Journal of Theoretical Politics 21(2): 161–186
DOI: 10.1177/0951629808100762
http://jtp.sagepub.com

MULTIPLE PRINCIPALS AND OVERSIGHT OF BUREAUCRATIC POLICY-MAKING

Sean Gailmard

ABSTRACT

I examine a model in which multiple legislative principals monitor a bureaucratic agent's implementation of a project. The principals can each perform oversight of the implementation to limit information asymmetries exploited by the agent. Oversight is costly to perform and due to information leakages between principals, oversight by one principal reveals information to all principals. Thus for some values of the audit costs, there is a collective action problem in monitoring among the principals: the multiplicity of principals can cause the level of this form of oversight to be underperformed relative to the principals' joint interests. Notably, the multiplicity of principals reduces their collective control over the agent even though they have common interests about the agent's actions, i.e. conflicting preference about agent actions are not necessary to attenuate accountability when there are multiple principals. Overall the results point out that the institutional structure of the overseeing body has an important effect on accountability, independent of the institutional structure of the overseen.

KEY WORDS • bureaucracy • common agency • congressional oversight • multiple principals

Institutions of bureaucratic policy-making have a decisive effect on the responsiveness of policy made through administrative channels to the policy goals and preferences of legislatures. Following the seminal research of Fiorina (1981), McCubbins et al. (1987, 1989), and Moe (1985, 1989), it is impossible (or at least unwise) to overlook political principals' desires to achieve accountability and responsiveness from agents – if necessary, at a cost of 'effectiveness' – as an important cause of bureaucratic structure. In addition, actions and behaviors within the legislature (both before and after legislation delegating to

Thanks go to Jeff Banks, Steve Callander, Tim Feddersen, Matt Jackson, Rod Kiewiet, Richard McKelvey, Tom Palfrey, Garry Young, and seminar paticipants at Michigan State University, Northwestern University, and the University of Chicago for helpful comments. As usual I am responsible for any errors. This research was generously supported by National Science Foundation Grant SES-0318347.

bureaucracies is passed) are themselves important for securing responsiveness of bureaucratic policy-making (Aberbach, 1990; Bawn, 1995, 1997; Epstein and O'Halloran, 1999; Huber and Shipan, 2002). Since the institutions by which political principals are organized affect the decisions they make, those institutions, in addition to institutions in the bureaucracy itself, exert an important influence on the accountability of agents to their political principals. In this article I use a formal model to develop and clarify the logic of this argument as it applies to legislative oversight of agencies.[1]

My argument builds on several important components of the legislative–bureaucratic nexus in the USA, among them that bureaucrats typically know more about what they do than legislators, and that legislators can use oversight to obtain outside information and limit this information asymmetry.[2] Moreover, bureaucrats care about much more than their budgets, such as reputations, relationships with superiors, and environmental stability.[3] Thus, focusing on budgets as the primary lever legislators can 'pull' to back up their oversight is not always appropriate. Finally, and most importantly in terms of the institutional argument in this article, bureaucrats face multiple legislative principals.[4] For example, bicameralism, the appropriations-authorization process, and entrepreneurial committees subject to fluid oversight jurisdiction all present multiple principals to bureaus.

Based on these components I develop a formal model with multiple principals independently attempting to influence a single common agent subject to both hidden actions and hidden information. The bureaucratic agent performs a project at some cost for legislative principals, who can each obtain costly information about the project's cost before offering the agent an incentive scheme to induce good performance (cf. Spencer, 1980; Bendor et al., 1987). The model shows that the multiplicity of legislative principals attenuates legislative control over bureaucracy, but the results still occupy a middle ground between 'congressional dominance' and the dominated Congress of Niskanen's (1971) monopoly bureau model.[5] Furthermore, since each principal in the model values the agency's project, the agency loss cannot be attributed to conflicts among them (as is the case with multiple principals as a cause of bureaucratic drift in McCubbins et al., 1987 and Hammond and

1. Volden (2002) makes a similar argument about the importance of the organization of political principals, focusing on the statutory discretion granted jointly by Congress and the president.

2. See especially Ogul (1976), McCubbins and Schwartz (1984), Banks (1989), Aberbach (1990), and Banks and Weingast (1992).

3. Fenno (1966), Wildavsky (1978), Kaufman (1981), Wilson (1989), and Golden (2000) all make this point in detail.

4. See e.g. Mitnick (1980), Wilson (1989), McCubbins et al. (1987, 1989), and Hammond and Knott (1996).

5. Spencer (1980), Miller and Moe (1983), and Bendor et al. (1987) make similar points that follow from depicting legislative principals as active participants in agency policy-making, rather than passive recipients of agency requests/demands.

Knott, 1996). It is due instead to the collective action problems created by the diffusion of control and oversight authority. By isolating the institutional structure of the overseeing body from value conflicts within it, the model supports the argument that the institutional structure of the overseeing body, as well as that of the overseen, is an important determinant of responsiveness and accountability.

In the model, any oversight of bureaucrats for which legislators have to use or trade resources – be it 'fire alarm' or 'police patrol', formal or informal, latent or manifest – may be underprovided (relative to the optimum for the principals jointly) due to collective action problems, which in turn arise from the multiplicity of principals. As Ogul (1976: 181) noted, while members of Congress desire more oversight at the collective level, individual incentives often dictate leaving it to someone else. In this sense the model supports a version of Ogul's argument that all oversight, be it latent or manifest, is subject to underprovision. On the other hand, for some model parameters the level of oversight captured here is efficient, so as McCubbins and Schwartz (1984) and Aberbach (1990) have each argued, oversight is not necessarily 'Congress's neglected function'. Thus, these apparently diametrically opposed views of oversight can in fact be understood as unified at a more fundamental level, in the sense of being equilibrium outcomes in one model.

Furthermore, the model shows that the more principals use their information networks for oversight purposes, and the more effective the oversight technology is, the worse the collective action problem will be. Both of these changes increase the benefit to all principals from any one principal's oversight activity. Since those benefits are not internalized in any individual oversight decision, more available surplus is lost. In addition, the model shows that a more effective oversight technology can in some cases reduce the principals' collective welfare, besides reducing the share of the available benefits that they capture in equilibrium. The reason is that more effective oversight by one principal can make it more attractive for others to free-ride, if the marginal value of additional oversight is sufficiently low. Taken together, the results show that because of a collective action problem among multiple principals, more effective oversight can both reduce the share of available surplus obtained in equilibrium and reduce the net surplus realized in equilibrium.

Conceptually, this article's point is ironically similar to that made by Ting (2003), who shows that redundancy in *bureaucracy* will not necessarily increase the effectiveness or reliability of the bureaucracy as a system, when strategic considerations and collective action problems are reckoned. In this article I show that redundancy in *oversight jurisdiction* will not necessarily increase the effectiveness of oversight and, therefore, the performance of bureaucracies, due to similar collective action problems among oversight bodies. My argument therefore provides conceptual foundations for an observation about overlapping oversight jurisdictions made by Dodd and Schott (1979), the continuing relevance of which is demonstrated by

O'Connell's (2006) review and analysis of issues in structuring oversight of intelligence agencies to ensure effective performance.

My formal argument builds on contract theory and the economic theory of incentives under asymmetric information. The seminal work on multiple principals offering optimal incentives subject to adverse selection is due to Martimort (1992) and Stole (1997); Bernheim and Whinston (1986) initiated the study of optimal incentives in common agency with moral hazard. Laffont and Tirole (1991, 1993) developed the hidden action/hidden information structure used in this article. Dixit (1996) used a common agency model with moral hazard to explain the low-powered financial incentives typically observed in public bureaus. McCubbins et al. (1987, 1989) and Hammond and Knott (1996) offered pioneering spatial frameworks to account for the multiplicity of principals attempting to influence bureaucracies, but the incentives, incomplete information, and oversight that drive this model are not present.

The rest of the article is organized as follows. Section 2 motivates the model and lays it out formally. Section 3 analyzes the multiple principals model with general beliefs about the agent's type. In Section 4, I investigate optimal auditing by the principals and welfare effects of parameter changes. Section 5 concludes. Proofs are contained in the Appendix.

1. The Model

This section first discusses some intuition and justification for the model's components and then introduces it formally. It is convenient to separate these for easier reference.

1.1 Modeling Issues

For the formal common agency approach, each principal must be able to offer its own incentive scheme to the agent. Budgets are a common focus in the discussion of legislative control of bureaucracy (e.g. Banks, 1989; Banks and Weingast, 1992; de Figueiredo et al., 1999; but see Ting, 2001), and of course, a legislature can offer only one budget to an agency. The perspective here is instead that the legislative principals each have at their disposal other, nonbudgetary incentives (e.g. perks or abuse at the hands of a committee), which they are free to offer as they see fit,[6] and which for simplicity the bureaucrat treats as equivalent to money.[7]

6. Baron (2000) uses a similar approach in a bilateral agency model of legislative organization.

7. This is not restrictive if, for example, the bureaucrat cares about budgets and contumely, is risk neutral in both, and has an additively separable utility function. Then the assumption amounts to a choice of units.

The political science literature (e.g. Fenno, 1966; Wildavsky, 1978; Kaufman, 1981; Wilson, 1989; Aberbach, 1990) documents that bureaucrats are interested in these sorts of incentives under the control of individual committees, not just budgetary incentives of committees acting in concert. Kaufman (1981) documents that avoiding embarrassment, of the kind congressional committees are capable of inflicting, is a key priority of public administrators. This is not only because of any intrinsic cost of it, cost to personal reputation, or loss of outside opportunities, but also because it makes it more difficult to manage and motivate subordinates within the bureau. Moreover, perhaps because it acts in part as a signal in environments with incomplete information, it damages the interaction of bureaus with other stakeholders, such as interest groups, other federal or state bureaus, and the Executive Office of the President at the federal level.

Legislative principals find it costly to provide these incentives, positive or negative, for the bureaucrat. It takes time from other legislative activity, fundraising, or case work. Since these activities have electoral payoffs, there is some opportunity cost of time taken from them.[8]

Cost-reducing activity ('effort') is costly for the bureaucratic agent. Effort spent on any project has opportunity costs – for example, if time constraints bind and it takes time from other valuable policy initiatives. Moreover, any organization, government bureau or not, has waste due to imperfect coordination that is difficult to identify and eliminate. In addition, projects involving subcontracting with private firms may involve incentive problems of their own that are costly to reduce. Thus, the assumption of costly effort need not imply that bureaucrats like to shirk or are lazy.

The oversight or outside information available to the legislative principals is costly. It may require valuable legislative concessions if provided by interest groups, or valuable time for case work, each of which have opportunity costs. Or, if formally provided by legislative support agencies like the Government Accountability Office at the federal level or presented in formal hearings, and especially if of a technical nature, it will take time to consume and distill.

Outside information gathered by one principal is observed by and beneficial to all principals. This formalizes Aberbach's (1990) observation that congressional committees keep tabs on other committees' oversight activities. Given that oversight increasingly occurs on the public record, a piece of intelligence about an agency is likely to make its way to multiple principals. On the other hand, some benefits of information networks can be more 'proprietary', such as the development of innovative policy on topical issues for which legislators can personally claim credit and enjoy publicity. These benefits are targeted specifically to legislative principals that invest in information gathering.

8. See Cameron and Rosendorff (1993) and Baron (2000) for a similar perspective on the cost to Congress of incentive or discipline measures.

Both principals in the model derive positive benefits (though not necessarily equal, and not necessarily greater than from other projects) from the bureaucrat's project. For example, if the Customs Service initiates a project stressing the integrity of agricultural imports, it could benefit both the agricultural committees and the counterterrorism-minded intelligence committees in the legislature, even though other Customs projects might have greater counterterrorism benefits to the intelligence committees. But arguably, conflicting values among the principals are key to bureaucratic politics. However, it would be less surprising to show that preference conflict among multiple principals attenuates their collective control over the bureaucracy. Since this model focuses on the multiplicity of principals itself, rather than conflict or differences among them, it strengthens the conclusions about how this institutional structure per se affects the accountability of bureaucrats, and how it interacts with oversight incentives.

1.2 Formal Structure

There are two principals P_1 and P_2 and one agent A. The situation here is one of 'intrinsic' common agency (Bernheim and Whinston, 1986), so the agent can contract either with both principals or neither. For simplicity I assume both principals are 'contracting' with the agent over a single project. Furthermore, the incentive scheme of a given principal can be based only on the report of the agent's efficiency parameter, and not features of the other principal's contract. This would be the case if, for example, one principal's incentive scheme was not observable by another, because it was not a matter of public record. This obviously does not apply to budgets but is more applicable to informal communication that seems at least as important in influencing agency behavior (see Subsection 2.1 for more discussion).

The agent can perform the project for the principals at cost $D = \theta - \xi$. The exogenous parameter $\theta \in [\underline{\theta}, \overline{\theta}] \subset \Re$ is A's privately known efficiency parameter, and $\xi \in \Re$ is A's 'effort'. P_i's prior belief is that θ is distributed according to some CDF F with strictly positive density f and a monotone hazard rate, and this is common knowledge. Total project cost D is publicly observable, but neither θ nor ξ is: all players know the total cost of the project, but the principals cannot disentangle how much is beyond the agent's control. Effort is costly for A, with the effort cost given by $e(\xi)$, $e' > 0$, $e'' > 0$, $e''' \geq 0$.[9] The function e is commonly known. Let $t_i(D)$ be the transfer function for P_i, representing the transfer to A given cost observation D. Let the agent's utility function be

$$u_A(\theta) = t_1(D(\theta)) + t_2(D(\theta)) - e(\xi(\theta)).$$

Thus by convention the principals pay the cost of the project and can independently offer perks and other benefits to the agent. I assume these t functions are differentiable and therefore focus on differentiable equilibria.

9. $e''' \geq 0$ helps to avoid complications with stochastic incentive schemes.

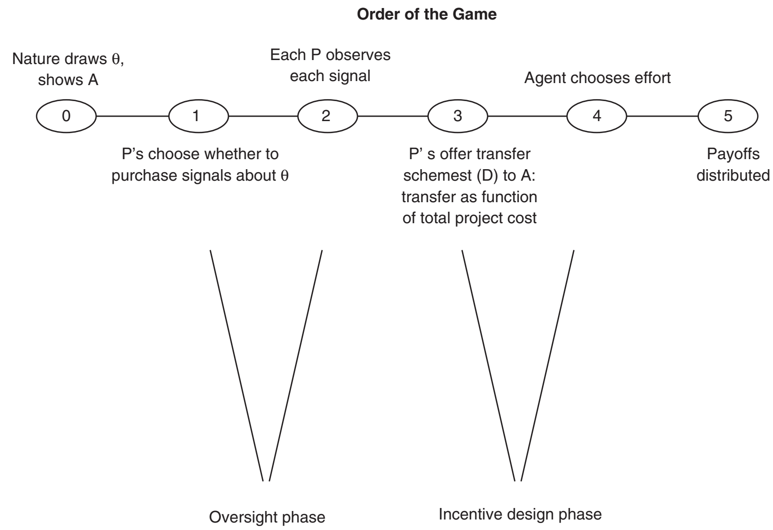

Figure 1. Timing of the Game

The project, if realized, has value $V_i > 0$ to P_i. The rewards $t_i(D)$ are costly for principals to offer, so P_i's utility is

$$u_i = \begin{cases} V_i - t_i(D) - s_i D & \text{if the project is performed} \\ 0 & \text{otherwise} \end{cases}$$

where $s_i > 0$ is the share of the project's cost borne by P_i. It makes no difference whether $\sum_i s_i = 1$, so the principals collectively internalize all of the project's cost, or if $\sum_i s_i < 1$, so they pass some on to the rest of the legislature (as with congressional committees in some models of legislative organization).

Each principal is able to audit the agent and thereby, with some probability, discover to what extent the cost D is beyond the agent's control (i.e. due to the efficiency parameter θ). The auditing technology for P_i is as follows. With some probability π an audit reveals A's true effciency θ. With probability $1 - \pi$ the audit reveals nothing and principals keep their prior beliefs. If both principals audit, their results are statistically independent conditional on the actual state θ (thus the chance that j's audit succeeds is unaffected by the success of i's audit). P_i can purchase an audit with cost C, regardless of its results.[10]As will be described shortly, the fixed cost of the audit may yield benefits for oversight

10. An alternative specification would allow principals to purchase audit success probabilities from a continuum according to a cost function. This would preserve the possible collective action problem among the principals, which would still become more severe with suitably parameterized decreases in the auditing cost function and increases in the importance of oversight, but would complicate the results.

activity that accrues to all overseeing principals, and private benefits of information gathering that accrue to the auditing principal only.

Denote by $p(s)$ the posterior beliefs held by the principals when $s=\{s_1,s_2\}, s_i\in\{\varnothing\}\cup(\underline{\theta},\overline{\theta})$, is the vector of audit signals to P_1 and P_2. Thus $p(\varnothing)\equiv p(\{\varnothing,\varnothing\})=F(\theta)$ (the prior); $p(\widehat{\theta})\equiv p(\{\varnothing,\widehat{\theta}\})=p(\{\widehat{\theta},\varnothing\})=p(\{\widehat{\theta},\widehat{\theta}\})$ is given by $\Pr(\theta)=1$ if $\theta=\widehat{\theta}$ and $\Pr(\theta)=0$ otherwise. So in the auditing stage before the 'contracting' stage, information leakage from one principal to another is complete.

The game proceeds as follows (see Figure 1). Nature draws A's type θ and shows A. The principals simultaneously decide whether to audit by choosing numbers $n_1=1$ if P_1 audits and 0 otherwise, and similarly for n_2. P_i and P_j observe the signal from P_i's audit. Then each principal is simultaneously allowed to select an incentive scheme to offer the agent, associating a reward $t_i(D)$ with each cost observation D. Following this the agent chooses an effort level ξ. Finally the game ends and payoffs are distributed according to the t_i functions chosen by the principals.

Let the choices of n_i for any continuation utilities be called an auditing game. Let $V_i(\theta)$ be the equilibrium utility to P_i in the common agency subgame when P_i begins that subgame knowing the true state θ. Let $\overline{V}_i(F)$ be the ex ante equilibrium expected utility in the common agency subgame when P_i begins that game with uncertain beliefs F about θ.[11]

P_i's audit has 'public oversight benefits', policy and efficiency benefits common to all principals, as well as 'private oversight benefits' to P_i only. The latter benefits capture (for instance) 'position taking' motivations for oversight, e.g. the electoral and news media value of crusading against waste, fraud, and abuse in bloated bureaucracies. Any such benefits are captured by the principal doing the crusading and do not spill over to other principals. Let $W_i^a(n),\ n=0,1$ or 2, be the expected utility to P_i from the entire game with oversight when $n=n_1+n_2$ audits are executed. Let $W^a(n)=W_1^a(n)+W_2^a(n)$. For general π, we have

$$W_i^a(2)=[\pi^2+2\pi(1-\pi)]\left[\int_{\underline{\theta}}^{\overline{\theta}} V_i(\theta)dF(\theta)\right]+(1-\pi)^2(\overline{V}_i(F));$$

$$W_i^a(1)=\pi\left[\int_{\underline{\theta}}^{\overline{\theta}} V_i(\theta)dF(\theta)\right]+(1-\pi)(\overline{V}_i(F));$$

$$W_i^a(0)=\overline{V}_i(F).$$

11. This assumes some predetermined selection from the continuum of equilibria entailing different splits by the principals of the cost of meeting the agent's constraints. Equilibrium selection is discussed further in Section 3.

Let $W_i^b(n_i)$ be the exogenous private oversight benefits, which depend only on P_i's audit. Let

$$W_i(n_1,n_2) = vW_i^a(n) + (1-v)W_i^b(n_i)$$

be the total benefit to P_i of the auditing decisions. Thus $v \in [0,1]$ parameterizes the importance of oversight in the committees' functions. Let $W(n_1,n_2) = W_1(n_1,n_2) + W_2(n_1,n_2)$.

2. Incentive Subgame with Multiple Principals

In this section, I analyze the common agency subgame played by the principals and the agent for any given beliefs and, therefore, audit results. This is a benchmark case to review the results when multiple principals each offer incentive schemes to an agent whose performance they all care about, and it defines continuation values of the larger game with auditing that will be analyzed later. Under both complete and incomplete information, the incentive subgame has multiple equilibria. In each case, the one analyzed is the most efficient (with respect to the welfare of all parties), in order to avoid biasing the findings toward inefficiency.

Under incomplete information, the agent with type θ makes a report $\widehat{\theta}$ of his type to each principal. The principals must each associate an effort level ξ and a transfer t_i with every reported type $\widehat{\theta}$. A combination of functions $(\xi(\widehat{\theta}), t_i(\widehat{\theta}))$ is called a *mechanism* for principal i. A widely used result in contract theory, the Revelation Principle, ensures that (once it is assumed that principals can commit to a mechanism) there is no loss of generality from restricting attention to mechanisms that are incentive compatible, i.e. that give A an incentive to truthfully report θ to each P.[12] The idea is simply that if A would benefit from dissembling about θ, a principal could just as well obtain a truthful report and offer a transfer to compensate the agent for reporting truthfully rather than dissembling, in which case the agent would be just as happy to report truthfully. This relies crucially on the assumption that principals follow through on promised transfers, but given commitment, truth-telling does not impose an additional restriction. The Revelation Principle should be read not as a substantive assertion that agents never dissemble, but rather as a technical assertion that any solution with dissembling is equivalent in terms of payoffs to some other solution without, and it is easier to analyze only the latter category.

Thus the principals choose their mechanisms in a noncooperative game. Equilibrium entails that an agent of type θ is better off reporting θ than any other

12. An important issue is how the revelation principle applies in mechanism design games with multiple principals. Given a mechanism by P_j, there is no loss of generality in restricting P_i to direct mechanisms. Martimort (1996) and Martimort and Stole (1999) discuss this in more detail.

type (incentive compatibility), the principals induce the same effort level from A for the reported type (coordination of effort levels), and each principal's ξ and t_i functions maximize her utility (in expectation with respect to the unknown agent type θ) given the functions chosen by the other principal. Note that the strategy space for each principal is a function space, so that analysis of the maximization problems relies on a different branch of optimization theory (to wit, optimal control theory) than the usual case where the strategy space is a vector space. However, in the present environment these differences are somewhat slight. Also note that in the incentive subgame, the principals necessarily have symmetric information about θ, because of information leakage in auditing at the oversight phase, and move simultaneously.

Since incentive compatibility is a constraint on the principals' choice of incentive schemes, it must be determined exactly what the constraint specifies. Let

$$\upsilon(\theta,\widehat{\theta}) \equiv t_1(\widehat{\theta}) + t_2(\widehat{\theta}) - e(\theta - D(\widehat{\theta}))$$

be A's rent as a function of the true type θ and the announced type $\widehat{\theta}$ and

$$U(\theta) \equiv \upsilon(\theta,\theta)$$

be the rent of type θ.

Incentive compatibility means that truthful reporting of θ is in A's best interest. Formally this requires that the rent $\upsilon(\theta,\widehat{\theta})$ is maximized by reporting θ when the type really is θ. A first order condition for this maximization is $\frac{\partial \upsilon(\theta,\widehat{\theta})}{\partial \widehat{\theta}} = 0$ for all values of θ. Using this in $\frac{dU(\theta)}{d\theta}$ gives the first order condition for incentive compatibility in a form that does not depend on transfers t_i:

$$\frac{dU(\theta)}{d\theta} = -e'(\xi(\theta)), \forall \theta \in [\underline{\theta},\overline{\theta}]. \tag{IC-1}$$

The 'natural' temptation facing A is to overstate θ. If principals believed such overstated reports they would offer transfers sufficient to compensate for relatively high effort, so low θ agents could obtain high transfers without having to exert high effort. The incentive compatibility constraint reflects that this temptation is neutralized if agents reporting higher intrinsic costs θ enjoy lower surplus.

Of course, IC-1 is a necessary but not sufficient condition for an (interior) optimal report of θ. The second order sufficient condition requires that the project's actual (post-effort) cost grow with the intrinsic cost θ, or

$$\frac{dD(\theta)}{d\theta} > 0. \tag{IC-2}$$

Together the first and second order incentive compatibility conditions IC-1 and IC-2 are sufficient to ensure that A's optimal report of θ is truthful. Incentive schemes must satisfy these conditions in order to be incentive compatible, so they specify the limitations on the effort the principals can extract from the agent and transfers that must be offered in order to maintain truthful reporting.

In addition to inducing truthful reports conditional on A participating, the incentive schemes must induce the participation of the agent in the first place.[13] The individual rationality constraint is

$$U(\theta) \geq 0, \forall\theta. \tag{IR}$$

Assuming the project is implemented, each principal wishes to choose a mechanism $\xi(\theta)$, $t(\theta)$ to maximize her own utility $V_i - t_i(D) - s_i D$, taking as given the mechanism of P_j and the *IC* and *IR* constraints. Combining the objective and constraints in one function (and noting that $t_i = U(\theta) - t_j + e(\xi(\theta))$ from the definition of $U(\theta)$) yields a Hamiltonian H_i for principal i:[14]

$$H_i = [V_i - s_i(\theta - \xi(\theta)) - e(\xi(\theta)) - U(\theta) + t_j(\theta - \xi(\theta))]f(\theta) - F(\theta)e'(\xi(\theta)). \tag{1}$$

The utility obtained in the solution for a given type θ is weighted by the density $f(\theta)$; this is a typical expected utility formulation. But when changing the effort demanded of type θ, P must also consider the effect of this change on the cost of ensuring truthful revelation from types below θ: offering high transfers for effort from high types θ, while optimal *conditional* on that θ, raises the temptation of lower types to claim to be that high type. The probability of a lower type than θ is exactly $F(\theta)$, so the cost of the incentive constraint in the Hamiltonian is weighted by this cumulative density.[15]

The key first order condition defining A's equilibrium effort is obtained from $\frac{\partial H_i}{\partial \xi} = 0$, which translates to

$$e'(\xi(\theta)) = s_i - t_j'(\theta - \xi(\theta)) - \frac{F(\theta)}{f(\theta)} e''(\xi(\theta)). \tag{2}$$

The term $t_j'(\theta - \xi(\theta))$ arises because there are multiple principals. Thus P_i can free-ride on the transfer offered by P_j when inducing effort from A, a fact that will be reflected (and mitigated) in P_j's equilibrium transfer function.

The lefthand side of Equation 2 is also restricted by IC-1. Using this fact and adding the first order conditions from H_1 and H_2,[16] it follows that equilibrium

13. Of course an agency created by Congress cannot literally decline as an organization to participate in the mechanism congressional overseers implicitly design. Nevertheless the people who work for the agency and actually possess its expertise and implement its projects might decline. These agents must be sufficiently satisfied with public service that they wish to participate in it, and thus an individual rationality constraint for A is an appropriate construct in the application of this model.

14. The Hamiltonian serves a similar purpose in optimal control problems that a Lagrangian serves when the choice space is a vector space (cf. Kamien and Schwartz, 1981).

15. The Hamiltonian reflects IC-1 but not IC-2. It can be verified that the first order solution obtained in this way does indeed satisfy IC-2.

16. If two equations must be satisfied in equilibrium, their sum must be as well.

effort under common agency with imperfectly informed principals is $e'(\xi^{CA}(\theta)) = s_1 + s_2 - \frac{2F(\theta)}{f(\theta)} e''(\xi^{CA}(\theta))$.[17]

The key comparisons are between the common agency effort levels and effort in baseline cases with (*i*) principals behaving cooperatively and holding complete information on θ (the 'first best' case), and (*ii*) principals behaving cooperatively but facing uncertainty about θ described by $F(\theta)$ (the 'second best' case). These benchmarks capture a unitary, centralized oversight structure with and without information asymmetries, respectively. The conditions yielding the equilibrium effort levels are as follows:

$$e'(\xi^{FB}(\theta)) = s_1 + s_2 \qquad \text{(First Best)}$$

$$e'(\xi^{SB}(\theta)) = s_1 + s_2 - \frac{F(\theta)}{f(\theta)} e''(\xi^{SB}(\theta)) \qquad \text{(Second Best)}$$

$$e'(\xi^{CA}(\theta)) = s_1 + s_2 - 2\frac{F(\theta)}{f(\theta)} e''(\xi^{CA}(\theta)). \qquad \text{(Multiple Principals)}$$

Thus the first best solution entails effort such that the marginal cost for a given type equals the principals' collective marginal benefit of reducing cost. In the second best solution the principals act cooperatively and balance the marginal benefit of effort against not just the marginal cost to the agent of exerting it (for which the principals must compensate), but also the marginal cost of inducing truthful revelation of θ by A.

To compare effort levels note that since $e'' > 0$ (the marginal cost of effort is increasing for A: a few cost savings are easy to find but they become progressively more difficult), $e'(\xi^{FB}(\theta)) > e'(\xi^{SB}(\theta)) > e'(\xi^{CA}(\theta))$, $\forall \theta > \underline{\theta}$. It follows, as is typical in common agency models with adverse selection (Martimort, 1992; Stole, 1997), that the equilibrium effort in the incentive game with multiple principals and asymmetric information is distorted downward relative to the second best effort, which itself is distorted downward relative to the first best effort.

PROPOSITION 1. *When principals believe θ is distributed according to $f(\theta)$, the equilibrium effort of type $\theta > \underline{\theta}$ is lower under multiple principals than the second best effort under a single principal with beliefs $f(\theta)$.*

17. Equilibrium also requires specifying the transfer functions implementing this effort, though they are of less immediate substantive interest. They are obtained by integrating the first order conditions from the principals' maximization problems, yielding (for $i = 1, 2$)

$$t_i(\theta - \xi^{CA}(\theta)) = \int_{\theta}^{\overline{\theta}} \{(e'(\xi^{CA}(\widehat{\theta})) + \frac{F(\widehat{\theta})}{f(\widehat{\theta})} e''(\xi^{CA}(\widehat{\theta})) - s_j) (1 - \frac{\partial \xi(\widehat{\theta})}{\partial \widehat{\theta}}) d\widehat{\theta}\} + t_i(\overline{\theta} - \xi^{CA}(\overline{\theta})).$$

The extra distortion under incomplete information is due to an externality among the principals. The information asymmetry on θ creates the tradeoff between efficiency and rent extraction common in bilateral models. This calls for downward distortion of the agent's effort, relative to the first best. If one principal's contract causes further downward distortion in effort, this leads to lower effort for other principals too. But since these other principals in turn are distorting effort downward, an increase in effort benefits them. This creates a negative externality across principals and leads to excessive downward distortion.[18]

This approach implies that principals trade less desirable policy outcomes (production levels) from inefficient agents for lower informational rents to efficient ones (cf. de Figueiredo et al., 1999). Therefore, principals will always be dissatisfied with the performance of bureaucrats; but provided some of each type of agent exist, only some of this dissatisfaction stems from the optimal tradeoff with information rents. Common agency implies two sources of dissatisfaction with agency policy: the tradeoff with information rents, and the existence of other principals. Only the former benefits any principal.

The analysis thus far covers the case with principals uninformed of the agent's type θ. When an audit is informative and the principals are informed of θ, the incentive subgame has an efficient equilibrium with agent effort equal to the first best, with $e'(\xi(\theta)) = s_1 + s_2$. The cost to the agent of an increment of effort is equated to the reduction that increment causes in the project cost that the principals must bear. The contractual externality among the principals is no longer present, since there is no information rent inducing them to distort the output of any type. At least in the incentive design phase, common agency results in perfect coordination among principals when information is complete; it is only under asymmetric information with respect to A that incentives are distorted under common agency.

Proposition 2. *When principals are perfectly informed of θ, the equilibrium effort of type θ is the same under multiple principals as the first best effort under a single principal.*

These results are based on differentiable t functions. With incomplete information, there are also equilibria with nondifferentiable t functions. The simplest transfer functions that illustrate this have each P offer a payment t^* when project cost D meets some target D^*, and $-\infty$ otherwise. A optimizes by accepting

18. This discussion only considers the incentive distortion from common agency when the project is actually executed, but there is another related distortion: the project will be executed for fewer θ values with multiple principals than with one principal. The extra downward distortion on effort under common agency means that there are more types θ such that the total project cost $D(\theta)$ is not worth it to the principals for a given total value of the project.

if $2t^* > e(\theta - D^*)$ and rejecting otherwise. Accepting agents simply meet the target cost. As θ falls below that of the type just willing to accept, agents capture more rent because less effort is needed to meet the target. The principals choose the cutoff type to economize on the slack left to the agent types that do produce. In equilibrium, the marginal cost of the rent will equal the lost surplus of the marginal agent type.

The key is that the principals' utility is higher with complete information than under differentiable or nondifferentiable equilibria with incomplete information. There would be no slack left to agents more efficient than the cutoff type; that payment would be retained by the principals. Moreover, unlike in the nondifferentiable equilibria, all agent types would produce some surplus for the principals. Thus, what follows is robust to consideration of nondifferentiable equilibria.

3. Noncooperative Auditing

Now I turn to a representation of oversight of bureaucrats by legislators. I focus on a form of oversight that, while simple, still captures many different types of relationships. The important point is that legislators must trade some resources, whether direct opportunity costs or legislative concessions, for the information in the audit.[19] Since an audit by either principal will mitigate the information problems both principals face with respect to the agent, and since audits are costly, the principals face a collective action problem about who will become informed. Oversight with multiple principals layers another collective action problem on top of the one described above for the incentive design phase. Even when the principals are informed of θ, so the collective action problem in incentive design disappears (Proposition 2), they still face a collective action problem in obtaining that information in the first place.

Two assumptions underlie all the analysis in this section. First, efficiency of auditing game equilibria is with respect to the joint utility of the principals. Since the article is about accountability and responsiveness of bureaus to these principals, the legislative principals are normatively privileged so this efficiency criterion is appropriate. Second, for all π and a given C, the auditing game may have multiple equilibria, including some in mixed strategies. Assume for a given C that the solution to the auditing game is from the set of Pareto dominant equilibria in that game. This simply stacks the deck in favor of efficiency with

19. This marks probably the most important departure between my conceptualization of oversight and that of McCubbins and Schwartz (1984). In their analysis, the principals' primary benefit of fire alarm oversight is that they can obtain information about bureaucrats nearly for free. One reconciliation of my framework with theirs is that θ is a type of information to which interest groups may not be privy, so that they would be unable to report it to the principals in a fire alarm.

respect to the coalition of principals in the auditing game, and therefore strengthens the inefficiency results. The results below would continue to hold for some other equilibrium selection, but more C values would lead to inefficiency.

The collective action problem in oversight is most stark for the case of perfectly informative audits so I begin with this case and then extend to any level of audit informativeness. Following this analysis I examine the welfare effects of increasing π, and discuss the benchmark case of integrated principals and no collective action problem.

3.1 Perfectly Informative Audits

Oversight by any one principal is perfectly informative when $\pi = 1$. In this case it is impossible for one principal's signal to be beneficial to the principals collectively, given purchase by the other principal. This induces, at the auditing stage, a one-of-two contributions game with complete information. The efficient outcome (with respect to the joint welfare of the principals) requires purchase either by one principal or by neither, depending on the audit cost C.[20]

P_i chooses $n_i = 1$ if and only if $C < W_i(1, n_j) - W_i(0, n_j)$. It is clear that for any strictly positive auditing costs, no equilibrium entails auditing by both principals with certainty: when $\pi = 1$, $W_i(1,1) = W_i(0,1)$. Whether any audits are done depends on $C \lesseqgtr W_i(1,0) - W_i(0,0)$, $i = 1, 2$.

For notational simplicity let $X_i(\pi) = W_i^a(1,0) - W_i^a(0,0)$ (the marginal public oversight benefit of auditing to P_i given that P_j does not audit) and $Y_i(\pi) = W_i^a(1,1) - W_i^a(0,1)$ (the marginal public oversight benefit of auditing to P_i given that P_j does audit). Let $Z_i(\pi,v) = W_i^b(1) - W_i^b(0)$ (the marginal private benefit of an audit to P_i). By definition $X_i(\pi) = \pi(\int_{\underline{\theta}}^{\overline{\theta}} V_i(\theta) dF(\theta) - \overline{V}_i(F))$ and $Y_i(\pi) = 0$ when $\pi = 1$.[21]

There are three ranges of C to consider, and one that leads to an inefficient level of auditing.

Proposition 3. *When $\pi = 1$, the level of auditing is inefficiently low with respect to the principals' joint welfare if and only if $C > \max_i(X_i + Z_i)$, and $C < X_1 + X_2 + Z_i$ for some i.*

20. The C region with efficient auditing in equilibrium has three equilibria: two are asymmetric, pure strategy equilibria and are efficient, while the symmetric, mixed equilibrium is inefficient. The region where the collective action problem arises has a unique, pure strategy auditing game equilibrium.

21. Arguments of these functions X, Y, and Z are included here for clarity but suppressed below because there is not much ambiguity about them.

Since only one audit can be useful to the principals, the inefficiency arises when the cost of an audit is less than its collective oversight benefit plus the private oversight benefit to some principal, but greater than the total (selfish) benefit to either principal. When $\pi = 1$ and C is either very small or very large, the level of auditing is what it would be if principals were integrated or accounted for the effects of their audits on each other. This proposition says that only for intermediate values of C can the collective action problem arise. The ones where it does arise are simply the ones for which auditing is selfishly irrational but still collectively beneficial.

3.2 Imperfectly Informative Audits and Symmetric Principals

In this subsection, I assume $V_1 = V_2$, $s_1 = s_2$, $W_1 = W_2$, and that the principals split equally the cost of the agent's information rent in the incentive subgame's most efficient equilibrium. This simplifies the analysis of general values of π and focuses attention only on the multiplicity of principals, rather than any differences between them. In this case all the marginal auditing benefits are equal across principals: $X_1 = X_2$, $Y_1 = Y_2$, and $Z_1 = Z_2$.

When the signal is imperfect in the sense that $\pi < 1$, the collective action problem in the previous subsection can still arise. However, an additional difficulty exists. Now it is possible that $W^a(1,1) > W^a(0,1)$, because of the probability $(1 - \pi)$ event that when only principal i audits, the audit is uninformative. Thus purchase by both principals can, in expectation, contribute to their collective welfare. That is, $Y = (\pi - \pi^2)(\int_{\underline{\theta}}^{\overline{\theta}} V(\theta)dF(\theta) - \overline{V}(F))$ and X is the same as above. Therefore, $X > Y$ when $\pi > 0$, and the marginal oversight benefit of the first audit exceeds that of the second. Further, if $\pi > \frac{1}{2}$, then $X > 2Y$, and if $\pi \leqslant \frac{1}{2}$, then $X \leqslant 2Y$.[22] Proposition 4 explicitly relates C to the efficiency of the oversight process.

PROPOSITION 4. *(i) With symmetric principals and* $\pi \in (\frac{1}{2},1]$*, the level of auditing is inefficiently low if and only if* $C \in [Y + Z,2Y + Z]$ *or* $C \in [X + Z,2X + Z]$*. (ii) With symmetric principals and* $\pi \in [0,\frac{1}{2}]$*, the level of auditing is inefficiently low if and only if* $C \in [Y + Z,2X + Z]$*.*

Proposition 3 (restricted to symmetric principals) then is a special case of this one; in that case $Y = 0$. With perfectly informative audits, only $C \in [X + Z,2X + Z]$ is a source of inefficiency. Note also that the welfare loss does not vanish as C approaches the boundaries of the C regions associated with inefficiency.

22. The regions where one audit is performed (whether one or two are efficient) have symmetric, inefficient equilibria in mixed strategies.

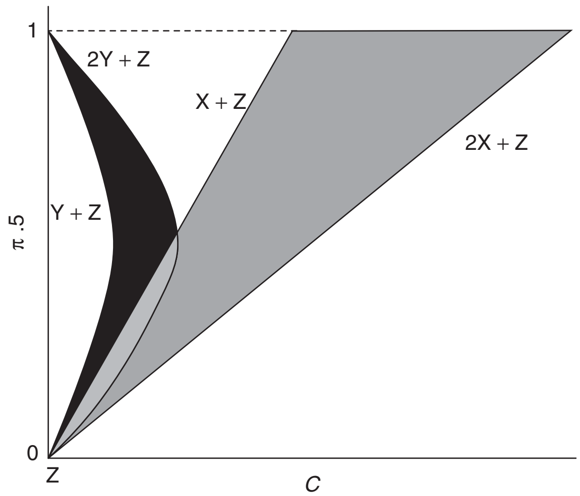

Figure 2. Efficiency of Auditing, given π and C

This result is illustrated in Figure 2. In the $C - \pi$ plane, the figure shows lines depicting when zero, one, or two audits are efficient and are in the most efficient auditing game equilibrium. In the large medium gray area on the right, one audit is efficient, but zero are performed. In the small light gray area, two audits are efficient, but zero are performed. In the black area, two audits are efficient and one is performed. In the white areas the level of auditing is efficient.

The regions of inefficiency depend on whether π is such that $X > 2Y$. When $X > 2Y$, π is relatively large. Then given one relatively effective audit and information leakage, a second one is not very valuable. On the other hand, the first audit is fairly valuable because it is likely to reduce the information asymmetry. This means there are audit costs where one audit is selfishly beneficial for a principal, but two audits are not even collectively beneficial enough to outweigh marginal cost. In other words, for an interval of Cs there are auditing game equilibria where one audit is efficient and one is performed. But if C falls far enough, two audits will again be efficient (the black region in Figure 2 for $\pi > \frac{1}{2}$); if it grows enough, no audits will be performed even though one would be efficient (the medium gray region in Figure 2 for $\pi > \frac{1}{2}$ – see Proposition 5).

When $X > 2Y$, the first audit is less likely to succeed, so a second audit is relatively more valuable. As C increases beyond $X + Z$ in this case, there is still a region where it is below $2Y + Z$ (the light gray region in Figure 2 for $\pi < \frac{1}{2}$). In that region no principal will audit, even though two audits are efficient. The induced auditing game for these Cs is simply a prisoners' dilemma. But the marginal benefit of the first audit exceeds that of the second, so as C increases enough to make two audits collectively inefficient, one audit can still be efficient even though none will be performed (medium gray region in the figure). On the other hand, even if C declines enough so one audit is selfishly rational, two audits may still be efficient (black region).

Thus collective action can weaken oversight as a tool of accountability. But a legislature faced with this problem, and not wishing to undo its multi-principal structure, could pursue several remedies. Two natural ones are increasing the effectiveness of the oversight technology and creating special oversight committees.

Proposition 5. *With symmetric principals and $\pi \in [0, 1)$, the range of Cs causing an inefficiently low level of auditing is increasing in π.*

A more effective auditing technology makes the collective action problem worse. As π increases, the benefit of an audit both to the auditing principal and to the other principal increases. But only the first benefit is taken into account in deciding to audit. As π increases the ignored externality also increases, making the collective action problem arise for more C values.

Proposition 6. *For $\pi \in (0,1)$, the range of Cs causing an inefficiently low level of auditing is increasing in v.*

Thus when oversight is a more important part of the committees' functions – say, by making them into special oversight committees – in a sense the collective action problem gets worse. As v increases, an audit becomes more valuable for the auditing principal as well as the other principal. The first effect is accounted for in a principal's audit decision, but the second is not. As v increases this ignored external benefit increases.

Aberbach (1990, Ch. 4) offers some suggestive evidence related to this. Oversight committees – for which oversight mattered most as part of the committee's function – tended to have less developed information networks and used them less effectively than non-oversight committees. In fact for the three groups of committees Aberbach presents (oversight, appropriations, substantive), oversight committees have the least well-developed information networks. New policy proposals tend not to originate in oversight committees, and thus their information networks are concentrated on oversight activity. This is not true for 'substantive' committees, which can use information networks for a variety of purposes, some more 'proprietary' than oversight (like the development of innovative policy).

Taking the previous two propositions together, some intuitive and practically important ways of eliminating the collective action problem in oversight do not necessarily have the desired effect. Both more effective audits and principals more concerned with oversight make the collective action problem worse.

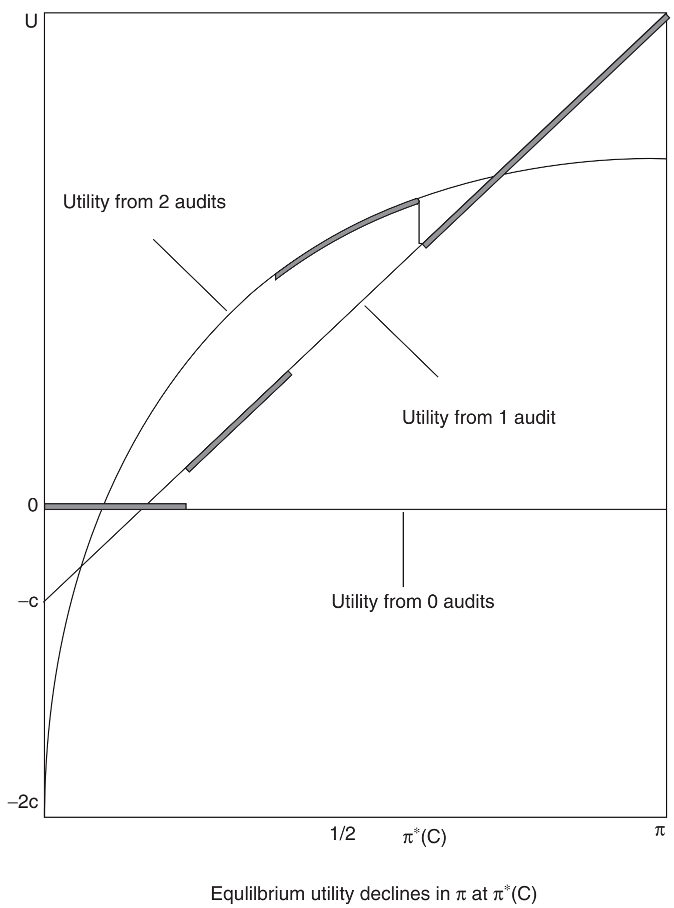

Figure 3. Utility from 0, 1, and 2 Audits as a Function of π (Equilibrium Utility shaded)

3.3 Welfare Effects of Audit Success Probability

Viewing the welfare gains from auditing as a pie, the previous results say that the share of the pie extracted declines in relative terms as π grows. It is also important to analyze the absolute size of the share extracted – that is, whether welfare can ever actually decrease as π grows. Interestingly in some cases increases in π are welfare reducing.

PROPOSITION 7. *If for any π the equilibrium selected in the induced auditing game is efficient, then for $C \in [Z, Y + Z]$ there is a $\pi^*(C)$ where the principals' equilibrium utility declines in π.*

The probability $\pi^*(C)$ is the $\pi > \frac{1}{2}$ such that for $\pi < \pi^*(C)$, two audits are efficient and two are performed, while for $\pi > \pi^*(C)$, two audits are efficient but one is performed.[23] At $\pi^*(C)$, an additional audit would therefore necessarily be welfare enhancing for the principals. But at $\pi^*(C) - \epsilon$, two audits are performed and are only infinitesimally less effective than at $\pi^*(C)$. This welfare dominates the equilibrium at $\pi^*(C)$.

Figure 3 illustrates this result. The solid curves depict the principals' collective utility from 0, 1, and 2 audits, and the bold segments trace out their utility in the most efficient auditing game equilibrium. For the indicated range of C values, there is a positive measure of πs above $\pi^*(C)$ such that collective welfare in equilibrium is higher at πs just below $\pi^*(C)$.

The discontinuity in equilibrium utility at $\pi^*(C)$ that causes this result cannot be avoided with a careful selection from the set of equilibria in the auditing game. In addition, since the auditing game equilibrium assumed in the proposition is as beneficial as possible to the principals jointly, the size of the jump in utility at $\pi^*(C)$ can only increase under a different equilibrium. For $\frac{1}{2} < \pi < \pi^*(C)$, the unique auditing game equilibrium entails two audits. For $\pi^*(C) < \pi < \pi^e$ (see Appendix, Proof of Proposition 7), the auditing game has three equilibria: two asymmetric ones where one principal audits with certainty, and a symmetric, mixed strategy equilibrium where each principal audits with some probability. The principals' utility in the mixed equilibrium is lower than their utility in the asymmetric equilibria. Moreover, it is obviously lower than their utility with two audits at π^*: the mixed equilibrium utility is an average over the utility from zero, one, and two audits. Thus, the same intuition as in Proposition 7 can be applied when the (inefficient) symmetric equilibrium is selected in this region. The key is the uniqueness of equilibrium below $\pi^*(C)$. Since Proposition 7 depends on moving from this equilibrium to the most efficient one above $\pi^*(C)$, moving from this equilibrium to a less efficient one above $\pi^*(C)$ will not help.[24]

If $\pi < \frac{1}{2}$, total welfare among the principals is clearly increasing in π. Extra audits are never welfare reducing in the auditing game's most efficient equilibrium. When $\pi < \frac{1}{2}$ increases in π never reduce the number of audits and sometimes increase it, while the probability that they are informative rises. Furthermore, if C is large enough, the principals' collective equilibrium welfare is at least weakly increasing in π.

23. There are πs less than $\frac{1}{2}$ with this property, but fixing C, one can only get further from them as π increases. When $\pi > \frac{1}{2}$, an increase in π lowers the marginal value of a second audit, so it is efficient only for lower C values. See Figure 2.

24. Furthermore, there are obviously other selections from the set of auditing game equilibria where increases in π reduce equilibrium utility for other values of π as well. For example, some equilibria have too many audits when $\pi > \frac{1}{2}$, and selecting these can reduce welfare for some increases in π.

3.4 Cooperative Auditing and Other Solutions to the Collective Action Problem

This collective action problem in oversight does suggest other remedies the principals might pursue to solve it, such as designing their own internal mechanisms. One might also think of principals in bilateral bargains, trading legislative concessions for oversight duties. However, this introduces a new aspect to the cost of oversight, in addition to probable information problems that plague efficient resolution of these bargains. The standard remedies of repeated play (which increases the benefits of long-term service on oversight committees, even when seniority would allow for a jump to 'substantive' committees) and an external enforcer in the legislative or party leadership also suggest themselves.

The upshot is that recognizing that there is a collective action problem to be solved helps to make the institutional structure of a legislature, and changes in that structure, more intelligible. For example, the Legislative Reorganization Act of 1946 led to a major expansion of committee staff as part of an effort to increase oversight (Smith and Deering, 1997). An institution-wide solution is difficult to understand without some sort of incentive problem at the individual level, and conventional explanations have emphasized that legislators benefit more from new policy proposals than from monitoring old ones (Oleszek, 2000). This model suggests another possibility, that a collective action problem led to an institution-wide effort to reduce oversight costs.

Assuming some effective remedy were found and the principals' decision problems were integrated as if they were a single principal, the problem is straightforward and useful as a benchmark. First, as noted in Section 3, the principals would face a second, not third, best situation in the incentive design phase. Second, while this would reduce the agent's effort distortion, oversight would still be beneficial, and oversight decisions would be fully efficient. There would not be any external benefit; all benefits would be accounted for in the oversight decision. For any π and C these decisions can be read off of Figure 2. The line $2X + Z$ and the curve $2Y + Z$ would determine when to purchase one and two audits respectively (Z would have to be modified to account for the new constituency).

Third, equilibrium utility would be continuous and increasing in π.[25] For all C, as π increases, the principals would switch from two audits to one just when the first audit was so effective that the second was inefficient. In general this equality will hold for any change in the number of audits; it is a simple consequence of efficiency, and ensures continuity.

4. Conclusion

The model in this article shows that multiplicity of legislative principals attenuates the control they have collectively over a bureaucratic agent. This is not because of

25. At least weakly: for C large enough, utility is constant in π because there is never an audit.

conflicts among the principals about goals or values they wish the bureaucrat to pursue, but rather is due to their organizational structure directly. Ideological conflicts among principals may certainly exist in reality and engender bureaucratic drift when they do. What this article clarifies is that multiplicity alone is sufficient to induce control loss, and conflicts among principals about bureaucrats' activities are not necessary. When multiple principals each offer incentives to a bureaucratic agent and can audit the agent's otherwise private information, it is possible for a collective action problem to make the level of auditing inefficiently low. On the other hand, this is by no means guaranteed. For some model parameters the level of auditing is efficient. Moreover, the more committees use their information for oversight purposes, and the more likely it is that the audit is informative, the worse the collective action problem will be. Finally, for some audit costs C, there is a range of audit success probabilities π where the principals' equilibrium utility is lower than it would be at lower values of π.

The model could be generalized by allowing non-zero correlation in audit results, different probabilities of informative audits for P_1 and P_2, and strategic auditors. It seems likely that these elements will complicate the exposition and alter the parameter ranges where different cases hold, but should not overturn the intuition.

Likewise, partial (rather than full) information leakage among principals could be treated. For example, one principal may only observe a (non-degenerate) garbling of another's audit results. Or, audit results per se may not leak at all, but may be partially revealed in the incentive subgame through the incentive schemes offered. Treating these cases will require a model of common agency with differentially informed principals. Moreover, this model did not consider strategic considerations in information transmission among the principals, which could also be interesting.

The model assumed simultaneous contract offers and audit decisions by the principals, but this is not responsible for the results. Martimort (1999) has shown that a similar distortion from common agency remains with sequential offers of incentive schemes. With sequential audits, the external benefits of a principal's audit would remain as well, given the information leakage. These external benefits would still increase with audit success probability and the importance of oversight. The non-monotonicity of equilibrium utility is also not tied to simultaneity. The result does not depend on coordination failures, an inefficiency that can be eliminated with sequential moves in public good provision problems. Instead it requires that the marginal benefit of a second audit decline in π when π is relatively large, and a selfish reckoning of when to audit or not. Given these features, then certain increases in π will make a second audit irrational for either principal before it becomes inefficient. For such increases in π there is a slightly smaller increase that leads in equilibrium to more but only slightly less effective audits. That is the key to the result, and it is not related to simultaneous moves.

The model developed here can apply, with some modification, to regulation as well – that is, conceiving of (multiple) bureaucrats themselves as principals and regulated firms as agents. Similar issues to those examined here (but different from those in other regulation models with auditing) arise when regulation is controlled by different agencies that can each solicit advice from third parties.

Appendix: Proofs

Proof of Proposition 3.

If $C > X_1 + X_2 + Z_i$, $i = 1,2$, there is no inefficiency: auditing is too expensive relative to its collective benefits to the principals, and neither principal will audit. If $C > \max_i(X_i + Z_i)$ and $C > X_1 + X_2 + Z_i$ for either $i = 1$ or 2, the audit should be purchased from a collective point of view, but will not be in equilibrium, which entails that neither principal audit. If $C > X_i + Z_i$ for either $i = 1$ or 2, the audit will be purchased by one principal, as efficiency requires.

Proof of Proposition 4.

(i) For $C \in [Y + Z, 2Y + Z]$ one principal will audit but an audit by both would be collectively beneficial. For $C \in [X + Z, 2X + Z]$ neither principal will audit, but by construction an audit by one would be collectively beneficial. If $C < Y + Z$, $C \in (2Y + Z, X + Z)$, or $C > 2X + Z$, the equilibrium level of auditing is the efficient level (two, one, or zero audits respectively). (ii) When $\pi \in [0, \frac{1}{2}]$, $X \leqslant 2Y$. This case adds the possibility that neither principal audits but two audits are collectively beneficial.

Proof of Proposition 5.

If $\pi > \frac{1}{2}$, the size of the range is $2X - X + 2Y - Y = X + Y = (2\pi - \pi^2)(\int_{\underline{\theta}}^{\overline{\theta}} V(\theta)dF(\theta) - \overline{V}(F))$. $\frac{\partial(X+Y)}{\partial\pi} = 2(\int_{\underline{\theta}}^{\overline{\theta}} v_i(\theta)dF(\theta) - \overline{V}(F))(1 - \pi) > 0$ for $\pi \in (\frac{1}{2}, 1)$. If $\pi \leqslant \frac{1}{2}$, the size of the range is $2X - Y = (\pi + \pi^2)(\int_{\underline{\theta}}^{\overline{\theta}} V(\theta)dF(\theta) - \overline{V}(F))$. $\frac{\partial(2X-Y)}{\partial\pi} = (1 + 2\pi)(\int_{\underline{\theta}}^{\overline{\theta}} V(\theta)dF(\theta) - \overline{V}(F)) > 0$.

Proof of Proposition 6.

Let $C_i^s = v[W_i^a(1, n_j) - W_i^a(0, n_j)] + (1 - v)[W_i^b(1) - W_i^b(0)]$ be the cost below which auditing is optimal for principal i. Let $C_i^p = v[W^a(1, n_j) -$

$W^a(0,n_j)] + (1-v)[W_i^b(1) - W_i^b(0)]$ be the cost below which auditing is collectively beneficial. $C_i^p - C_i^s = v[W_j^a(1,n_j) - W_j^a(0,n_j)]$ is increasing in v.

Proof of Proposition 7.

For any such C, let $\pi^*(C)$ be the $\pi \geqslant \frac{1}{2}$ such that for $\pi < \pi^*(C)$ two audits are efficient and two are performed, while for $\pi \geqslant \pi^*(C)$ two audits are efficient while one is performed. At $\pi^*(C)$, an additional audit would necessarily add a discrete welfare gain: the utility of 1 and 2 audits is equal at some $\pi^e > \pi^*(C)$, the utilities from 1 and 2 audits as a function of π are continuous and increasing in π, and the utility from 1 audit increases faster in π than the utility from 2 audits when $\pi \geqslant \frac{1}{2}$. But an outcome with two audits at $\pi^*(C)$ is arbitrarily well approximated at $\pi^*(C) - \varepsilon$ for ε small enough: at $\pi^*(C) - \varepsilon$ there is a discrete gain in welfare compared to $\pi^*(C)$ because of the extra audit, and an infinitesimal loss in welfare due to the less informative audits.

REFERENCES

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

Banks, Jeffrey S. (1989) 'Agency Budgets, Cost Information, and Auditing', *American Journal of Political Science* 33: 670–99.

Banks, Jeffrey S. and Barry R. Weingast (1992) 'The Political Control of Bureaucracy under Asymmetric Information', *American Journal of Political Science* 36: 509–24.

Baron, David (2000) 'Legislative Organization with Informational Committees', *American Journal of Political Science* 44: 485–505.

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

Bawn, Kathleen (1997) 'Choosing Strategies to Control the Bureaucracy: Statutory Constraints, Oversight, and the Committee System', *Journal of Law, Economics, and Organization* 13: 101–26.

Bendor, Jonathan, Serge Taylor and Roland van Gaalen (1987) 'Politicians, Bureaucrats and Asymmetric Information', *American Journal of Political Science* 31: 796–828.

Bernheim, B. Douglas and Michael Whinston (1986) 'Common Agency', *Econometrica* 54: 923–42.

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

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

Dixit, Avinash (1996) *The Making of Economic Policy: A Transaction Cost Politics Perspective*. Cambridge, MA: MIT Press.

Dodd, Lawrence C. and Richard L. Schott (1979) *Congress and the Administrative State*. New York: Macmillan.

Epstein, David and Sharyn O'Halloran (1999) *Delegating Powers*. New York: Cambridge University Press.

Fenno, Richard (1966) *The Power of the Purse: Appropriations Politics in Congress*. New York: Little, Brown.

Fiorina, Morris P. (1981) 'Congressional Control of the Bureaucracy: A Mismatch of Incentives and Capabilities', in Lawrence Dodd and Bruce Oppenheimer (eds) *Congress Reconsidered*, 2nd edn., Washington, DC: Congressional Quarterly Press.

Golden, Marissa Martino (2000) *What Motivates Bureaucrats? Politics and Administration During the Reagan Years.* New York: Columbia University Press.

Hammond, Thomas and Jack Knott (1996) 'Who Controls the Bureaucracy? Presidential Power, Congressional Dominance, Legal Constraints, and Bureaucratic Autonomy in a Multi-institutional Model of Policy Making', *Journal of Law, Economics, and Organization* 12: 119–66.

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

Kamien, Morton I. and Nancy L. Schwartz (1981) *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*. Amsterdam: North Holland.

Kaufman, Herbert (1981) *The Administrative Behavior of Federal Bureau Chiefs*. Washington, DC: The Brookings Institution.

Laffont, Jean-Jacques and Jean Tirole (1991) 'Privatization and Incentives', *Journal of Law, Economics, and Organization* 6: 1–32.

Laffont, Jean-Jacques and Jean Tirole (1993) *A Theory of Incentives in Procurement and Regulation.* Cambridge, MA: MIT Press.

Martimort, David (1992) 'Multiprincipaux avec Anti-selection', *Annales d'Economie et de Statistiques* 28: 1–38.

Martimort, David (1996) 'Exclusive Dealing, Common Agency, and Multiprincipals Incentive Theory', *RAND Journal of Economics* 27: 1–31.

Martimort, David (1999) 'Renegotiation Design with Multiple Regulators', *Journal of Economic Theory* 88: 261–93.

Martimort, David and Lars Stole (1999) 'The Revelation and Taxation Principles in Common Agency Games', mimeo: University of Chicago Graduate School of Business.

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

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

McCubbins, Mathew D., Roger G. Noll and Barry R. Weingast (1989) 'Structure and Process, Politics and Policy', *Virginia Law Review* 75: 431–82.

Miller, Gary J. and Terry M. Moe (1983) 'Politicians, Bureaucrats, and the Size of Government'. *American Political Science Review* 77: 297–322.

Mitnick, Barry (1980) *The Political Economy of Regulation*. New York: Columbia University Press.

Moe, Terry (1985) 'Control and Feedback in Economic Regulation: The Case of the NLRB', *American Political Service Review* 79: 1094–116.

Moe, Terry (1989) 'The Politics of Bureaucratic Structure', in John E. Chubb and Paul E. Peterson (eds), *Can the Government Govern?* pp. 267–329. Washington, DC: Brookings Instution.

Niskanen, William (1971) *Bureaucracy and Representative Government*. Chicago: Aldine–Atherton.

O'Connell, Anne Joseph (2006) 'The Architecture of Smart Intelligence: Structuring and Overseeing Agencies in the Post-9/11 World', *California Law Review* 94: 1655–744.

Ogul, Morris (1976) *Congress Oversees the Bureaucracy: Studies in Legislative Supervision*. Pittsburgh: University of Pittsburgh Press.

Oleszek, Walter (2000) *Congressional Procedures and the Policy Process.* Washington, DC: Congressional Quarterly Press.

Smith, Stephen and Chris Deering (1997) *Committees in Congress.* Washington, DC: Congressional Quarterly Press.

Spencer, Barbara (1980) 'Outside Information and the Degree of Monopoly Power in a Public Bureau', *Southern Economic Journal* 47: 228–33.

Stole, Lars (1997) 'Mechanism Design under Common Agency', mimeo: University of Chicago Graduate School of Business.

Ting, Michael (2001) 'Pulling the Trigger: On the Resource-Based Control of Bureaus', *Public Choice* 106: 243–74.

Ting, Michael (2003) 'A Strategic Theory of Bureaucratic Redundancy', *American Journal of Political Science* 47: 272–92.

Volden, Craig (2002) 'A Formal Model of Delegation under Separation of Powers', *American Journal of Political Science* 46: 111–33.

Wildavsky, Aaron (1978) *The Politics of the Budgetary Process*. New York: Little, Brown.

Wilson, James Q. (1989) *Bureaucracy: What Government Agencies Do and Why They Do It.* New York: Basic Books.

SEAN GAILMARD is an assistant professor of political science at the University of California Berkeley. His research interests include problems of accountability in administrative government, American political institutions, and formal political theory. [email: gailmard@berkeley.edu]