{cases} \tag{3}$$

Utilizing Equation (3), the following straightforward derivations yield the optimal strategy for the Agency.

If $\min[b_P + x_P, \sigma b_P + x_O] \geq \tau/q$, then the Agency should report whenever

$$-1/q \geq -\frac{\pi_O \tau}{q} - (1 - \pi_O)x_A,$$

$$q \geq \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A},$$

and if $b_P + x_P \geq 1/q$, the Agency should report whenever $q \geq 1/x_A$. Indeed, $a_A = 1$ is a dominant action for any report with $q \geq 1/x_A$.²¹ Accordingly, the Agency’s optimal strategy can be written as follows:

$$a_A^* = \begin{cases} 1 & \text{if } q \geq 1/x_A, \\ 1 & \text{if } q \geq \max \left[\tau / \min[b_P + x_P, \sigma b_P + x_O], \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} \right], \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

We are now in a position to examine the comparative statics of the Agency’s optimal behavior. There are two cases of interest, depending on whether the Outsider is more or less conservative than the Principal. If $b_P + x_P < \sigma b_P + x_O$, then the Principal is more conservative than the Outsider and, conversely, if $b_P + x_P > \sigma b_P + x_O$, then the Outsider is more conservative than the Principal. We label these two cases ‘Aggressive Outsider’ and ‘Conservative Outsider,’ respectively.

To make the presentation of the remaining results parsimonious, we assume throughout what follows that q_r is drawn from a continuously differentiable CDF $F(q)$ (with probability density function f) possessing positive support for all $q > 1$.

Assumption 2 (Higher-quality reports always possible). *For all $q > 1, f(q) > 0$.*

With Assumption 2 in hand, we can now proceed to answer several questions within the model. First, under what circumstances does the presence of the Outsider affect the Agency’s behavior? Second, what are the determinants of the proportion of reports that are forwarded by the Agency in response to the presence of the Outsider?

When does stovepiping alter the agency’s behavior? Equation (4) clearly indicates a necessary condition for the Outsider’s presence to have any effect on the number of reports forwarded by the Agency. Specifically, since $1/x_A > \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A}$ for all $\tau > 1, \pi_O > 0$, and $x_A > 0$, the following must hold for the Agency to forward reports in the presence of the Outsider that would not be forwarded if the Outsider were not present:

$$1/x_A > \tau / \min[b_P + x_P, \sigma b_P + x_O]. \quad (5)$$

The following proposition summarizes the implications of Inequality (5).²²

Proposition 1. *The Outsider has an effect on the optimal reporting behavior of the Agency if:*

1. the Agency is sufficiently satisfied with the status quo policy (x_A is close enough to zero);
2. the Principal is sufficiently interested in making new policy (b_P is large enough);
or
3. O and P are both sufficiently displeased by the status quo (x_P and x_O are large enough).

The first conclusion of Proposition 1 is intuitive: when the agency is sufficiently satisfied with the status quo policy, the presence of a stovepipe will lead to an increase in the number of reports that are forwarded to the Principal simply because the choice of not forwarding the report to the Principal is no longer sufficient to guarantee that the status quo policy will remain in effect. So long as the Outsider is less efficient at policy implementation than the Agency (i.e. $\tau > 1$), the installation of a stovepipe will lead to direct reduction in the Agency’s expected payoff from withholding reports from the Principal.

The second and third conclusions of Proposition 1 are similar to the first conclusion but rely on a slightly different foundation. Specifically, as the Principal and the Outsider both become increasingly displeased with the status quo and/or the Principal is sufficiently interested in making new policy for the sake of making new policy, the Agency realizes that the minimal quality level that the Principal and Outsider will find sufficient cause to make new policy is decreasing as well. Since the Agency has an incentive to be ‘in the room’ whenever it knows that a new policy is being made, the increasing willingness of the Principal to make policy increases the willingness of the Agency to forward reports to the Principal.

Proposition 1 compares the frequency of reporting under two starkly different institutional structures (i.e. in the presence of versus the absence of any Outsider). We now turn to the question of comparing institutional structures that, in general, differ on the margin: holding the presence of the stovepipe constant, we can examine the effect of changes in the preference and technology parameters of the model.

Comparative statics of the Agency’s behavior. In order for the comparative statics of the Agency’s behavior to depend on the structure of the Outsider’s involvement, the Outsider need be sufficiently competent (i.e. that τ not be too large). Otherwise, the minimum level of quality that the Agency need worry about depends solely on the Principal’s preferences.²³ The maximum level of relative inefficiency of implementation in the absence of the Agency is given by the following inequality:

$$\begin{aligned}
 0 &< \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} - \tau / (b_P + x_P), \\
 \tau &< \frac{b_P + x_P}{\pi_O(b_P + x_P) + (1 - \pi_O)x_A}.
 \end{aligned}
 \tag{6}$$

Aggressive Outsider. Supposing that $b_P + x_P < \sigma b_P + x_O$ and that Inequality (6) is satisfied, the reports of interest are those for which

$$q \in \left(\frac{1 - \pi_O \tau}{(1 - \pi_O)x_A}, \frac{1}{x_A} \right).
 \tag{7}$$

This is the set of reports that, in equilibrium, the Outsider’s presence causes the Agency to forward to the Principal when the Outsider is sufficiently competent that Inequality (6) is satisfied. This set is no-empty for all $\pi_O > 0$ and vanishes when $\pi_O = 0$. Invoking Assumption 2 then yields the following comparative static.

Proposition 2. *When the Outsider is more aggressive than the Principal and the Outsider is not too incompetent (i.e. inequality (6) is satisfied), the ex ante probability that the Agency forwards a report to the Principal is:*

1. increasing with respect to the probability that the Outsider observes the report, π_O ’ and
2. increasing with respect to the Outsider’s relative incompetence, τ .

Supposing that $b_P + x_P < \sigma b_P + x_O$ and that inequality (6) is *not* satisfied, the reports of interest are those for which

$$q \in \left(\frac{\tau}{b_P + x_P}, \frac{1}{x_A} \right). \tag{8}$$

As above, this is the set of reports that, in equilibrium, the Outsider’s presence causes the Agency to forward to the Principal when the Outsider is sufficiently incompetent that inequality (6) is not satisfied. This set is clearly empty for sufficiently large values of τ : in such cases the Outsider’s presence has no effect on the Agency’s reporting behavior. Invoking Assumption 2 then yields the following comparative static.

Proposition 3. *When the Outsider is more aggressive than the Principal and the Outsider is sufficiently incompetent (i.e. inequality (6) is not satisfied), the ex ante probability that the Agency forwards a report to the Principal is decreasing with respect to the Outsider’s relative incompetence, τ .*

The effect of the Outsider’s relative inefficiency (τ) is displayed in Figure 1. The figures graphs the minimal quality level that the Agency would forward to the Principal as a function of τ . All three possible regions are illustrated. When τ is low (i.e. close to 1), the effect of increasing τ is to decrease the minimal quality level (increasing the frequency of reports being forwarded), as described in Proposition 2. For intermediate values of τ , increasing τ increases the minimal quality level that is forwarded (decreasing the frequency of reports being forwarded), as described in Proposition 3. Finally, for sufficiently large values of τ , the Agency’s behavior is insensitive to further increases in τ because the Principal will never act upon a report unless it is forwarded by the Agency.

Before continuing to the case of a conservative Outsider, note that, regardless of the competence of the Outsider (i.e. the value of τ), the set of reports that the Agency will forward to the Principal is, quite intuitively, increasing in the Agency’s dissatisfaction with the status quo (i.e. x_A).

Conservative Outsider. When $b_P + x_P > \sigma b_P + x_O$, the analog to inequality (6) is

$$\tau < \frac{\sigma b_P + x_O}{\pi_O(\sigma b_P + x_O) + (1 - \pi_O)x_A}. \tag{9}$$

Also, when $b_P + x_P > \sigma b_P + x_O$, inequality (9) also identifies when the agency’s behavior is affected directly by the preferences of the Outsider. Thus, the following Propositions are immediate.




Figure 1. The effect of outsider policy making efficiency on agency reporting.

Proposition 4. *When the Outsider is more conservative than the Principal and the Outsider is not too incompetent (i.e. inequality (9) is satisfied), the ex ante probability that the Agency forwards a report to the Principal is:*

1. increasing with respect to the probability that the Outsider observes the report, π_O ; and
2. increasing with respect to the Outsider’s relative incompetence, τ .

The intuition behind Proposition 4 can be understood as follows. When the Outsider is not ‘too incompetent,’ there is a positive probability that the Principal will receive and act upon a report forwarded only by the Outsider, O . Thus, in this region, increases in the likelihood of the report being received by the Outsider commensurately increase the downside risk the Agency incurs by withholding a report and accepting that policy might be made without the Agency in the room by increasing the probability of such an event. Analogously, in this region, decreases in the Outsider’s competence increase the downside risk for the Agency of withholding a report by increasing the *ex post* cost of such an occurrence. The next proposition deals with situations in which the Outsider is sufficiently incompetent that inequality (9) is not satisfied.

Proposition 5. *When the Outsider is more conservative than the Principal and the Outsider is sufficiently incompetent (i.e. inequality (9) is not satisfied), the ex ante probability that the Agency forwards a report to the Principal is:*

1. decreasing with respect to the Outsider’s relative incompetence, τ ;
2. increasing with respect to the similarity of the Principal and Outsider’s preferences, σ .

The conclusions of Proposition 5 can be understood in the following way. When the Outsider is sufficiently incompetent, there will exist reports of low quality that the

Agency would forward *if* the Agency believed that the Principal would act upon them in the Agency's, *but* the Principal would not be rational to act upon them if the Agency is not in the room (because the Outsider is too incompetent to justify action over inaction in the Agency's absence). The first conclusion of the proposition follows from the fact that this set of reports grows as τ increases (i.e. as the Outsider becomes less competent). As for the second conclusion, we are considering the case of a conservative Outsider: there are some reports that the Outsider will not forward to the Principal that the Principal would act upon even if the Outsider were the only one to forward them. Accordingly, this set of reports decreases, implying that the probability that the Outsider forwards a report to the Principal increases, as σ increases, as this makes the Outsider less conservative (i.e. more aggressive). Thus, the Agency's downside risk from withholding a report is increasing as σ increases.

Implications for institutional design. Before continuing on to extend the model so as to incorporate endogenous information collection, consider the implications of the model so far for the 'institutional design' problem faced by the Principal if he or she is allowed to select an Outsider. Specifically and for simplicity, presume that the Principal can choose τ . That is, suppose that the Principal has a set of possible outsiders to select from, each with different levels of implementation ability (i.e. competence). Furthermore, and in line with one of the two endogenous information extensions we consider later in the article, modify the game above so that the Outsider obtains the report *only if* the Agency also obtains the report.²⁴

The combination of Propositions 4 and 5 suggests that the Principal has an interior, optimal level of incompetence. Specifically, given the presumption that the Outsider receives the report only if the Agency also receives it *and* that the probability of the Agency receiving the report is independent of the chosen level of τ ,²⁵ the level of incompetence of the Outsider that maximizes the Principal's equilibrium expected payoff is strictly greater than $\tau = 1$ (i.e. less competent than the Agency). In this extension/modification of the model, the Principal has an incentive to leverage his or her the ability to commit himself to inefficient implementation, which the Agency does not like, as a mechanism to force the Agency's hand on reports of marginal quality that the Agency would not forward if it did not fear the cost of 'not being in the room.' This cost is increasing in τ , and Proposition 5 indicates that any optimal level of τ from the Principal's (*ex ante*) perspective must satisfy inequality (9).

3. Stovepipes and endogenous information collection

The fundamental strategic effect of stovepiping is to increase the Agent's incentive to report what it knows. It may therefore seem that P is uniformly better off with a stovepipe than without. However, this intuition ignores the effects of stovepiping on the Agent's incentives to acquire information in the first place. In this subsection we consider two extensions that partially endogenize the information conditions in the game. Both of these extensions show that the simple intuition that stovepipes are beneficial for the Principal is misguided when information is endogenous.²⁶

Throughout this section we make the simplifying normalization that $x_P = x_O = 0$: the Principal and Outsider’s interests in making new policy are measured completely by b_P .²⁷

3.1. Costly information processing in the agency

The first extension relates to information acquisition by the Agent when the Outsider’s information is given. Just as the stovepipe can induce the Agency to forward reports that it would not otherwise forward, the stovepipe can induce the Agency to acquire more information in the first place. At the same time, if the stovepipe does not impose too large a cost on the Agency when it does not report, the stovepipe can actually reduce the Agency’s incentive to acquire information about q .

Specifically, we suppose now that the Agency can choose whether to discover the report r ’s quality, q_r . As above, the Agency believes that the distribution of the quality is characterized by a cumulative distribution function, F . If the Agency chooses to discover the quality q_r , then it incurs an opportunity cost denoted by $k > 0$. The Agency can then decide whether to forward the report to the Principal as in the analysis reported in the previous section. If the Agency choose not to discover the quality, then it cannot forward the report to the Principal.²⁸ Regardless of the Agency’s decision, the Outsider learns q with probability π_O , as in the earlier analysis. This assumption removes any strategic effects of the Agency’s decision on the likelihood that the Outsider will forward the information to the Principal.

Let $q^*(\tau, b_P, \sigma, \pi_O, x_A)$ denote the minimal quality level of the set of reports that the Agency would forward to the Principal. Recalling Equation (4) and substituting our assumption that $x_P = x_O = 0$, this threshold is given by

$$q^*(\tau, b_P, \sigma, \pi_O, x_A) \equiv \max \left[\frac{\tau}{\min[b_P, \sigma b_P]}, \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} \right], \tag{10}$$

which, whenever the context is clear, we will denote simply by q^* .

Aggressive Outsider: $\sigma \geq 1$. If the Outsider is at least as aggressive as the Principal in pursuit of policy change, then the reporting threshold presented in Equation (10) can be rewritten as

$$q^* \equiv \max \left[\frac{\tau}{b_P}, \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} \right].$$

Applying sequential rationality on the part of the Agency implies that the Agency’s expected payoff from discovering the report’s quality is

$$\bar{v}_A^1 = - \left[\int_1^{\tau/b_P} x_A dF + \int_{\tau/b_P}^{q^*} \left(\frac{\pi_O \tau}{q} + (1 - \pi_O)x_A \right) dF + \int_{q^*}^{\infty} \frac{1}{q} dF \right] - k,$$

and the Agency’s expected payoff from not discovering the report’s quality is

$$\bar{v}_A^0 = - \left[\int_1^{\tau/b_P} x_A dF + \int_{\tau/b_P}^{\infty} \left(\frac{\pi_O \tau}{q} + (1 - \pi_O)x_A \right) dF \right].$$

Thus, the Agency's net expected payoff from processing the report and discovering q_r can be written as

$$\Delta_A(\tau, b_P, \sigma, \pi_O, x_A, k) \equiv \int_{q^*}^{\infty} \left(\frac{\pi_O \tau - 1}{q} + (1 - \pi_O)x_A \right) dF - k. \tag{11}$$

Conservative Outsider: $\sigma < 1$. If the Outsider is no more aggressive than the Principal in pursuit of policy change, then the reporting threshold presented in Equation (10) can be rewritten as

$$q^* \equiv \max \left[\frac{\tau}{\sigma b_P}, \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} \right].$$

Note, however, that duplicating the derivations for the case of an aggressive Outsider results in the following expression of the Agency's net expected payoff from processing the report and discovering q_r :

$$\Delta_A(\tau, b_P, \sigma, \pi_O, x_A, k) \equiv \int_{q^*}^{\infty} \left(\frac{\pi_O \tau - 1}{q} + (1 - \pi_O)x_A \right) dF - k,$$

which is identical to Equation (11). All differences between the Agency's optimal behavior when the Outsider is aggressive and that when the Outsider is conservative are conveyed through the differences between q^* in the two cases (i.e. the differences due to the change in which constraint binds: τ/b_P or $\tau/(\sigma b_P)$) and, accordingly, we can omit the distinction between the two cases in the analysis that follows.

To be clear, note that if Δ_A is positive, the Agency's optimal response is to incur the cost k to discover q_r while, if Δ_A is negative, the Agency's optimal choice is to decline to process the report. Note that while the comparative statics of Δ_A depend crucially and unsurprisingly on those of q^* , there are now direct effects of π_O and x_A as well in the sense that these parameters affect the Agency's decision to acquire information even if they have no effect on the frequency with which the Agency would forward an already-processed report to the Principal. Let $k^*(\tau, b_P, \sigma, \pi_O, x_A)$ be defined as the opportunity cost that leaves the Agency indifferent with respect to whether it processes the report or not:

$$k^* \equiv \int_{q^*}^{\infty} \left(\frac{\pi_O \tau - 1}{q} + (1 - \pi_O)x_A \right) dF.$$

We can now directly describe the effect of various parameters on the Agency's willingness to gather and process information for potential forwarding to the Principal. That is the subject of the following proposition.

Proposition 6. *Suppose that $\pi_O < 1$. Then, the maximum cost the Agency should incur to discover the quality of the report, k^* , is*

1. *non-increasing in the minimal report quality the Agency would forward to the Principal (q^*);*
2. *increasing in the Agency's dissatisfaction with the status quo (x_A);*

3. may be increasing or decreasing in the probability that the Outsider observes the report (π_O).

Proof. The first conclusion requires a straightforward application of Leibniz’s rule:

$$\begin{aligned} \frac{\partial k^*}{\partial q^*} &= \frac{\partial}{\partial q^*} \int_{q^*}^{\infty} \left(\frac{\pi_O \tau - 1}{q} + (1 - \pi_O)x_A \right) dF, \\ &= - \left(\frac{\pi_O \tau - 1}{q^*} + (1 - \pi_O)x_A \right) f(q^*), \\ &= \begin{cases} 0 & \text{if } \frac{\tau}{\min[b_P, \sigma b_P]} < \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} \\ \left(\min[b_P, \sigma b_P] \frac{1 - \pi_O \tau}{\tau} - (1 - \pi_O)x_A \right) f(q^*) \leq 0 & \text{if } \frac{\tau}{\min[b_P, \sigma b_P]} \geq \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} \end{cases} \end{aligned} \tag{12}$$

where inequality (12) is implied by the following derivation:

$$\begin{aligned} \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A} &\leq \frac{\tau}{\min[b_P, \sigma b_P]}, \\ \left(\min[b_P, \sigma b_P] \frac{1 - \pi_O \tau}{\tau} - (1 - \pi_O)x_A \right) &\leq 0, \\ \left(\min[b_P, \sigma b_P] \frac{1 - \pi_O \tau}{\tau} - (1 - \pi_O)x_A \right) f(q^*) &\leq 0. \end{aligned}$$

The second conclusion follows from a more involved application of Leibniz’s rule:

$$\begin{aligned} \frac{\partial k^*}{\partial x_A} &= \frac{\partial}{\partial x_A} \int_{q^*}^{\infty} \left(\frac{\pi_O \tau - 1}{q} + (1 - \pi_O)x_A \right) dF, \\ &= \int_{q^*}^{\infty} \frac{\partial}{\partial x_A} \left(\frac{\pi_O \tau - 1}{q} + (1 - \pi_O)x_A \right) dF - \frac{\partial q^*}{\partial x_A} \cdot \left(\frac{\pi_O \tau - 1}{q^*} + (1 - \pi_O)x_A \right) f(q^*), \\ &= (1 - \pi_O)(1 - F(q^*)) > 0. \end{aligned} \tag{13}$$

The third conclusion follows from a similar application of Leibniz’s rule: letting

$$G_F(q^*) \equiv \int_{q^*}^{\infty} q^{-1} dF,$$

partial differentiation of k^* with respect to π_O yields the following:

$$\begin{aligned} \frac{\partial k^*}{\partial \pi_O} &= \frac{\partial}{\partial \pi_O} \int_{q^*}^{\infty} \left(\frac{\pi_O \tau - 1}{q} + (1 - \pi_O)x_A \right) dF, \\ &= \int_{q^*}^{\infty} \left(\frac{\tau}{q} - x_A \right) dF - \frac{\partial q^*}{\partial \pi_O} \cdot \left(\frac{\pi_O \tau - 1}{q^*} + (1 - \pi_O)x_A \right) f(q^*), \\ &= \tau G_F(q^*) - x_A(1 - F(q^*)). \end{aligned} \tag{14}$$

To establish the relevant conclusion of the proposition, note that $G_F(q^*) \in (0, 1/q^*)$ and $G'_F \leq 0$, but more importantly, note that q^* is independent of F , since q^* is based on an

interim calculation after q is realized. Thus, for any $\varepsilon > 0$ and $q^* > 1$, there exists a CDF F^+ satisfying our assumptions such that $\min[F^+(q^*), G_{F^+}(q^*)] > 1 - \varepsilon$ and there exists a CDF F^- satisfying our assumptions such that $\max[F^-(q^*), G_{F^-}(q^*)] < \varepsilon$. These facts imply that, fixing $x_A \in (0, 1)$ and $\tau > 1$, one can find a CDF F^+ such that the right-hand side of Equation (14) is positive, and a CDF F^- such that the right-hand side of Equation (14) is negative, as was to be shown. \square

Proposition 6 can be interpreted in the following way. The first conclusion, that the Agency is *less willing* to pay for reports when *those forwarded by the Agency* have higher expected quality is interesting and somewhat counterintuitive. Accordingly, it is illustrative to understand its basis. As the proof of the proposition lays out, the Agency's willingness to pay for information is invariant to q^* when q^* is determined 'by the Agency,' which is the case when $\frac{\tau}{\min[b_P, \sigma b_P]} < \frac{1 - \pi_O \tau}{(1 - \pi_O)x_A}$.

The second conclusion is unsurprising: as the Agency's disutility from the status quo policy increases, both the net payoff from forwarding any given report to the Principal, and the probability of observing a report of sufficiently high quality to warrant forwarding, increase as well. The complementary effects directly translate into an increased willingness on the Agency's part to pay to acquire information in the hopes of generating an actionable report.

Finally, the third conclusion is closer in spirit to an existence result than it is to a positive characterization: the effect of increased involvement by the Outsider (higher values of π_O) on the Agency's willingness to acquire information is ambiguous from an *ex ante* perspective. This ambiguity follows from the range of possible distributions of q . If few reports are of very high quality (i.e. a higher proportion of the mass of F is just less than x_A^{-1}), then increasing the probability that the Outsider gets the report will increase the Agency's incentive to screen the report as well. If, on the other hand, many reports were of very high quality (i.e. greater than x_A^{-1}), then the Outsider's involvement is, with a very high probability, a substitute for that of the Agency from the Agency's perspective. Accordingly, in such a case increasing the probability that the Outsider gets the report will decrease the Agency's incentive to invest in information acquisition.

3.2. Agency sabotage

In the previous extension we considered an institutional environment in which the stovepipe's access to the information was beyond the Agency's control. A closely related institutional possibility would allow the Agency to preclude both itself and the Outsider from the report. For instance, the Agency might choose to 'sabotage' the report by destroying the evidence before it can be processed. Specifically, we suppose now that before the game laid out in Section 1 is played, the Agency has a chance to permanently destroy the report (which is equivalent to imposing $c_A = c_O = 0$ in the formulation of Section 1). If the Agency chooses not to sabotage, then the game proceeds as in Section 1 whereas, if the Agency does sabotage, then *neither* the Agency nor the Outsider observes the report. We maintain our other assumptions about players' information and payoffs, so that the choice by the Agency to sabotage and destroy the report implies that the status quo policy will be retained, resulting in a payoff of $-x_A$ for the Agency.

Aggressive Outsider: $\sigma \geq 1$. If the Outsider is at least as aggressive as the Principal in pursuit of policy change, then we can utilize sequential rationality and the arguments Section 3.1 to write the Agency’s expected payoff from *not sabotaging* the report’s quality as

$$v_A^1 = - \left[\int_1^{\tau/b_P} x_A dF + \int_{\tau/b_P}^{q^*} \left(\frac{\pi_O \tau}{q} + (1 - \pi_O)x_A \right) dF + \int_{q^*}^{\infty} \frac{1}{q} dF \right],$$

and the Agency’s expected payoff from sabotaging the report as

$$v_A^0 = -x_A.$$

Continuing with the logic of Section 3.1, we can define the net payoff from sabotage as

$$\begin{aligned} \Sigma &= v_A^0 - v_A^1 \\ &= \pi_O \int_{\tau/b_P}^{q^*} \left(\frac{\tau}{q} - x_A \right) dF + \int_{q^*}^{\infty} \left(\frac{1}{q} - x_A \right) dF \end{aligned}$$

The first result is perhaps the most informative about the nature of the Agency’s strategic calculation: the Agency will sabotage *only* if it is the most conservative player (i.e. $q^* > \tau/b_P$). This is because, in essence, the Agency need fear the Outsider’s involvement only when there are reports such that the Outsider would forward them to the Principal, the Principal would act upon them with only the Outsider in the room, and the Agency would not forward them to the Principal in the Outsider’s absence.

Proposition 7. *The Agency will sabotage the information collection process only if it is more conservative than both the Principal and the Outsider. That is, sabotage is the optimal choice for the Agency only if $q^* \geq \max[\tau / \min[b_P, \sigma b_P]]$.*

Proposition 7 additionally provides a useful baseline in terms of deriving the comparative statics of the benefit from sabotage.²⁹ In particular, noting that

$$\frac{\partial q^*}{\partial \pi_O} = \begin{cases} \frac{1-\tau}{(1-\pi_O)^2 x_A} < 0 & \text{if } \frac{\tau}{\min[b_P, \sigma b_P]} < \frac{1-\pi_O \tau}{(1-\pi_O)x_A}, \\ 0 & \text{if } \frac{\tau}{\min[b_P, \sigma b_P]} \geq \frac{1-\pi_O \tau}{(1-\pi_O)x_A}, \end{cases}$$

direct and repeated applications of Leibniz’s rule yield the following:

$$\begin{aligned} \frac{\partial \Sigma}{\partial \pi_O} &= \begin{cases} \int_{\tau/b_P}^{q^*} \left(\frac{\tau}{q} - x_A \right) dF + \frac{1-\tau}{(1-\pi_O)^2 x_A} \pi_O f(q^*) \left(\frac{\tau-1}{q^*} \right) & \text{if } \frac{\tau}{\min[b_P, \sigma b_P]} < \frac{1-\pi_O \tau}{(1-\pi_O)x_A}, \\ \int_{\tau/b_P}^{q^*} \left(\frac{\tau}{q} - x_A \right) dF > 0 & \text{if } \frac{\tau}{\min[b_P, \sigma b_P]} \geq \frac{1-\pi_O \tau}{(1-\pi_O)x_A}, \end{cases} \\ \frac{\partial \Sigma}{\partial \tau} &= \begin{cases} \pi_O \left(\int_{\tau/b_P}^{q^*} q^{-1} dF + (x_A - b_P) f(\tau/b_P) \right) & \text{if } \tau/b_P < q^*, \\ \pi_O (x_A - b_P) f(\tau/b_P) < 0 & \text{if } \tau/b_P \geq q^*, \end{cases} \\ \frac{\partial \Sigma}{\partial b_P} &= \frac{\pi_O \tau}{b_P^2} (b_P - x_A) f(\tau/b_P) > 0, \\ \frac{\partial \Sigma}{\partial x_A} &= -\pi_O (1 - F(\tau/b_P)) < 0. \end{aligned}$$

These derivations are the basis of the following proposition.

Proposition 8. *If $q^* > \max[\tau / \min[b_P, \sigma b_P]]$, the Agency's incentive to sabotage the information collection process:*

- *is increasing with respect to the Principal's desire to promulgate a new policy (b_P);*
- *is decreasing with respect to the Agent's dissatisfaction with the status quo policy (x_A);*
- *is for some parameter values increasing, and for others decreasing, with respect to the probability that the Outsider will observe an unsabotaged report (π_O); and*
- *is for some parameter values increasing, and for others decreasing, with respect to the Outsider's competence (τ).*

The first and second conclusions of Proposition 8 are intuitive: each simply says that the Agency's desire to destroy information (and ensure that Principal is not allowed to make new policy by choosing $a_P = 1$) declines as the Agency and Principal's preferences become more similar. Accordingly, the third and fourth conclusions are the most interesting. Each of these illustrates an ambiguity that is particularly relevant if one considers the implicit institutional design problem for the Principal as described earlier (see 'Implications for institutional design' in Section 2). In a nutshell, both the third and fourth conclusion indicate that, if information collection is endogenous, a policy-seeking Principal should consider the details of the informational environment (i.e. the distribution of qualities, F) and the players' preferences (i.e. b_P , x_P , and x_A) when considering both how aggressively the Principal wants the Outsider to be (this is analogous to π_O) and how competent the Outsider should be with respect to policy implementation (i.e. τ).

Taken together, the results on endogenous information show that the welfare effects of stovepiping for the Principal are more subtle than they first appear. The results of Section 2 demonstrate that the 'first-order' effects of stovepiping is to increase P 's information, both directly through the stovepipe's reporting and indirectly through the altered incentives of the agent to transmit its information. But stovepiping also has 'second-order' effects when information in the Agency is endogenous. Stovepiping can in some cases (depending on the effectiveness of the stovepipe when the Agency remains ignorant) decrease the Agency's incentive to incur costs to obtain information, and stovepiping can increase the Agency's incentive to sabotage available information to ensure that no policy change can be made. These 'second-order' effects work against the 'first-order' effects of stovepiping by decreasing the aggregate information available in the policy-making process, to the detriment of the Principal.

4. Conclusion

When bureaucrats have the first cut at decision-relevant information and can censor it, they exercise agenda-setting power with respect to hierarchical principals. Principals, for their part, can employ several alternative channels to obtain the requisite information despite this agenda setting. Stovepipes and other outside providers of information provide such a channel. While these alternative information sources may be less 'conservative' (status quo inclined) than line bureaucrats, they typically do not have the substantive expertise or organizational resources that line bureaucrats have to formulate or implement

a decision. Thus, while stovepipes may be more willing than line agents to provide information to a decision-making principal, their information may not be as high in quality as the line agents'.

The model in this article captures this dynamic, and its effect on the information available to principals in policy making. The model reveals two effects whereby stovepipes can lead to greater provision of information to a decision maker. First, the stovepipe provides information to the principal in situations where the line agent does not, because the line agent considers the information too weak to be actionable. Second, the stovepipe induces the line agent to report information to the principal in situations where the line agent would not report if there were no stovepipe. This is precisely because the stovepipe's information is of lower quality than the line agent's, and given that policy will be changed in such instances, the line agent would rather its higher-quality information be used as the basis for it. The line agent would rather not change policy at all in these cases, but with a more aggressive (inclined to change the status quo) stovepipe and principal, this is not an option. Thus, the model reveals that in these senses, stovepipes and kitchen cabinets increase the information available to the principal to make a decision, and therefore the principal's utility.

At the same time, these effects assume that the information available to both the line agent and the outsider is given. When this information is endogenous, the effects of stovepipes and kitchen cabinets on the principal's utility are more subtle. First, if the stovepipe's information, although of lower quality than the line agent's, is reasonably good, the line agent may have less incentive to incur any costs to process the information the principal needs. In such cases, the stovepipe creates a redundant channel of information, and the line agent has an incentive to free ride on its provision of information to the principal. Ironically, to avoid such perverse incentives on the agent, the principal prefers a relatively incompetent stovepipe or kitchen cabinet. Second, the line agent may have an opportunity to 'sabotage' the decision-relevant information, or destroy it before either the agent itself or the stovepipe can act on it. The agent has an incentive to sabotage the information when it is sufficiently conservative. In such cases the principal is better off with no stovepipe at all; in that case, the agent passes at least some information to the principal, rather than censoring it entirely.

Acknowledgments

We gratefully acknowledge helpful conversations with Scott Ainsworth, Keith Dougherty, Justin Fox, Jeff Gill, Greg Huber, Kris Kanthak, Skip Lupia, Gary Miller, Burt Monroe, Matthew Stephenson, Andy Whitford, and Alan Wiseman, as well as the comments of audience members at the 2010 Conference on Public Policy and Political Economy at Stanford Law School, the American Politics Research Workshop at Yale University, the Political Economy series at Stanford Graduate School of Business, and the Department of Political Science at the University of Georgia. Of course, the authors are responsible for all remaining errors.

Notes

1. Or, in other words, the 'deciders.'
2. The office, which has since been relabeled the Northern Gulf Affairs Office, is located in the Office of the Under Secretary of Defense, Policy, in the Department of Defense.

3. A similar option is for principals to consult an informal network of advisors distinct from the formally designated information-processing units in the agency staff. This is a generalized version of the ‘kitchen cabinets’ sometimes used by presidents to help them define policy options and assess their consequences.
4. This should not be construed as an empirical argument that all line agencies do in fact have useful analytic capacity; rather, it is simply a recognition that if they do not, the principal has a different problem than the one we analyze in this article.
5. There are, however, three closely related models in the literature. First is that of ‘whistle-blowing’ (Ting, 2008). Whistleblowing involves transmission of information from a line agent directly to a principal when an intermediary would prefer to censor it. Stovepiping occurs when the line agent prefers to censor information reported to the principal, but another agent transmits it directly. Second is the model of three-tier hierarchical agency models originating with Tirole (1986). However, in Tirole’s model, the line agent exerts effort on behalf of the principal, and the intermediary can choose to report information about the line agent’s productive capacity to the principal. The crucial differences in our model are that (i) both agents have the ability to communicate with the principal, not just the intermediary; and (ii) our baseline model is not one of ‘effort’ or moral hazard, but rather provision of verifiable information. Nevertheless, our results share with Tirole’s the theme that, in hierarchical organizations, the agency problem with one agent interacts with the agency problem with another agent. Third is the model of ‘advocates’ (Dewatripont and Tirole, 1999). Like our model, theirs analyzes effects of multiple senders of verifiable signals. However, in our model the conflict between information providers is not as stark as the advocates on opposite sides of an issue of Dewatripont and Tirole, and one provider is unambiguously higher quality than the other in our case. Our primary focus is on the effects of an alternative, but lower-quality channel of information on the provision of it by the higher-quality channel, rather than on the effects of conflicts of interest on the incentives to acquire information as in Dewatripont and Tirole (1999).
6. Among other possible reasons, we focus on the strategic and rational aspects of stovepiping, and do not wish to argue that these aspects alone determined all behavior in the case, to the exclusion of behavioral heuristics and judgment biases.
7. We define a specific notion of ‘conservative’ below.
8. Type II errors are also possible, as we detail below when specifying utility functions.
9. See, for example, Bendor (1985), Heimann (1998), and Ting (2003).
10. To push this a bit farther (and allude to two assumptions of the model that are formally presented below), the crucial assumptions here are that O ’s reporting decision is pivotal (in the sense of affecting any player’s payoff) only when A does not report *and* that O observes the quality level, q_r , with certainty regardless of A ’s action.
11. The minimal quality being equal to 1, in addition to simplifying the algebra underneath the model, allows one to think of the quality of the report as the ratio of the information contained in the report to the information possessed by the Principal prior to reading the report.
12. We use the term “precision” in both its formal and intuitive senses (for they coincide here): the precision of a random variable is simply the reciprocal of its variance. Thus, the precision of policy making is inversely proportional to the uncertainty about the policy’s consequences. This uncertainty can be thought of as representing the quality of the report to the degree that higher-quality reports provide more information that can be used to foresee and insure against potential consequences of changing public policy.
13. This assumption is important in that it greatly simplifies the Agent’s calculation of the probability that the Outsider also possesses a report that the Agent has observed (and conversely for the Outsider when considering whether the Agent has observed a report that the Outsider possesses).

14. For our purposes, this approach to the Principal's payoff function is intended to capture the possibility that the Principal would rather 'do nothing' without sufficient information. We could capture this just as easily by assuming that there is a direct cost of making a new policy.
15. In addition to being a key technical assumption, it is also a substantive one, for it distinctly limits the degree to which one can think of r as being generated through the Agency's information collection efforts about an underlying 'state of nature.' In such a model, for example, low values of q_r might imply to the Outsider that the Agency did not observe the report.
16. Indeed, if $x_P + b_P > 1$, then $a_P = 1$ is a strictly dominant strategy for the Principal.
17. As opposed, for example, to the microfoundations of information collection, as studied by Patty (2009).
18. Specifically, if Assumption 1 is violated, then each player's optimal choice is dictated by weak dominance, which is essentially a decision-theoretic concept.
19. Note that it is not necessary to consider O 's inferences about q_r conditional on $a_A = 0$. If O observes r then there is no additional information conveyed to O due to $a_A = 0$. If O does not observe r , then $a_A = 0$ does convey information, but it is irrelevant because O does not itself have a report to provide P .
20. These expected payoffs are calculated conditional upon the Agency observing r : thus, π_A plays no role. Clearly, an *ex ante* expected payoff calculation would incorporate this probability. But the observation of r by the Agency is by assumption outside of the Agency's control and is independent of all other aspects of the game.
21. If $b_P + x_P < 1/q$, the Principal will not act upon the report. Since forwarding the report is costless, it is weakly dominant for the agency to forward the report in such cases.
22. While the proof of this Proposition is omitted, it follows from inspection of Equation (5), it is important to recall that it utilizes our assumption that $\tau > 1$.
23. Specifically, even though the Agency will rationally forward some reports that would not be forwarded in the absence of the Outsider, the minimal level of quality that the Agency would prefer to forward if the Principal were going to act upon that quality is lower than the minimal level of quality that the Principal would act upon in the Agency's absence. Accordingly, sequential rationality implies that the Agency's behavior in such realms is sensitive to the Principal's preferences.
24. This modification makes the Principal's calculations more straightforward and it is straightforward to verify that this has no effect on any of the players' incentives, as only the Agency has nontrivial calculations to make in the first place, and those occur only if the Agency receives the report.
25. Something that will of course not be true in the equilibria of the endogenous information extensions examined later in the article.
26. This point resembles one made by Dewatripont and Tirole (1999). In their model of information acquisition by advocates, they show that if agents can censor 'bad news' then it might be preferable for the Principal to rely on one conduit of information rather than two competing ones, because then the Principal can at least get partial information when it is in the Agency's interest to pass it along.
27. This is without loss of generality because b_P can be interpreted as already capturing the opportunity cost of maintaining the status quo for both P and O .
28. This assumption could be relaxed in transparent ways, but creates a clear baseline for analysis: the report represents 'raw information' that is unintelligible to non-experts prior to being processed in some fashion.
29. For reasons of parsimony, we omit the presentation of an explicit 'microfoundation' for the verbal depictions of the behavioral implications of our comparative statics analysis. We omit this because it is straightforward and well understood how to extend this type of model so as to

have (in this case) the Agency face a randomly and exogenously determined cost of sabotage so as to more properly motivate the use of phrases such as ‘more likely to sabotage.’

References

- Bendor J (1985) *Parallel Systems: Redundancy in Government*. Berkeley, CA: University of California Press.
- Bendor J, Taylor S and Gaalen RV (1985) Bureaucratic expertise versus legislative authority: a model of deception and monitoring in budgeting. *American Political Science Review* 79: 1041–1060.
- Carpenter DP (2010) *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton, NJ: Princeton University Press.
- Dewatripont M and Tirole J (1999) Advocates. *Journal of Political Economy* 107: 1–39.
- Dreyfuss R and Vest J (2004) The lie factory. *Mother Jones*, 26 January. Available at: <http://www.motherjones.com/politics/2004/01/lie-factory> (accessed 26 March 2013).
- Heimann CFL (1998) *Acceptable Risks: Politics, Policy, and Risky Technologies*. Ann Arbor, MI: University of Michigan Press.
- Hersh SM (2003) Selective intelligence. *The New Yorker*, 12 May. Available at: <http://www.newyorker.com/archive/2003/05/12/030512fa-fact> (accessed 26 March 2013).
- Hersh SM (2004) *Chain of Command*. New York, NY: HarperCollins.
- Niskanen W (1971) *Bureaucracy and Representative Government*. Chicago, IL: Aldine.
- Patty JW (2009) The politics of biased information. *Journal of Politics* 71: 385–397.
- Ting MM (2003) A strategic theory of bureaucratic redundancy. *American Journal of Political Science* 47: 274–292.
- Ting MM (2008) Whistleblowing. *American Political Science Review* 102: 249–267.
- Tirole J (1986) Hierarchies and bureaucracies: on the role of collusion in organizations. *Journal of Law, Economics, and Organization* 2: 181–214.
- Weber M (1946) Economy and society. In: Gerth H and Mills CW (eds), *Max Weber: Essays in Sociology*. New York, NY: Oxford University Press.