A STRATEGIC THEORY OF BUREAUCRATIC REDUNDANCY[1]

Michael M. Ting

Department of Political Science and SIPA
Columbia University
and
Center for Basic Research in the Social Sciences
Harvard University

mmting@latte.harvard.edu

February 27, 2003

[1]I wish to thank Jonathan Bendor, Daniel Carpenter, Joseph Cuffari, Sean Gailmard, Tom Hammond, and three anonymous referees for their particularly helpful comments and discussions. Seminar participants at the University of Chicago, Columbia University, Carnegie-Mellon University, and panel participants at the 2001 Annual Meeting of the Midwest Political Science Association also provided useful feedback. This paper was written at the University of North Carolina at Chapel Hill, and I thank the Department of Political Science there for their support.

Abstract

Do redundant bureaucratic arrangements represent wasteful duplication or a hedge against political uncertainty? Previous attempts at addressing this question have treated agency actions as exogenous, thus avoiding strategic issues such as collective action problems or competition. I develop a game-theoretic model of bureaucratic policy-making in which a political principal chooses the number of agents to handle a given task. Importantly, agents have policy preferences that may be opposed to the principal's, and furthermore may choose their policy or effort levels. Among the results are that redundancy can help a principal achieve her policy goals when her preferences are not aligned with the agents'. And contrary to the claims of supporters of streamlining bureaucracies, redundancy may even increase efficiency under some conditions. But redundancy is unnecessary if even a single agent has preferences relatively close to the principal's, or if the agent's jurisdiction can be terminated.

1. Introduction

Observers have long linked the effectiveness of government agencies to institutional design. Beginning with the seminal work of Landau (1969), redundant bureaucratic structures have been a prominent and recurring part of this discussion. Proponents have argued that redundancy improves the chances of some part of an organization succeeding in its task, and thus reduces the likelihood of failure. Opponents have questioned the efficiency of such arrangements, and have called for eliminating "wasteful duplication" and "overlap" in the bureaucracy. Others have also pointed out that increasing the number of components can lead to unpredictable interactions between them, ultimately hindering organizational effectiveness. To date, however, no equilibrium theory of redundancy and its alternatives has confronted these issues. As James Q. Wilson has summarized, "The problem, of course, is to choose between good and bad redundancies, a matter on which scholars have made little progress" (2000, p. 274).

In the field of reliability engineering, redundancy theory (*e.g.*, Barlow and Proschan 1965) establishes conditions under which functionally similar components can be added to improve a system's *ex ante* performance. For example, the United States Navy has long insisted on twin-engined designs for its sea-based aircraft because of the greater risks associated with flying over water. If these components have independently distributed probabilities of performing a task successfully, then their combination will be more likely to achieve a success together than either acting alone. Two components that are each independently 50% reliable will then be as effective as one that is 75% reliable. Of course, the tendency toward infinite redundancy is tempered by decreasing marginal returns and the costs of introducing additional components.

Bureaucratic redundancy theory extends this basic logic to political agencies, arguing that principals choose (or acquiesce to) multiple agents in order to increase organizational effectiveness. The American political system is replete with redundant arrangements, both across agencies (*e.g.*, Chisholm 1989) as well as within them (*e.g.*, Sapolsky 1972).[1] A few examples illustrate their scope:

- Welfare policy has long been administered by a patchwork of overlapping programs

[1] This paper speaks to both possibilities, so I adopt the neutral language of principals and agents throughout.

(some created in part by federalism), many of which embody different "theories" for addressing poverty.[2]

- Each branch of the military has its own "air force." Each service has a somewhat differentiated role; for example, only the Air Force has long-range bombers. However, they also perform many common tasks, such as the support of ground troops in battle.

- In the Department of Justice, the Office of Professional Responsibility and the Office of the Inspector General are jointly responsible for investigating internal malfeasance. The former was established internally in 1975 as part of a series of post-Watergate reforms, and the latter by Congress in 1989.

The arguments for redundancy are amenable to formalization, and more recent theories have explicitly treated the political principal's choice of agents as an optimization decision. Bendor (1985) establishes that the theory holds even if agents have non-independent success probabilities (though the benefits of redundancy are reduced if success probabilities are positively correlated). Non-independence might obtain, for instance, if units or agents share technical interdependencies. Heimann (1993, 1997) examines the implications of Type I and Type II errors and finds that redundant systems can prevent Type II errors, but raise the probability of Type I errors, while serial (non-redundant) systems have the opposite effect.[3]

An important assumption of these and other theories of redundancy is that the reliability of each agent is either given exogenously or determined by technical interdependencies. Perhaps due to the literature's roots in engineering, there are no *strategic* interdependencies, whereby individual agents might choose effort levels (hence their "reliability") based on their preferences and the actions of others.[4] This assumption has been questioned on two fronts. First, numerous historical studies of redundancy identify significant strategic behavior by component agencies—in some cases anticipated by the system's designers (*e.g.*, Maass 1951, Armacost 1969). Second, the primary intellectual rival to redundancy theory, known as *normal accident* theory (Perrow 1984, Sagan 1993), suggests that a system's reliability can depend on complex interaction among its components. Often these interactions raise the

[2]See Bendor (1985) and Salamon (1978) for a discussion of redundancies in welfare policy in the 1960s.

[3]Lohmann and Hopenhayn (1998) also examine this issue in a somewhat different context.

[4]A substantial literature in economics also examines the impact of different structures on organizational effectiveness. Many of these models similarly assume only technical (if any) interdependencies in agency effectiveness. Examples include Sah and Stiglitz (1986) and Sobel (1992). For a partial exception, see Kremer (1993).

likelihood of failure, in which case the system is best controlled by centralizing authority, rather than by increasing redundancy. While this theory is developed informally, its critique of the redundancy theory view of agent reliability naturally draws attention to other sources of interdependencies.

Given the formalization of redundancy theory and the implicit recognition of strategic interdependence in other aspects of the literature, it is puzzling—and counter to the trend of models of bureaucratic politics—that no strategic theory of redundancy has yet emerged. This omission has serious consequences because agencies frequently have incentives for subverting their principals' wishes, both because of divergent policy preferences (especially in presidential systems like the United States) and imperfect or incomplete information. Further, it is not obvious how strategic actors would affect the performance of redundant systems. Two common intuitions are of particular interest here. First, redundancy may cause collective action problems, where individual agents free ride off the efforts of others. Second, as Kaufman (1976) argues, redundancy may be indirect. An agent with exclusive jurisdiction over a task may face pressure from the possibility of another agent replacing it in the future. Thus "latent" competition may over time serve many of the same goals as redundant agents.

This paper connects the classical redundancy arguments with these strategic intuitions through a simple game-theoretic model. The critical feature of the model is that agents are treated as strategic actors in their own right, thus requiring both principal and agent choices to be derived endogenously. In light of the theoretical advances in the study of bureaucratic politics, such an effort naturally entails certain compromises. In order to remain as close as possible to the original arguments on organizational redundancy, the model is highly stylized, and abstracts away from many institutional features. However, it is also tractable and can be expanded in various ways to incorporate greater detail.

In the basic game, there is a single policy domain with two discrete outcomes, which may be labeled "good" and "bad." The principal always wishes to attain the former outcome, while agents have possibly divergent preferences over the outcomes. The principal chooses the number of agents, while each agent chooses an unobservable and costly effort level or policy, which is operationalized as a probability of "succeeding." As is standard in models of redundancy, a good outcome results if at least one agent succeeds. The optimal level of redundancy therefore depends on the agents' incentives to exert high effort levels in the presence of other agents. Several extensions to the basic model are also developed. First, a repeated version of the game examines conditions under which latent redundancy can achieve the principal's goals more efficiently. Additionally, two extensions examine the effects of different policy technologies, one in which good outcomes require more than one agent to succeed, and the other with policy externalities (*i.e.*, a form of technical interdependence).

The results of the model provide a basis for understanding the organizational structures that political principals would either design or accept. Some of the predictions are quite intuitive. Agents' policy choices impose externalities on their redundant partners in equilibrium, even when the policy technology is *technically* statistically independent (thus satisfying a basic condition for classical redundancy). Typically these take the form of a collective action problem: all agents exert less effort than they would if acting alone. However, in the case where a good outcome requires multiple agents to succeed, positive externalities can occur.

More surprisingly, the effect of the negative externalities on the collective outcome (and thus the principal's choice of organizational structure) varies. A redundant system tends to help the principal to achieve her goals when the agents' policy preferences are far from her own. But redundancy does not help her otherwise. When a "friendly" agent exists the collective action problems are particularly acute, and additional agents can *reduce* the chances of a good outcome. Moreover, in a repeated setting where players do not discount future payoffs too heavily, latent redundancy works. The principal can induce even a relatively unfriendly agent to choose her ideal policy, simply by terminating it if it does not perform well. In this setting, agents essentially compete against themselves.

The model naturally contributes to a substantial body of work concerned not only with redundancy, but with the institutional determinants of bureaucratic effectiveness more generally.[5] It also adds a new argument to theories of bureaucratic delegation. Much of this literature has emphasized the importance of preference convergence in delegation decisions, although the results here demonstrate some conditions under which a principal will delegate even to unfriendly agents.[6] Finally, the model's focus on collective action in a highly general hierarchical relationship suggests applications to non-government organizations as well.

The results will also be useful for developing empirical tests of the relationship between principal and agent preferences and organizational design. To date, an extensive literature

[5] There are numerous developments of this theme; for examples see Moe (1989), Wilson (2000), and Carpenter (2001).

[6] For discussions on bureaucratic delegation, see Calvert, McCubbins, and Weingast (1985), Kiewiet and McCubbins (1991), Epstein and O'Halloran (1994, 1999), Bawn (1995), and Gailmard (2002a).

based on case studies has covered a range of issues related to redundancy and normal accidents (*e.g.*, Lerner 1986, Rochlin, LaPorte, and Roberts 1987, LaPorte and Consolini 1991). By contrast, some of the hypotheses regarding jurisdictional assignments developed here will be testable on a broader scale, in some cases using readily available data.

Before proceeding, I note two caveats which, given the foundation laid by this work, may be viewed as opportunities for generalization. First, while the policy technologies used here maintain maximum comparability with existing theories of redundancy, other formulations are of course plausible. For example, in Brehm and Gates (1997), agents have the option of "sabotage" in addition to working or shirking. The issue is a general one, as agent effort may translate into outcomes in a variety of ways depending on the policy domain (*e.g.*, Esteban and Ray 2001, Hirshleifer 1983, Holmstrom 1982).

Second, while the model approaches redundancy from the perspective of a "reliability-maximizing" principal, other rationales and explanations for redundancy exist. Perhaps most significantly, entrepreneurial agents can move unilaterally into new policy areas. In this case, the model suggests how a principal would react to such jurisdictional shifts, and in turn the kinds of policy expansion that agents will undertake. Another possibility is that redundancy arises from political compromise (*e.g.*, Moe 1989). In this case, agencies with different preferences or procedures may be used to buy the support of opposing interests. Next, borrowing an analogy from private markets, redundancy may induce efficiency-improving competition among agents (Niskanen 1971, Miller and Moe 1983, Donahue 1991, Osborne and Gaebler 1992).[7] Finally, a principal may use redundancy to solve adverse selection, as opposed to moral hazard problems.[8] Most of these explanations are not necessarily incompatible with the approach taken here, but would entail non-trivial changes in the underlying model.

The paper proceeds as follows. In the next section, I motivate the theory with an example that contrasts strategic and non-strategic versions of redundancy. Section 3 lays out the direct redundancy model, which is the basis of the games developed here. Section 4 derives the main results for this model. Latent redundancy is explored in Section 5. Section 6 develops the extensions on collaborative environments and policy externalities. Section 7

[7] Indeed, in a later work, Landau (1991) himself invokes a market analogy for why redundant agents should be more efficient.

[8] Examples might include principals confronting a Condorcet Jury problem (*e.g.*, Austen-Smith and Banks 1996), or attempting to monitor corruption (*e.g.*, Montinola 2002).

proposes some empirical applications and extensions for future research and concludes.

2. Strategic and Non-Strategic Redundancy: An Example

To illustrate very simply why a strategic perspective on redundancy matters, consider the following example. A principal can choose up to two agents to share jurisdiction over a single task. The task might be the gathering of intelligence that potentially involves the "turf" of both civilian and military agencies (*e.g.*, Federal Bureau of Investigation, Central Intelligence Agency, National Security Agency, military intelligence, State Department). Each agent may individually succeed or fail. For simplicity, assume that the probability of success, or *reliability*, is $r \in (0,1)$ (the example generalizes easily to heterogeneous reliability rates). This probability is independent across agents, which may reflect the fact that agents use different methodologies in their work.[9] The overall outcome of the task is "good" for the principal if either agent succeeds in gathering the crucial intelligence, and "bad" otherwise. The principal receives one unit of utility for a good outcome, and zero for a bad one.

In a non-strategic environment, each agent's reliability is unaffected by the presence of other agents. Thus, the probability of a good outcome is simply r with one agent, and $r + (1-r)r > r$ with two. The redundant agent therefore raises the probability of a good outcome, and the principal's expected utility, by $(1-r)r$.

Now suppose that agents are strategic; that is, they receive payoffs from outcomes and choose effort levels. Agents receive one unit of utility for a good outcome, and zero otherwise. Each agent can now choose either to *work* or *shirk*. Its reliability, if it should work, is r; otherwise, it is zero. Working imposes a cost $c \in (0,r)$, while shirking costs nothing. Thus, each agent shares the principal's desire for a good outcome, and would exert effort if acting alone. However, each also understands that its effort is wasted if the other agent succeeds. The following payoff matrix summarizes the game.

Figure 1: A Simple Redundancy Game

[9] As will be clear in the subsequent analysis, the independence of agent effort can contribute greatly to the success of redundant systems. This implies—contrary to the argument of some reformers—that agencies should *not* always collaborate or share information in their efforts.

	Work	Shirk
Work	$2r-r^2-c, 2r-r^2-c$	$r-c, r$
Shirk	$r, r-c$	$0, 0$

Notice that the outcome of the non-strategic model, (Work, Work), is a Nash equilibrium if and only if $2r-r^2-c \geq r$, or $r \in [\frac{1}{2}-\frac{\sqrt{1-4c}}{2}, \frac{1}{2}+\frac{\sqrt{1-4c}}{2}]$. In other words, r must be moderate and c low for both agents to work. When r is low, the marginal contribution of the second agent is too low to justify the cost, while a high r leaves little room for improvement. In both cases, agents face a collective action problem, and the Nash equilibrium consists of one agent working and the other shirking.

The game extends easily to more agents. Generally, with $n > 1$ agents, all agents will work in a Nash equilibrium only if $1-(1-r)^n-c \geq 1-(1-r)^{n-1}$. It is then straightforward to demonstrate that for a given c, the maximum number of agents that will work in a pure strategy equilibrium is given by the highest n satisfying $c \leq r(1-r)^{n-1}$. Beyond this threshold, additional agents have no effect on the probability of a good outcome. This expression also implies that as costs increase, the number of working agents must decrease.

This example shows that the classic redundancy model is not robust to the introduction of strategic agents. Even with the simplest possible generalization to a game theoretic formulation, the prediction of redundancy theory is a special case occurring only under a certain set of parameter values. In this particular example, agents face the familiar collective action problem. In the games developed in the following sections, other effects of strategic interaction between principals and agents will also become evident.

3. Direct Redundancy: A Model

The direct redundancy game greatly generalizes the game-theoretic model of the preceding example. It addresses the basic question raised by redundancy theory: in a given period, how many agents should a principal choose? The model also forms the basis of the extensions developed in Sections 5 and 6.

Environment and Players. The game describes a single period of policy-making under uncertainty between two kinds of players: a principal (P) and up to $N > 0$ agents, denoted A1, …, AN. P may be construed as any actor with authority over the number of units to assign to a particular task. Such actors are found at different points throughout the executive and legislative branches. For example, Congress, department heads, governors, and the president are all empowered to some degree to create bureaucratic agencies or assign their jurisdictions. Alternatively, P can represent an agency head deciding the number of divisions to assign to a given job. I therefore use the general term 'agent' to refer to any unit that is given jurisdiction over the task in question.

P chooses n, the number of agents, where $0 \leq n \leq N$, but faces a moral hazard problem because she cannot dictate their policy choices directly. Each agent Ai shares jurisdiction over the task in question, and can independently set an unobservable effort level or policy, $\phi_i \in [0, 1]$. Denote the n-element vector of policies ϕ. The order in which agents are created is fixed, with A1 first, A2 second, *etc.* Thus, P cannot give A2 jurisdiction over the policy if $n = 1$. This may reflect the possibility that an "incumbent" agent already exists for the given task.

All players are interested in an observable outcome $x \in \{0, 1\}$ (where $x = 1$ corresponds to the "good" outcome of the previous section). Outcomes are determined as follows. Each Ai's policy results in a *success* with probability ϕ_i, and a *failure* otherwise. If any agent succeeds, then $x = 1$; otherwise, $x = 0$. Thus,

$$\Pr\{x = 1\} = 1 - \prod_{j=1}^{n} (1 - \phi_j). \tag{1}$$

This policy technology is standard in theories of redundancy, but is relaxed in Section 6.1. It is simple but highly general, requiring mainly that some agent's effort be exerted to achieve the good outcome (otherwise, the good outcome is the status quo and no agents would be necessary). Thus, x may be considered a summary statistic for whether an agency's policy outcome was satisfactory to P; for example, whether air fatalities decrease after a change in Federal Aviation Administration policy, or whether Food and Drug Administration-approved drugs have few unpredicted side effects. Note however that there is only one type of failure in this framework. There are also no policy externalities; that is, each agent's *own* probability of succeeding is independent of the other agents' actions. This assumption is relaxed in Section 6.2, but the subsequent analysis will show that agents can generate externalities endogenously through their strategic behavior.

Payoffs. All preferences are common knowledge. P has linear preferences over policy outcomes, receiving x, and therefore desires policies that maximize the likelihood that $x = 1$.

P additionally pays a cost $k \geq 0$ for each agent. This represents a fixed level of resources (or "budget") that must be committed before an agent can participate in policy-making.[10] Thus, P's expected utility is:

$$u(\phi) = 1 - \prod_{j=1}^{n}(1 - \phi_j) - nk.$$

Each Ai also has linear preferences over outcomes, receiving $w_i x$, where $w_i \in \Re$. Additionally, setting policy incurs a quadratic cost $c_i(\phi_i) = m_i \phi_i^2$, where $m_i > 0$. Unless otherwise noted, I assume throughout that $\frac{w_i}{m_i} \geq \frac{w_j}{m_j}$ for all $i < j$, so that sequentially prior agencies desire the highest effort or policy levels. As will be clear in the analysis below, this is exactly the sequence that P would choose if it could. Ai's expected utility is then:

$$v_i(\phi) = w_i \left[1 - \prod_{j=1}^{n}(1 - \phi_j)\right] - c_i(\phi_i). \tag{2}$$

Note that since the product term of (2) is linear in ϕ_i and $c_i(\cdot)$ is convex, $v_i(\cdot)$ is concave in ϕ_i. The expression makes clear that each agent's ideal policy is endogenous, except in the trivial case ($w_i \leq 0$) where it has a dominant strategy of not exerting effort. In general, it will balance marginal policy utility with marginal cost, and when $n > 1$, it may shirk and let another agent undertake the costly effort.

Sequence. Game play proceeds as follows.

(1) *Agent Selection.* P chooses the number of agents n.

(2) *Policy-Making.* Each agent Ai simultaneously chooses an unobservable policy $\phi_i(n)$.

(3) *Policy Outcome.* Nature randomly determines outcome x according to (1).

4. Direct Redundancy: Results

A subgame-perfect Nash equilibrium of this game consists of an optimal number of agents, n^*, and N n-vectors ($n = 1, \ldots, N$) of policies, each denoted $\phi^*(n)$. When P is indifferent between different values of n, it is assumed that the tie is broken in favor of the lowest number of agents. Denote Ai's optimal policy choice in the n-agent subgame $\phi_i^*(n)$, and let $\Phi^*(n)$ represent the equilibrium probability that $x = 1$ in that subgame.

[10] By imposing fixed agent costs that are not linked with the cost of policy, the model focuses on the policy-making, and not the budgeting side of the principal-agent relationship.

4.1. An Example: Up to Two Agents

To build some intuition I begin by examining the reduced game where $N = 2$. In this case I also relax the assumption that $\frac{w_1}{m_1} \geq \frac{w_2}{m_2}$, so that second agent can be less "friendly" to P's interests than the first. Thus, the case corresponds well to situations in which principals face short-term constraints on creating or choosing agents.

Policy Choice. Solving backwards, when $n = 1$, A1 simply balances marginal policy utility and marginal cost. Differentiating (2), its ideal policy is:

$$\phi_1^*(1) = \begin{cases} 0 & \text{if } w_1 \leq 0 \\ \frac{w_1}{2m_1} & \text{if } w_1 \in (0, 2m_1) \\ 1 & \text{if } w_1 \geq 2m_1. \end{cases}$$

As intuition suggests, equilibrium policy is increasing in A1's utility for a good outcome (w_1), and decreasing in its marginal cost (m_1). It will be useful to define $r_1 = \phi_1^*(1)$, and r_i for all other Ai in an analogous manner. This is the policy that each Ai would choose if acting alone, and is thus an analog of the (non-strategic) reliability of an agent.

Now consider the subgame in which $n = 2$. It will be convenient to denote with a subscript $-i$ parameters belonging to the agent that is not Ai. Then since $v_i(\cdot)$ is concave, the following first-order condition is sufficient to characterize Ai's best response:

$$w_i(1-\phi_{-i}) - 2m_i\phi_i = 0. \tag{3}$$

Simplifying and taking feasibility constraints into account, we obtain Ai's best response:

$$\phi_i^*(2) = \begin{cases} 0 & \text{if } w_i \leq 0 \\ \min\{\frac{w_i(1-\phi_{-i})}{2m_i}, 1\} & \text{if } w_i > 0. \end{cases} \tag{4}$$

Note that at an interior solution, each agent's best response is linear and decreasing in the other's policy. Corner solutions are possible under three circumstances. If $w_i \leq 0$, then exerting no policy effort is a dominant strategy and $\phi_i^*(2) = 0$. If w_i is sufficiently large and $\phi_{-i} < 1$, then Ai will assure itself of $x = 1$ by choosing $\phi_i^*(2) = 1$. Finally, if $\phi_{-i} = 1$, then Ai is unable to affect the outcome and chooses $\phi_i^*(2) = 0$. Otherwise, for moderate values of w_i and w_{-i}, an interior solution exists. Combining expressions, the following comment derives the equilibrium policies.

Comment 1. Policy with Redundant Agents. If $N = 2$, the equilibrium policies are:

$$\phi_i^*(2) = \begin{cases} 0 & \text{if } w_i \leq 0, \text{or } w_i \in (0, 2m_i) \text{ and } w_{-i} \geq 2m_{-i}, \\ & \text{ or } i = 2 \text{ and } w_1 \geq 2m_1 \\ \frac{w_i(2m_{-i} - w_{-i})}{4m_1 m_2 - w_1 w_2} & \text{if } w_i \in (0, 2m_i) \text{ and } w_{-i} \in (0, 2m_{-i}) \\ \frac{w_i}{2m_i} & \text{if } w_i \in (0, 2m_i) \text{ and } w_{-i} \leq 0 \\ 1 & \text{if } w_i \geq 2m_i \text{ and } w_{-i} \leq 2m_{-i}, \text{or } i = 1 \text{ and } w_1 \geq 2m_1. \end{cases} \quad \blacksquare$$

Proof. Proofs of all comments and propositions are in the Appendix. ■

While this equilibrium is not unique, it is very nearly so. When $w_i \geq 2m_i$ for both agents (with the inequality strict for at least one), there are two equilibria: $\phi_1^*(2) = 1$ and $\phi_2^*(2) = 0$, and $\phi_1^*(2) = 0$ and $\phi_2^*(2) = 1$. To maintain continuity with the $n = 1$ subgame, I simply impose the former as the solution.[11]

Table 1: Redundant and Non-Redundant Systems

$w_1 = 1.1$, $m_1 = 1$, $m_2 = 2$

n	Case	$\phi^*(n)$	Total Cost	$\Phi^*(n)$
1	(a)-(c)	(0.55)	0.3025	0.55
2	(a) $w_2 = 3$ (b) $w_2 = 2$ (c) $w_2 = 1$	(0.234, 0.574) (0.379, 0.310) (0.478, 0.130)	0.7148 0.3365 0.2628	0.674 0.572 0.546

As Table 1 illustrates, the comparative statics of these subgames behave in a sensible manner. When policy choices are interior, Ai's policy is decreasing in m_i and increasing in w_i. Each agent's equilibrium policy is also increasing in the *other* agent's marginal costs (m_{-i}) and decreasing in its marginal utility for a good outcome (w_{-i}). Finally, the addition of A2 always weakly reduces A1's policy choice. Thus, to some degree increasing the number of agents introduces a collective action problem. But this reduction does not necessarily reduce the probability that $x = 1$. The following comment characterizes the effect of increasing the number of agents on $\Phi^*(n)$.

Comment 2. Redundancy and Effectiveness. $\Phi^*(2) < \Phi^*(1)$ if and only if $r_1 \in (\frac{1}{2}, 1)$ and $r_2 \in (0, \frac{4m_1(w_1 - m_1)}{w_1^2})$. ■

To understand this result, note that the addition of A2 *always* weakly reduces A1's

[11] In the "knife-edge" case where $w_i = 2m_i$ for both agents, there is a range of equilibria, including the two identified here. I impose the same solution in this case.

incentive to produce policy. Because of increasing marginal costs, this reduction is most acute when A1 would choose a "high" policy. Further, when w_2 is sufficiently low A2's success probability will be too small to offset the decrease in A1's effort. Thus, as case (c) of Table 1 illustrates, the collective action problem has a particularly serious bite when the agents' policy preferences are far apart.

Comment 1 also allows us to evaluate the effect of redundancy on policy efficiency. Let the measure of efficiency be $\frac{\Phi^*(n)}{\sum_i c_i(\phi_i^*(n))}$, or the increment in the probability of a good outcome per unit cost. Then efficiency drops from 1.82 to approximately 0.94 in case (a) of Table 1. Here the high cost of A2's effort reduces efficiency and indeed produces some "wasteful duplication." But when A2 chooses a relatively low policy level (as in case (c)), efficiency actually *increases*. This occurs because low policies are more efficient (since marginal costs are increasing), so "low" policies by both agents incur lower marginal costs than a "high" policy by A1.[12] The next result formalizes this intuition.

Comment 3. Inefficient Redundancy. If $r_1 \in (0, \frac{1}{2})$ and $r_2 \in [\frac{w_1(2m_1 - w_1)}{4m_2(m_1 - w_1)}, 1)$, then two agents are less efficient than one. ■

Thus inefficiency will tend to occur if A2 is "friendly" but uses a costly policy technology.[13]

Finally, it is worth noting from a welfare perspective that if $w_i > 0$ for all Ai, each agent benefits from the addition of a redundant partner. This is because each agent then prefers an outcome of one to zero, and chooses a strictly positive policy in equilibrium. As a result, even if Ai chose the same policy that it would have in isolation, it would do strictly better.

The Optimal Number of Agents. Given the policy responses, P simply chooses n to maximize the $\Pr\{x = 1\}$, net of the costs of adding new agents, k. P will thus prefer a redundant system to a non-redundant one if $k \leq \Phi^*(2) - \Phi^*(1)$. Returning to case (a) of Table 1, the marginal benefit of adding A2 is 0.124. Thus if $k < 0.124$, P chooses two agents, and if $k \in [0.124, 0.55)$, P chooses one. In case (b), the lower value of w_2 decreases the marginal value of the second agent, and thus redundant agents are chosen only if $k < 0.022$.

Clearly, a redundant system would not be chosen if $k \geq 0.5$. Nor would it be chosen

[12] An alternative measure of policy efficiency is $\sum_i \phi_i^*(n) / \sum_i c_i(\phi_i^*(n))$, or the cost per unit of policy produced. Then in case (a) of Table 1, efficiency drops to only 1.13 with two agents. Because this statistic does not take the effect of duplicated efforts into account, the drop in efficiency is less pronounced.

[13] Note that in the general model, where it is assumed that $\frac{w_1}{m_1} \geq \frac{w_2}{m_2}$, the condition of Comment 3 holds only if $m_2 \geq \frac{m_1(2m_1 - w_1)}{2(m_1 - w_1)}$.

if it would *reduce* $\Pr\{x = 1\}$. Comment 2 is therefore a special case (where $k = 0$) of the following result, which characterizes n^*.

Proposition 1. If $N = 2$,

$$n^* = \begin{cases} 0 & \text{if } k \geq \max\{\phi_1^*(1), \frac{\phi_1^*(2)+(1-\phi_1^*(2))\phi_2^*(2)}{2}\} & \text{(i)} \\ 1 & \text{if } k \in [\phi_1^*(2) + (1-\phi_1^*(2))\phi_2^*(2) - \phi_1^*(1), \phi_1^*(1)) & \text{(ii)} \\ 2 & \text{if } k < \min\{\phi_1^*(2) + (1-\phi_1^*(2))\phi_2^*(2) - \phi_1^*(1), \frac{\phi_1^*(2)+(1-\phi_1^*(2))\phi_2^*(2)}{2}\}. & \text{(iii)} \end{cases} \quad \blacksquare$$

Thus, n^* depends on the relative effectiveness of the agents, and is decreasing in k. Interestingly, there are two cases in which n^* does not decrease from 2 to 1 to 0 as k increases. First, interval (iii) is empty when $\phi_1^*(2) + (1-\phi_1^*(2))\phi_2^*(2) < \phi_1^*(1)$; *i.e.*, when two agents do worse than one (see Comment 2). Second, interval (ii) may also be empty if A1 is an "unfriendly" agent (*i.e.*, r_1 is low). In this event, P chooses two agents if w_2 is sufficiently high relative to k, and none otherwise.

An alternative way of stating Proposition 1 would be with respect to w_i and m_i. This would be considerably more complicated because there are numerous corner conditions to be taken into account. However, to get a sense of the comparative statics with respect to these parameters, the following figure illustrates the relationship between w_1 and n^* for $w_2 \in (0, m_2)$.

[Figure 2 about here.]

In the figure, the solid lines represent the cost thresholds that make P indifferent between two values of n. It is straightforward to verify in general that these cutlines intersect at some $\hat{w}_1 > 0$. These lines are gray where the cutpoint is irrelevant to P's decision. For example, for $w_1 < \hat{w}_1$, the two cutlines determining whether P prefers one agent to zero or two are irrelevant, because one agent is P's least-preferred option (this situation corresponds to the case where interval (ii) in Proposition 1 is empty). The figure also indicates that the equilibrium relationship between agent preferences and n^* is not always monotonic. An increase in w_1 often makes one agent more desirable, and it is easily shown that for w_1 sufficiently high, $n^* = 1$. But for some moderate values of k, n^* can increase from zero to two and then decrease to one as w_1 increases. Thus the following, somewhat rougher monotonicity conditions hold: A1 will be part of a redundant system only if it is sufficiently "unfriendly," and it will not be part of a redundant system if it is sufficiently "friendly." A similar analysis holds for agent costs.

4.2. The General Case

I now establish the principal results of the general game, where N is arbitrary (and possibly infinite).

Policy Choice. Solving backwards, since $v_i(\cdot)$ is concave, the following first-order condition characterizes Ai's best response:

$$w_i \prod_{j \neq i}(1-\phi_j) - 2m_i\phi_i = 0.$$

And thus,

$$\phi_i^*(n) = \begin{cases} 0 & \text{if } w_i \leq 0 \\ \min\{\frac{w_i \prod_{j \neq i}(1-\phi_j)}{2m_i}, 1\} & \text{if } w_i > 0. \end{cases} \tag{5}$$

This expression generalizes (4). As in the two-agent case, redundancy introduces a collective action problem, thus causing all agents to choose (weakly) lower policies when $n > 1$ than they would in isolation. It follows immediately that success probabilities in this game are lower than in the non-game theoretic version, where all agents simply choose r_i. This is illustrated in Figure 3, which compares success probabilities between the game theoretic and classical redundancy models under two configurations of $\{r_i\}$. Additionally, $\phi_i^*(n)$ is interior unless $w_i \leq 0$ or w_i is sufficiently large and $\phi_j < 1$ for $j \neq i$. Finally, each agent's optimal policy response in any subgame is unique and pure: (5) generalizes straightforwardly to any profile of mixed strategies, and $\phi_i^*(n)$ is single-valued for any failure probability collectively implied by the other agents' strategies.

[Figure 3 about here.]

Closed forms for equilibrium policies are considerably more difficult to derive when $n > 2$, since the system of equations defined by (5) is non-linear. However, the main results of the model do not require such a characterization, and many of the comparative statics carry over directly from the $n = 2$ subgame. The following comment establishes some of these.

Comment 4. Equilibrium Characteristics.

(i) (Uniqueness) The subgame-perfect Nash equilibrium is unique.

(ii) (A1's policy) $\phi_1^*(n)$ is decreasing in n and increasing in r_1.[14]

[14] This result depends in part on the assumption that the "friendliest" agent is introduced first. If this assumption were dropped, the same result holds for that agent's effort after its introduction.

(iii) (System effectiveness and r_1) $\Phi^*(n)$ is increasing in r_1.[15] ■

Part (i) assumes (as in the $N = 2$ case) that the agents coordinate on the equilibrium where A1 chooses $\phi_1^*(n) = 1$ when $r_i \geq 1$ for more than one agent. Intuitively, it obtains because each agent's best response (5), while not linear, is sufficiently "flat." Thus, much as n linearly independent hyperplanes in n-space intersect at a point, the best response functions of all agents meet at a unique policy vector. Part (ii) can be extended easily (if tediously) to show that $\phi_i^*(n)$ is increasing r_i for all Ai.

Table 2 provides examples of equilibria for four multiple-agent subgames, corresponding to 'Model 1' of Figure 3. Within and across equilibria, policies are increasing in r_i (*i.e.*, decreasing in costs and increasing in utility for a successful outcome). Additionally, the effect of each additional agent tends to be decreasing (although as 'Model 2' of Figure 3 illustrates, this is not always the case for low n). Finally, redundancy can increase efficiency from the agents' perspective (though not the principal's). Just as in Comment 3, this happens because the addition of each new agent usually forces down the policies of incumbent agents. Given similar cost functions (m_i), the new vector of (lower) policies then incur lower marginal costs. In Table 2, for example, total policy costs *decrease* as n increases, even while $\Phi^*(n)$ rises for $n \geq 3$.

Table 2: Many Agents

$w_1 = 1.5$, $w_2 = 1.0$, $w_3 = 1.0$, $w_4 = 0.75$; $m_i = 1$ for all Ai

n	$\phi^*(n)$	Total Cost	$\Phi^*(n)$
1	(0.75)	0.563	0.750
2	(0.60, 0.20)	0.400	0.680
3	(0.47, 0.21, 0.21)	0.307	0.668
4	(0.41, 0.20, 0.20, 0.14)	0.270	0.677

Table 2 and Figure 3 illustrate two counterintuitive phenomena that are central to the main results of this section. First, $\Phi^*(n)$ is not necessarily monotonic in n. In the table and 'Model 1' of Figure 3, the addition of A2 and A3 actually reduces $\Phi^*(n)$, and $\Phi^*(n) < \Phi^*(1)$ for $n \leq 10$. Second, while A1's policy is always decreasing in n, other agents' efforts are not

[15] Interestingly, this result does not necessarily hold for agents other than A1. Related to the discussion of Proposition 2 below, $\Phi^*(n)$ is increasing in r_i ($i \neq 1$) over any range for which $\phi_1^*(n) < 0.5$, and decreasing in such r_i over any range for which $\phi_1^*(n) > 0.5$.

necessarily so. That is, in some cases agents other than A1 face less of a collective action problem—and thus increase their effort—as n increases.

These phenomena are actually closely related through the agents' equilibrium behavior. To see why they occur, observe that (5) implies that for each Ai:

$$\phi_i^*(n)(1-\phi_i^*(n)) = r_i(1-\Phi^*(n)). \tag{6}$$

Two implications follow directly from (6). First, since $\phi_i^*(n)(1-\phi_i^*(n)) \leq \frac{1}{4}$, $\Phi^*(n)$ is bounded from below by $1-\frac{1}{4r_1}$. Second, and more importantly, the equilibrium ratio between $\phi_i^*(n)(1-\phi_i^*(n))$ and $\phi_j^*(n)(1-\phi_j^*(n))$ is constant for all n. While this expression does not fully characterize equilibrium policies, it is a powerful necessary condition that pins down a set of policy vectors that can occur in equilibrium. These are characterized by solving (6):

$$\phi_i^*(n) \in \left\{ \frac{1-\sqrt{1-4r_i(1-\Phi^*(n))}}{2}, \frac{1+\sqrt{1-4r_i(1-\Phi^*(n))}}{2} \right\}. \tag{7}$$

The set of possible policies can be narrowed further using the symmetry of these roots around 0.5. Since $r_1 \geq r_i$ for all i, A1 chooses the highest equilibrium policy. By (5), if $\phi_1^*(n) \geq 0.5$, then $\phi_i^*(n) < 0.5$ $(i \neq 1)$. Thus, in equilibrium $\phi_2^*(n), \ldots, \phi_n^*(n)$ must correspond to the "low" root in (7), and only A1 can choose a policy higher than 0.5.

Since ϕ_1^* is decreasing in n, these observations imply that $\Phi^*(n+1) > \Phi^*(n)$ if and only if $\phi_i(n+1) < \phi_i(n)$ for all Ai. Further, if $\phi_1^*(n+1) \geq 0.5$, then by (6) the $n+1th$ agent caused A2, …, An to *increase* their policy levels. This in turn implies that $\phi_1^*(n+1)(1-\phi_1^*(n+1)) \geq \phi_1^*(n)(1-\phi_1^*(n))$, and hence $\Phi^*(n+1) < \Phi^*(n)$. Likewise, if $\phi_1^*(n) \leq 0.5$, then adding an agent results in all agents reducing their policies, and again by (6), $\Phi^*(n+1) > \Phi^*(n)$.

Combining these derivations with a few auxiliary results yields the following proposition, which is a generalization of sorts of Comment 2. $\Phi^*(n)$ is decreasing in n until A1 chooses a policy at or below 0.5, after which it is increasing. Consequently, the probability that $x = 1$ is maximized either at $n = 1$ or $n = N$.

Proposition 2. Non-Monotonicity of System Effectiveness. $\Phi^*(n) < \Phi^*(n-1)$ for $n < \tilde{n}$ and $\Phi^*(n+1) \geq \Phi^*(n)$ otherwise, where $\tilde{n} = \begin{cases} \min\{n|\phi_1^*(n) \leq 0.5\} & \text{if } \phi_1^*(N) \leq 0.5 \\ N+1 & \text{otherwise.} \end{cases}$ ■

This result implies that $\tilde{n} > 1$ if the conditions of Comment 2 are satisfied, and that $\tilde{n} = 1$ only if $\Phi^*(2) > \Phi^*(1)$.

The primary import of Proposition 2 is that collective action problems are especially harsh when one agent (necessarily A1) is "friendly," in the sense that $r_1 > 0.5$. To see why, it is helpful to think in terms of the dynamics as the $n+1$-*th* agent is added to an n-agent equilibrium. Since A1 chooses the highest policy, it stands to gain the most by reducing its effort in response to a new agent. Agents A2-An compensate for this reduction by increasing their effort. For $n < \tilde{n}$, these adjustments plus the contribution of An+1 do not offset the reduction in A1's effort. But beyond $\tilde{n}$, A1's policy is relatively cheap at the margin, causing it to shirk less. Consequently, the new agent's contribution more than compensates for the incumbent agents' reduced efforts.

In many cases, $\tilde{n}$ is easy to calculate. For example, $\tilde{n} = 1$ if and only if $r_1 \leq 0.5$ (and thus $\phi_1^*(n) \leq 0.5$ for all n). Otherwise, $\tilde{n} > 1$ (*e.g.*, $\tilde{n} = 3$ in Table 2). Under these conditions, the above derivations provide a way to characterize a simple lower bound on $\tilde{n}$. Recalling that $\phi_1^*(n) = r_1 \prod_{j \neq 1}(1-\phi_j^*(n))$, we see that $\phi_1^*(n) \leq 0.5$ if $\prod_{j\neq 1}(1-\phi_j^*(n)) \leq \frac{1}{2r_1}$. Using expression (7), this condition is equivalent to:

$$\prod_{j=2}^{n}\left(1 - \frac{1-\sqrt{1-4r_j(1-\Phi^*(n))}}{2}\right) \leq \frac{1}{2r_1}. \tag{8}$$

Since $\Phi^*(n) \geq 1-\frac{1}{4r_1}$ and the left-hand side of (8) is decreasing in n, (8) implies:

$$\tilde{n} \geq \min\left\{n \;\middle|\; \frac{r_1}{2^{n-2}}\prod_{j=2}^{n}\left(1+\sqrt{1-r_j/r_1}\right) \leq 1\right\}.$$

The Optimal Number of Agents. Generally, P's optimal choice of the number of agents solves the following problem:

$$n^* = \max_n\{\Phi^*(n) - nk\}.$$

Because $\Phi^*(n)$ may be non-monotonic, n^* depends on $\tilde{n}$. In the simplest case, if $\tilde{n} = 1$, then $\Phi^*(n)$ is increasing n. The relation between $\Phi^*(n)$ and n is roughly "concave" (especially for large n), so P approximately balances marginal benefit with marginal cost. If $\tilde{n} > 1$, then a non-redundant system becomes more attractive because the marginal benefit of the first few agents after A1 is negative. Proposition 3 uses these facts to establish the main result of the direct redundancy model, which relates n^*, r_i, and k.

Proposition 3. Optimal Redundancy.

(i) n^* is non-increasing in k.

(ii) If $r_1 > \frac{1}{2}$, and $\sum_{i=2}^{N} r_i < 2 - \frac{1}{r_1}$ or $\prod_{i=2}^{N}(1 - r_i) > \frac{1}{r_1} - 1$, then $n^* = 1$.

(iii) If $r_1 \in (0, \frac{1}{2}]$ and $r_2 > 0$, then $n^* > 1$ for k sufficiently low. ■

Part (i) of the result generalizes Proposition 1 in a fairly obvious way: extra agents are less appealing as their costs increase. But part (ii) shows that if A1 is "friendly" and the other agents are collectively unfriendly, then $n^* = 1$ *regardless* of k. This is a direct consequence of a high value of $\tilde{n}$, as no number of new agents can compensate for the losses imposed by agents prior to $\tilde{n}$. Finally, part (iii) provides some conditions under which low costs do matter—*i.e.*, when agents are unfriendly and $\tilde{n} = 1$.

These results contrast usefully with those of the non-game theoretic formulation of redundancy. In the latter, new agents *always* raise the probability of a good outcome and redundancy is strictly decreasing in k. But in the game studied here, this is only true when there is no agent inclined to choose a high policy or effort level. Otherwise, redundant agents may contribute little if anything to the probability of success. Thus, despite the greater difficulty in deriving closed form solutions when $n > 2$, the primary intuition about the desirability of redundant systems remains the same.

5. Latent Redundancy

In the previous section the principal had no leverage over agents other than the (limited) option to give them common jurisdiction over the task. However, in many cases her leverage is likely to be considerably greater. Typically principals are in a position to assess performance over time, and may also be able to terminate the jurisdictions of agencies or wayward bureaucrats.[16] The latent redundancy model shows how this feature may affect bureaucratic effectiveness with a repeated game in which P faces a single, replaceable agent in each period. The results are suggestive in nature, as a more complete analysis would give P the ability to add or terminate multiple agents in each period.[17]

As a baseline for comparison, consider the following repeated variant of the direct redundancy game. P chooses a set of irreplaceable agents to open the game. In each period, P pays the cost k for each agent, the agents choose policy, and Nature reveals outcomes.

[16] One prominent example is the 1905 Transfer Act, analyzed by Carpenter (2001). The law transferred jurisdiction of 46 million acres of forest reserves from the General Land Office to the U.S. Forest Service.

[17] For more dedicated analyses of agency termination, see Kaufman (1976), Carpenter (2000), Lewis (2002), and Carpenter and Lewis (2002). For a theoretical and empirical analysis of political appointees see, *e.g.*, Chang, Lewis, and McCarty (2001).

Because P cannot replace agents, it is easy to see that—regardless of whether the game is finitely or infinitely repeated—she would choose n^*, and in each period Ai would choose $\phi_i^*(n^*)$. Thus the equilibrium for this repeated game is identical to that of the single-period direct redundancy game.

5.1. The Repeated Game

The latent redundancy game is infinitely repeated, with future payoffs discounted by a common factor $\delta \in (0, 1)$. Each period is identical to the $n = 1$ subgame of the direct redundancy game, except in the following respect: at the beginning of each period, P may terminate the incumbent agent and replace it with a new one of her choosing. So as not to "rig" the results excessively in favor of superior agent performance, I assume that agents face no costs from being terminated.[18] As before, each agent costs k in each period. Once terminated, an agent cannot return to the game. P may therefore potentially induce better policy performance by making her future choice of agents contingent upon observed performance.[19]

To maintain comparability with the direct redundancy game, I assume that the set of replacement agents in each period is time-invariant, with preferences identical to agents $\{A1, \ldots, AN\}$ in the direct redundancy game. This reflects the possibility that P has relatively little control over the formation of agent preferences; for example, a newly promoted division chief might have the same professional background as her predecessor.[20]

There are many Nash equilibria of this game. For this reason, I focus on a general class of intuitive, sequentially rational equilibria. The class is defined by two rules. First, there is a *termination rule*, defined by the set $\{M_i\}$, such that P tolerates M_i failures from the incumbent agent with preferences identical to Ai before terminating it and choosing a new one. Second, the *replacement rule* is a sequence $\{\rho_\tau\}_{\tau=1}^{\infty}$ ($\rho_\tau \in \{Ai\}\ \forall \tau$) specifying a deterministic order in which new agents are selected. These two simple rules thereby capture a wide range in both the requirements for termination (including the ability to discriminate across agent types), as well as the choice of replacements. If, out of equilibrium, P does not terminate the agent or chooses the wrong agent, the agent simply chooses its one-shot best

[18]In an employment context, this is equivalent to assuming that a bureaucrat's job prospects outside the agency are as good as those within the agency.

[19]The game therefore resembles a simplified version of the Banks and Sundaram (1998) agent retention model. In their model, principals face adverse selection as well as moral hazard problems, but agents may "live" for only two periods.

[20]While this is perhaps the natural assumption, another reasonable possibility is that new agents' preferences are randomly generated.

policy (*i.e.*, r_i) in each period for as long as it is not replaced.

The existence of such equilibria is easy to demonstrate. As an example, consider the termination rule $\{M_i = 0\}\ \forall i$ and replacement rule $\{\rho_\tau = \text{A1}\}\ \forall \tau$. P will then terminate an agent of type A1 each period. Since that agent has no control over its future, it simply chooses its one-shot best response. Note that P never chooses A2 because its one-shot best response is not as good as A1's, although she is indifferent in equilibrium between terminating an agent and continuing with it.

Of particular interest in this section is the equilibrium with optimal agent termination and replacement strategies for P. Aside from being a natural focal point, this equilibrium accords with the role that political principals play in institutional design. Since principals are to some degree responsible for establishing "rules" for subordinates, it is reasonable to conjecture that if any player *could* coordinate play on a particular equilibrium, it would be P. The following comment characterizes this optimal equilibrium.

Comment 5. The optimal termination rule is $\{M_i = 1\}\ \forall i$, and $\{\rho_\tau = \text{A1}\}\ \forall \tau$. ■

The best equilibrium for P is thus the one in which type-A1 agents, whose preferences most closely aligned with P, are always chosen. Moreover, that agent is subjected to a harsh termination rule, where a single failure results in its replacement. The result is intuitive because A1-type agents are the most willing to choose "high" policies, and harsh termination rules maximize their incentives to reduce the risk of failure. Crucial to the result is P's ability to draw a fresh A1-type agent costlessly after each termination. Changing this assumption would cause P to prefer more forgiving termination rules.

5.2. Equilibrium Policies

Given Comment 5, the characterization of optimal policies is straightforward. Let ϕ_i^{j*} represent Ai's optimal policy choice given that it has j failures before termination.

Proposition 4. Policy Under Latent Redundancy. In all periods, equilibrium policy is:

$$\phi_1^{1*} = \begin{cases} 0 & \text{if } r_1 \leq 0 \\ \min\{\frac{1-\sqrt{1-2\delta r_1}}{\delta}, 1\} & \text{otherwise.} \end{cases} \quad ■$$

Clearly, $\phi_1^{1*} \geq \phi_1^*(1)$, with the inequality strict when $r_1 < 1$. Thus, the ability to terminate agents raises policy and thereby limits the extent to which additional agents will be desired. This result goes much further, however. For δ sufficiently high, the threshold for the agent to choose a policy of 1—in other words, to do exactly as P wants—is quite low. In the direct redundancy game, $\phi_1^*(1) = 1$ for $r_1 \geq 1$, but in the latent redundancy game, $\phi_1^{1*} = 1$ if $r_1 \geq 1 - \frac{\delta}{2}$. Alternatively, if $r_1 \geq \frac{1}{2}$, then for any $\delta \geq 2(1-r_1)$, P can attain her ideal outcome with only a single agent.

6. Extensions

6.1. $s \times n$ Systems

In the previous sections, the assumed policy technology required only one success for a good outcome to result. In many applications, however, more than a single success is required: airliners may require more than one working engine to fly, and drug interdiction requires the effective collaboration of multiple agencies. To use the example of Section 2, one can imagine that because of resource or jurisdictional constraints, no agency can uncover all of the required information unilaterally, but that a combination of two effective agencies can. This type of system is easily incorporated into the framework developed in Section 4.

Let the minimum number of successes required to achieve $x = 1$ be $s \geq 1$. Such systems are often referred to as "$s \times n$" or "s-by-n" systems. Clearly, the number of agents n ($n \leq N$) must be at least s for $x = 1$ to result. In the simplest case, where the reliability r is constant across agents, the probability of a good outcome in a non-strategic environment is:

$$\Pr\{x = 1\} = \sum_{i=s}^{n} \binom{n}{i} r^i (1-r)^{n-i}.$$

As in the direct redundancy game (where $s = 1$), raising the number of agents strictly increases the probability that $x = 1$.

In the strategic environment studied here, reliability rates will continue to vary across agents, and the probability of each outcome will depend on agents' policy choices in more complex ways than previously. As a result, the probability of success is less straightforward to calculate than in the previous sections. To begin, it will be convenient to define the probability that exactly q agents other than Ai succeed as follows:

$$\mu_{-i}^q(\phi) = \sum_{j_1 \neq i} \sum_{j_2 \neq i, j_1} \cdots \sum_{j_q \neq i, j_1, \ldots, j_{q-1}} \left[\prod_{k=j_1}^{j_q} \phi_k \prod_{l \neq i, j_1, \ldots, j_q} (1 - \phi_l) \right]. \tag{9}$$

Agent Ai's objective function then generalizes from (2) as:

$$v_i(\phi) = w_i \left[\phi_i \mu_{-i}^{s-1}(\phi) + \sum_{j=s}^{n} \mu_{-i}^{j}(\phi) \right] - c_i(\phi_i).$$

Policy Choice. Clearly, $\mu_{-i}^{q}(\phi)$ does not depend on ϕ_i, and as a result Ai can only affect the outcome if its effort is pivotal, which occurs when exactly $s-1$ agents succeed. Ai's best response thus generalizes from (5) as follows:

$$\phi_i^{s*}(n) = \begin{cases} 0 & \text{if } w_i \leq 0 \\ \min\{\frac{w_i \mu_{-i}^{s-1}(\phi)}{2m_i}, 1\} & \text{if } w_i > 0. \end{cases} \quad (10)$$

The Optimal Number of Agents. While a closed-form solution for the optimal number of agents, n^{s*}, is clearly more difficult to derive than in Section 4, the following partial characterization is easily demonstrated.

Proposition 5. Necessary Redundancy. If $r_i < 1$ for all Ai ($i \leq N$), then $n^{s*} = 0$ or $n^{s*} > s$. ■

This result essentially reverses the central intuition of the direct redundancy model. Just as agents create negative externalities (by the collective action problem) when $s = 1$, agents can create positive externalities when $s > 1$. When $n = s$, all externalities are positive, because any effort is wasted unless all other agents also contribute. Thus, best responses are increasing (and linear) in other agents' efforts. This is easily seen in equations (9) and (10), as the last product in (9) is empty for $n = s$.

The proof uses these facts to show that if no agent would ever be willing to choose a policy of 1 (*i.e.*, $r_i < 1$), then all agents must choose 0. At an equilibrium, P therefore chooses either $n^{s*} = 0$ or, if k is sufficiently low, some n^{s*} strictly greater than the minimum technically necessary. An easily proved corollary of this result is that if all agents are perfectly reliable ($r_i = 1$), then $\Pr\{x = 1\} = 1$ when $n = s$, and thus $n^{s*} = s$ if $k < \frac{1}{s}$, and $n^{s*} = 0$ otherwise.

Table 3 below gives an example of the interplay of positive and negative externalities in some $2 \times n$ systems. Since $r_i < 1$ for all Ai, agents choose zero policies when $n = 2$. When n increases to 3 and then to 4, agents individually choose higher policies because of the positive externalities induced by the presence of new agents. When A5 is introduced, however, the number of agents becomes large relative to s, and the probability that only a single agent succeeds ($\mu_{-i}^{1}(\phi)$) drops, while the probability that each agent's effort will be unnecessary rises. As in the direct redundancy game, this creates negative externalities. Thus, the individual efforts of A1-A4 are lowered, but the collective reliability of the system still increases.

Table 3: $2 \times n$ **Systems** $w_1 = 1.6$, $w_2 = 1.4$, $w_3 = 1.2$, $w_4 = 1.0$, $w_5 = 0.8$; $m_i = 1$ $\forall$Ai		
n	$\phi^{s*}(n)$	$\Phi^{s*}(n)$
2	(0, 0)	0
3	(0.315, 0.286, 0.253)	0.197
4	(0.345, 0.308, 0.266, 0.222)	0.322
5	(0.340, 0.300, 0.257, 0.213, 0.169)	0.376

6.2. Policy Externalities

Up to this point it has been assumed that the policies chosen by agents did not impose any *technical* interdependencies or externalities on the effectiveness of other agents' actions. Externalities were exclusively generated endogenously, by the agents' strategic incentives. As normal accident theory points out, however, technical externalities may be common in more complex task environments. Here I explore some implications of these externalities for redundant bureaucratic structures.

For simplicity, I return to the case examined in Section 4.1, where $N = 2$ and the assumption that $\frac{w_1}{m_1} \geq \frac{w_2}{m_2}$ is relaxed (*i.e.*, $r_1 < r_2$ is allowed). Outcomes are determined in the same way as in the direct redundancy model, but with the generalization that the effectiveness of each agent may now depend on the effort of the other, as follows. The vector of success probabilities is given by $\pi = \phi \cdot \mathbf{\Psi}$, where $\mathbf{\Psi}$ is an $n \times n$ matrix with elements $\psi_{ji} \in \Re$. Ai's probability of succeeding is thus π_i (as opposed to ϕ_i), where:

$$\pi_i = \begin{cases} 1 & \text{if } \sum_j \phi_j \psi_{ji} \geq 1 \\ 0 & \text{if } \sum_j \phi_j \psi_{ji} \leq 0 \\ \sum_j \phi_j \psi_{ji} & \text{otherwise.} \end{cases}$$

Accordingly, $\Pr\{x = 1\} = 1 - \prod_{j=1}^{n}(1 - \pi_j)$. P's expected utility is thus given by:

$$u(\phi) = 1 - \prod_{j=1}^{n}(1 - \pi_j) - nk,$$

and Ai's expected utility by:

$$v_i(\phi) = w_i \left[1 - \prod_{j=1}^{n}(1 - \pi_j)\right] - c_i(\phi_i).$$

I refer to $\mathbf{\Psi}$ as the *weighting matrix*. This matrix operationalizes the externalities that agents' activities may impose on one another, thus causing success probabilities to be non-independent. If $\mathbf{\Psi}$ is the identity matrix, $\mathbf{I}$, then $\pi = \phi$. Agents' success probabilities are then independent, and the model is identical to that of Section 4.1. But in contrast with equation (3), a non-trivial weighting matrix may force an agent to consider the consequences of other agents' policy choices on its outcomes, and to consider the impact of its choice on other agents' outcomes. For example, if $w_2 > 0$ and $\psi_{21} > 0$, then A2 has an incentive to *increase* its policy choice compared to a situation in which $\psi_{21} = 0$. This happens because A2's effort will help A1 to succeed, thus increasing the marginal value of A2's effort. Similarly, if $\psi_{ji} < 0$, then any $\phi_j > 0$ reduces the probability that Ai succeeds.

Policy Choice. Solving backwards, if $n = 1$ then there are no externalities and policy is as determined in Section 4.1. If $n = 2$, then by the concavity of $v_i(\cdot)$, the following first-order condition is sufficient to characterize Ai's equilibrium policy at an interior solution:

$$w_i[\psi_{i1} + \psi_{i2} - (\psi_{12}\psi_{21}+\psi_{11}\psi_{22})\phi_{-i} - 2\psi_{i1}\psi_{i2}\phi_i] - 2m_i\phi_i = 0.$$

Simplifying, Ai's best response is given by:

$$\phi_i^*(2) = \frac{w_i[\psi_{i1} + \psi_{i2} - (\psi_{12}\psi_{21}+\psi_{11}\psi_{22})\phi_{-i}]}{2(m_i+w_i\psi_{i1}\psi_{i2})}.$$

Even with externalities, each agent's best response is linear and decreasing in the other's policy choice. Solving the system, Ai's unique interior equilibrium policy is:

$$\phi_i^*(2) = \frac{w_i[2(m_{-i}+w_{-i}\psi_{-i1}\psi_{-i2})(\psi_{i1}+\psi_{i2}) - w_{-i}(\psi_{12}\psi_{21}+\psi_{11}\psi_{22})(\psi_{-i1}+\psi_{-i2})]}{4(m_{-i}+w_{-i}\psi_{-i1}\psi_{-i2})(m_i+w_i\psi_{i1}\psi_{i2}) - w_i w_{-i}(\psi_{12}\psi_{21}+\psi_{11}\psi_{22})^2}.$$

Note that the conditions under which a corner solution occurs are no longer trivial. Even if $w_i < 0$, it is possible that $\phi_i^*(2) > 0$ because Ai may wish to impose a negative externality on Aj's ($j \neq i$) effort.

The effects of the weighting matrix on the comparative statics of the $n = 2$ subgame are generally straightforward. If $\mathbf{\Psi}$ is symmetric and agents have identical preferences and costs, then equilibrium policies are identical. Generally, for weighting matrices "close" to $\mathbf{I}$, the comparative statics resemble those of Section 4.1. As $\mathbf{\Psi}$ diverges from $\mathbf{I}$, however, agent incentives become increasingly distorted. Consider the example in Table 4, where all cost and utility parameters are as in Table 1. Since $\psi_{11} = 0.6$ and $\psi_{21} = 0.4$, A2 has almost as much control over A1's output as A1 itself. Thus A2 chooses a much higher policy than in Table 1, while A1 chooses a lower policy. But since A2 faces higher costs, its inclusion reduces the efficiency of equilibrium policies relative to the single agent case.

Table 4: Policy Externalities

$w_1 = 1.1,\ m_1 = 1,\ m_2 = 2;\ \mathbf{\Psi} = \begin{pmatrix} 0.6 & 0 \\ 0.4 & 1 \end{pmatrix}$

n	Case	ϕ^*	π	Total Cost	$\Phi^*(n)$
1	(a)-(c)	(0.55)	(0.55)	0.3025	0.55
2	(a) $w_2 = 3$	(0.125, 0.621)	(0.323, 0.621)	0.7871	0.744
	(b) $w_2 = 2$	(0.178, 0.462)	(0.291, 0.462)	0.4583	0.619
	(c) $w_2 = 1$	(0.244, 0.261)	(0.251, 0.261)	0.1959	0.446

The Optimal Number of Agents. P's decision is essentially similar to that in Section 4.1, in that Proposition 1 holds. However, when $\mathbf{\Psi} \neq \mathbf{I}$, the marginal value of the second agent, $\Phi^*(2) - \Phi^*(1)$, is different. In case (a) of Table 4, P chooses to have two agents if $k < 0.194$, as opposed to $k < 0.124$ in case (a) of Table 1. Adding A2 increases $\Phi^*(n)$ more in Table 4 because $\psi_{21} > 0$ and $\psi_{22} = 1$, thus raising the return to effort of an agent that already desires a high policy. A1's effort is correspondingly reduced, but not by enough to reduce $\Phi^*(n)$ overall. A similar analysis holds for case (b) in both tables. By contrast, in case (c), A1 is the high demander of policy, and its incentive to produce is inhibited by the low values of ψ_{11} and ψ_{12}. Consequently, P does even worse with two agents here than it did without externalities.

7. Discussion

7.1. Empirical Implications

While the models of bureaucratic politics in this paper are quite simple, they predict the assignment of agents to tasks as a function of some readily measurable variables. In particular, two tests of the relationship between Congress and federal bureaucracies would complement the many cases examined thus far in the literature. First, the direct redundancy model predicts that redundant structures will be chosen more often as agency preferences coincide less with Congress'.[21] Measures of agency preferences can be constructed in two

[21] One caveat is that the model is more naturally suited to "valence" than "spatial" issue spaces, though it is possible in some cases to consider r_i as an agent's spatial preference over $[0, 1]$.

ways. In policy areas where political appointees play an important role and preferences have a "spatial" component, data on the preferences of presidents or enacting coalitions (for agencies that are relatively "insulated," such as independent commissions) may be used. In other areas, survey data of career officials will be more appropriate.

Second, the latent redundancy model additionally predicts that redundant structures will be more prevalent as the ability of Congress (or the President) to affect personnel decreases. The substitutability of personnel within an agency could be measured by the extent to which career or political appointees staff the organization's leadership. Alternately, it could be measured by the level of specialization or expertise required for its personnel, which may be crudely estimated from the distribution of civil service ranks within the agency.

For both types of tests, the onset of civil service reform suggests some interesting possibilities. In the U.S., the 1883 Pendleton Act initiated an extensive transition from political to career appointees in executive agencies. The law shifted both bureaucrats' preferences and the ability of principals to replace them. Both the direct and latent redundancy models make predictions about the re-allocation of tasks that should result.

These tests can also be performed at the intra-agency level. For example, agencies like the FDA frequently consult a variable number of advisory boards before making decisions. Managers in law enforcement agencies routinely confront choices over the number of divisions that will be delegated a task. These and other questions about the design of reliable organizations can be subjected to systematic empirical scrutiny with this model or its extensions as a basis.

7.2. Conclusions

Theories in the wake of Landau's contribution to bureaucratic design and performance have steadily formalized and expanded upon the original model. The model presented here makes two additional contributions. First, consistent with more recent theoretical work in bureaucratic politics, it explicitly casts agents as strategic actors with preferences over policy outcomes. The primary insight for organizational design is that strategic interdependencies can play an important role, even in an environment with no technical interdependencies. Second, it examines a set of common policy environments that differ from the standard setting covered in Sections 3 and 4. Among these are $s \times n$ systems and settings where agents face the risk of termination. The resulting models therefore capture a variety of important features shaping the interaction of agents with political principals.

The incentives posed by the strategic environment suggest some serious limits on the amount of redundancy principals desire, when compared to its non-strategic counterpart. Redundant structures tend to help most when the set of agents available for a task are relatively "unfriendly," or disinclined to choose policies that P would like. Here the collective action problems are not serious enough to hurt aggregate policy production, which is increasing in n. But if P has access to a friendly agent, then adding agents will tend not to help performance. As Proposition 2 establishes, policy production will first decrease, and then increase, as n rises. This effect greatly reduces the average value of new agents, and moreover may make a single agent optimal in environments where n is constrained to be small. Finally, if the principal can terminate agents for poor outcomes, she can achieve her *ideal* outcome with only moderately friendly agents if the discount factor is sufficiently high. Simply stated, principals can make agents compete against themselves with relative ease.

It is useful to reconsider the link between redundancy and normal accident theories in light of these results. By developing a theory of strategic interdependencies among system components, the models developed here begin to bridge the two. As the example of Section 2 first illustrated, the basic intuitions of both theories may be correct, depending on the parameter values assumed. But despite the constant presence of interactive complexity—arising from either strategic interaction or policy externalities—redundancy is not always rendered undesirable. In fact, redundancy becomes necessary when the agents' interaction creates positive externalities, as in $s \times n$ systems.

While the models developed here were intended to capture the essential features of redundancy in a parsimonious manner, they can be usefully extended in two directions. First, the theory should be generalized to a broader range of policy problems by encompassing different policy technologies. Two examples suggested by the literature are "three-state" outcomes that incorporate Type I and Type II errors, and the problem of agent sabotage. Other extensions might examine policy technologies used in the economic literature on collective action.[22]

Second, the theory can potentially also speak to other aspects of bureaucratic politics. One example is the issue of agent incentives raised by the latent redundancy model. In many contexts, principals might assume a more active role in structuring these incentives.

[22]For example, Hirshleifer's (1983) "best shot" technology, whereby the highest of the agents' effort levels determines the collective outcome.

Within an agency, this might entail the design of employment contracts for bureaucrats. In some contexts, principals may also encounter common agency problems (*e.g.*, Dixit 1995, Gailmard 2002b). Another example is the delegation of authority to agents, which is greatly simplified by the present models. The decision to delegate may depend on the agents' private information, thus shifting the problem from moral hazard to adverse selection. Finally, the occasional charges of wasteful redundancy suggest a range of questions about budgets and efficiency. Instead of a fixed cost for each agent, the budget may be modeled as a choice that limits each agent's feasible policy set (*e.g.* Silver 1996, Ting 2001). In equilibrium, the budget would be tied to the principal's anticipation of the agent's costs. The principal would then bear the marginal costs of policy and also care explicitly about policy efficiency.

The models developed here therefore move the foundations of redundancy theory from reliability engineering to game theory. In so doing, they link redundancy with modern theories of bureaucratic politics and collective action, and furthermore establish a framework for considering a variety of new issues in organizational design.

APPENDIX

Proof of Comment 1. Since $v_i(\cdot)$ is concave, either the first-order condition $w_i(1-\phi_{-i})-2m_i\phi_i = 0$ (3) uniquely characterizes the solution, or the optimal policy is 0 (1) if ϕ_i satisfying (3) is less than (greater than) 0 (1). Note as a result that if $w_i \leq 0$, Ai's dominant strategy is to choose $\phi_i^*(2) = 0$. There are nine cases. (i) If $w_1 \leq 0$ and $w_2 \leq 0$, then $\phi_1^*(2) = 0$ and $\phi_2^*(2) = 0$. (ii) If $w_1 \leq 0$ and $w_2 \in (0, 2m_1)$, then $\phi_1^*(2) = 0$ and A2's best response is interior: $\phi_2^*(2) = \frac{w_2}{2m_2}$. (iii) If $w_1 \leq 0$ and $w_2 \geq 2m_2$, then $\phi_1^*(2) = 0$ and A2's best response is on the corner: $\phi_2^*(2) = 1$. (iv) If $w_1 \in (0, 2m_1)$ and $w_2 \leq 0$, then by symmetry with case (ii), $\phi_1^*(2) = \frac{w_1}{2m_1}$ and $\phi_2^*(2) = 0$. (v) If $w_1 \in (0, 2m_1)$ and $w_2 \in (0, 2m_2)$, then substituting A2's best response into A1's, we obtain $\phi_1^*(2) = \frac{w_1(2m_2 - w_2)}{4m_1m_2 - w_1w_2}$ and $\phi_2^*(2) = \frac{w_2(2m_1 - w_1)}{4m_1m_2 - w_1w_2}$. (vi) If $w_1 \in (0, 2m_1)$ and $w_2 \geq 2m_2$, then an interior solution is impossible: $\phi_1^*(2) = 0$ and $\phi_2^*(2) = 1$. (vii) If $w_1 \geq 2m_1$ and $w_2 \leq 0$, then by symmetry with case (iii), $\phi_1^*(2) = 1$ and $\phi_2^*(2) = 0$. (viii) If $w_1 \geq 2m_1$ and $w_2 \in (0, 2m_2)$, then by symmetry with case (vi), $\phi_1^*(2) = 1$ and $\phi_2^*(2) = 0$. (ix) If $w_1 \geq 2m_1$ and $w_2 \geq 2m_2$, then two corner solutions are possible, where $\phi_i = 0$ and $\phi_{-i} = 1$. I choose the solution $\phi_1^*(2) = 1$ and $\phi_2^*(2) = 0$. ■

Proof of Comment 2. There are four trivial cases. First, if $w_1 \leq 0$, then $\phi_1^*(1) = 0$ and $\Phi^*(2) \geq \Phi^*(1)$. Second, if $w_1 \geq 2m_1$, then $\Phi^*(n) = 1$ for any n. Third, if $w_1 \in (0, 2m_1)$ and $w_2 \leq 0$, $\Phi^*(n) = \frac{w_1}{2m_1}$ for any n. Fourth, if $w_1 \in (0, 2m_1)$ and $w_2 \geq 2m_2$, then $\phi_2^*(2) = 1$ and $\Phi^*(2) \geq \Phi^*(1)$.

I therefore focus on the fifth case, where $w_1 \in (0, 2m_1)$ and $w_2 \in (0, 2m_2)$. Here, $\Phi^*(1) = \phi_1^*(1) = \frac{w_1}{2m_1}$ and $\Phi^*(2) = 1-(1-\phi_1^*(2))(1-\phi_2^*(2))$. By Comment 1 the equilibrium policy is interior and thus $\Phi^*(2) < \Phi^*(1)$ iff:

$$\frac{w_1(2m_2 - w_2)}{4m_1m_2 - w_1w_2} + \frac{w_2(2m_1 - w_1)}{4m_1m_2 - w_1w_2} - \frac{w_1w_2(2m_1 - w_1)(2m_2 - w_2)}{(4m_1m_2 - w_1w_2)^2} < \frac{w_1}{2m_1}.$$

This expression simplifies to: $4m_1m_2(2m_1 - 3w_1) + w_1^2(w_2 + 4m_2 - \frac{w_1w_2}{2m_1}) < 0$. Clearly, a necessary condition for this to hold is: $\epsilon = 4m_2w_1^2 - 12m_2m_1w_1 + 8m_2m_1^2 < 0$, which obtains iff $w_1 \in (m_1, 2m_1)$. If in addition $w_1^2(w_2 - \frac{w_1w_2}{2m_1}) < -\epsilon$, then $\Phi^*(2) < \Phi^*(1)$. Let $\hat{w}_2 = \frac{-\epsilon}{w_1^2(1-\frac{w_1}{2m_1})} = \frac{8m_1m_2(w_1-m_1)}{w_1^2}$, and note that $\hat{w}_2 > 0$ for $w_1 \in (m_1, 2m_1)$. Thus a necessary and sufficient condition for $\Phi^*(2) < \Phi^*(1)$ is $w_1 \in (m_1, 2m_1)$ and $w_2 \in (0, \frac{8m_1m_2(w_1-m_1)}{w_1^2})$, or equivalently, $r_1 \in (\frac{1}{2}, 1)$ and $r_2 \in (0, \frac{4m_1(w_1-m_1)}{w_1^2})$. ■

Proof of Comment 3. The conditions of the comment imply that $w_1 \in (0, m_1)$ and $w_2 \in [\frac{w_1(2m_1 - w_1)}{2(m_1 - w_1)}, 2m_2)$. Thus, equilibrium policies are interior. Thus one agent is more efficient than two if: $\frac{\phi_1^*(1)}{c_1(\phi_1^*(1))} > \frac{\phi_1^*(2) + (1 - \phi_1^*(2))\phi_2^*(2)}{c_1(\phi_1^*(2)) + c_2(\phi_2^*(2))}$.

Substituting the appropriate expressions from Comment 1 and simplifying, this condition can be re-stated as: $8m_1m_2(m_1 - m_1\frac{w_2}{w_1} + w_2 - \frac{w_1}{2}) + w_1w_2(w_1 - 2m_1 - 2m_2) < 0$. Since $m_i > 0$ and $w_1 < 2m_1$, a sufficient condition for this to hold is: $m_1 - m_1\frac{w_2}{w_1} + w_2 - \frac{w_1}{2} \leq 0$. Since $m_1 > w_1$ by assumption, this obtains if $w_2 \geq \frac{w_1(2m_1 - w_1)}{2(m_1 - w_1)}$. ■

Proof of Proposition 1. Denote by k_{ij} ($2 \geq i > j \geq 0$) the cutpoint such that P prefers $n = i$ to $n = j$ iff $k < k_{ij}$. Thus, k_{21} is characterized by $\phi_1^*(2) + (1 - \phi_1^*(2))\phi_2^*(2) - 2k > \phi_1^*(1) - k$, which implies $k_{21} = \phi_1^*(2) + (1 - \phi_1^*(2))\phi_2^*(2) - \phi_1^*(1)$. Likewise, k_{10} is characterized by $\phi_1^*(1) - k > 0$, which implies $k_{10} = \phi_1^*(1)$, and k_{20} is characterized by $\phi_1^*(2) + (1 - \phi_1^*(2))\phi_2^*(2) - 2k > 0$, which implies $k_{20} = \frac{\phi_1^*(2) + (1 - \phi_1^*(2))\phi_2^*(2)}{2}$. Note that $k_{ij} \in [0, 1]$ for all i, j.

Of the six weak orderings of $\{k_{ij}\}$, only two are feasible: (i) $k_{21} \leq k_{20} \leq k_{10}$, and (ii) $k_{10} \leq k_{20} \leq k_{21}$ (it is easily verified that the others imply a non-transitive ordering over n). These imply three possible values for n^*. First, in both cases, $n^* = 0$ if $k \geq \max\{k_{10}, k_{20}\}$. Second, in case (i), $n^* = 1$ if $k \in [k_{21}, k_{10})$. In case (ii), $n = 1$ cannot be optimal. Finally, in both cases, $n^* = 2$ if $k < \min\{k_{21}, k_{20}\}$. ■

Proof of Comment 4. (i) The unique equilibrium for $n \leq 2$ is characterized in Section 4.1. Additionally, using the equilibrium selection rule assumed in the text, if $r_i = 1$ for any Ai, then the unique equilibrium is $\phi_1^*(n) = 1$ and $\phi_j^*(n) = 0$ for all $j \neq 1$. Further, all agents Aj for which $w_j \leq 0$ have a dominant strategy of $\phi_j^*(n) = 0$. I therefore consider only the case in which $n \geq 3$ and $w_i \in (0, 2m_i)$ for all Ai (*i.e.*, interior solutions obtain for all agents). For notational convenience, I drop references to n when discussing $\phi_i^*(n)$ and $\Phi^*(n)$.

I first derive a necessary condition (NC) for an equilibrium. Manipulating (5), $\phi_i^*(1 - \phi_i^*) = r_i(1 - \Phi^*)$. As in (7), let $\bar{\phi}_i(\Phi)$ and $\underline{\phi}_i(\Phi)$ denote the upper and lower roots of this equation for arbitrary $\Phi \in [1 - \frac{m_1}{2w_1}, 1]$ (note that (5) implies $\Phi^* \geq 1 - \frac{m_1}{2w_1}$). Clearly, $\bar{\phi}_i(\cdot)$ is increasing and $\underline{\phi}_i(\cdot)$ is decreasing. So, given success probability Φ, Ai's best response must be either $\bar{\phi}_i(\Phi)$ or $\underline{\phi}_i(\Phi)$. Note that $\bar{\phi}_i(\Phi)$ cannot be an equilibrium choice for A2, ..., An, since this would imply $\phi_i^* > \phi_1^*$, generating an obvious contradiction. Thus any Φ implies the following policy choices: $\bar{\phi}_1(\Phi)$ or $\underline{\phi}_1(\Phi)$ for A1, and $\underline{\phi}_i(\Phi)$ for A2, ..., An. Now let $\gamma^n(\Phi) = 1 - \prod_{i=2}^{n}(1 - \underline{\phi}_i(\Phi))$ be the value of $\Pr\{x = 1\}$ implied by $\{\underline{\phi}_i(\Phi)\}_{i=2}^{n}$. A1's best response to $\gamma^n(\Phi)$ is $\phi_1^{n*}(\Phi) = r_1(1-\gamma^n(\Phi))$. Since $\gamma^n(\cdot)$ is decreasing, $\phi_1^{n*}(\cdot)$ is increasing. NC is then: a policy vector is an equilibrium only if $\phi_1^{n*}(\Phi) \in \{\bar{\phi}_1(\Phi), \underline{\phi}_1(\Phi)\}$.

Note that $\bar{\phi}_i(\cdot)$ is increasing and $\underline{\phi}_i(\cdot)$ is decreasing, implying $\gamma^n(\cdot)$ is decreasing and $\phi_1^{n*}(\cdot)$ is increasing. All are continuous over $[1-\frac{1}{4r_1}, 1]$.

I claim that $\gamma^n(\cdot)$ satisfies the following 'single crossing' (SC) property: if $\gamma^n(\Phi) < \underline{\phi}_1(\Phi)$, then $\frac{d\gamma^n}{d\Phi} > \frac{d\underline{\phi}_1}{d\Phi}$, implying that if $\gamma^n(\hat{\Phi}) < \underline{\phi}_1(\hat{\Phi})$ for some $\hat{\Phi}$, then $\gamma^n(\Phi) < \underline{\phi}_1(\Phi)$ for all $\Phi < \hat{\Phi}$. To show this, suppose otherwise there exist Φ' and Φ'' such that $\Phi' < \Phi''$, $\gamma^n(\Phi'') < \underline{\phi}_1(\Phi'')$, and $\gamma^n(\Phi')-\gamma^n(\Phi'') > \underline{\phi}_1(\Phi')-\underline{\phi}_1(\Phi'')$. Let $y = \frac{\underline{\phi}_1(\Phi')}{\underline{\phi}_1(\Phi'')}$. It is easily verified that $y > 1$ and $y > \frac{\underline{\phi}_i(\Phi')}{\underline{\phi}_i(\Phi'')}$ for $i > 1$; thus $1-\prod_{i=2}^n(1-y\underline{\phi}_i(\Phi'')) > 1-\prod_{i=2}^n(1-\underline{\phi}_i(\Phi')) = \gamma^n(\Phi') > \gamma^n(\Phi'')$. Additionally it is straightforward (if tedious) to show that $\frac{1-\prod_{i=2}^n(1-y\underline{\phi}_i(\Phi''))}{\gamma^n(\Phi'')} < y$. Therefore $\frac{\gamma^n(\Phi')}{\gamma^n(\Phi'')} < y$ and $\gamma^n(\Phi')-\gamma^n(\Phi'') < \underline{\phi}_1(\Phi')-\underline{\phi}_1(\Phi'')$: contradiction.

There are two cases. First, if $\gamma^n(1-\frac{m_1}{2w_1}) \geq 1-\frac{m_1}{w_1}$, then $\phi_1^{n*}(1-\frac{m_1}{2w_1}) \leq \frac{1}{2}$. Since $\underline{\phi}_1(1-\frac{m_1}{2w_1}) = \frac{1}{2}$ and $\underline{\phi}_1(1) = 0$, there exists a unique $\tilde{\Phi}$ such that $\phi_1^{n*}(\tilde{\Phi}) = \underline{\phi}_1(\tilde{\Phi})$. Now consider $\bar{\phi}_1(\cdot)$. By SC, $\gamma^n(\Phi) > \underline{\phi}_1(\Phi)$ for all $\Phi > 1-\frac{m_1}{2w_1}$. Since $\bar{\phi}_1(\Phi) = 1-\underline{\phi}_1(\Phi)$ and $\phi_1^{n*}(\Phi) = r_1(1-\gamma^n(\Phi))$, SC implies $\phi_1^{n*}(\Phi) < \bar{\phi}_1(\Phi)$ for all Φ. Thus $\phi_1^{n*}(\Phi) \neq \bar{\phi}_1(\Phi)$ for all Φ. Second, if $\gamma^n(1-\frac{m_1}{2w_1}) < 1-\frac{m_1}{w_1}$, then $\phi_1^{n*}(1-\frac{m_1}{2w_1}) > \frac{1}{2}$. Clearly, $\phi_1^{n*}(\Phi) \neq \underline{\phi}_1(\Phi)$ for all Φ. Now consider $\bar{\phi}_1(\cdot)$. Since $\bar{\phi}_1(1-\frac{m_1}{2w_1}) = \frac{1}{2}$, $\phi_1^{n*}(1-\frac{m_1}{2w_1}) > \bar{\phi}_1(1-\frac{m_1}{2w_1})$. Further, $\bar{\phi}_1(1) = 1 > \phi_1^{n*}(1) = r_1$. By SC (and again using the facts that $\bar{\phi}_1(\Phi) = 1-\underline{\phi}_1(\Phi)$ and $\phi_1^{n*}(\Phi) = r_1(1-\gamma^n(\Phi))$), there exists a unique $\tilde{\Phi}$ such that $\phi_1^{n*}(\tilde{\Phi}) = \bar{\phi}_1(\tilde{\Phi})$.

Combining results, there exists a unique policy vector $\phi^* \in \{(\bar{\phi}_1(\tilde{\Phi}), \underline{\phi}_2(\tilde{\Phi}), \ldots, \underline{\phi}_n(\tilde{\Phi})), (\underline{\phi}_1(\tilde{\Phi}), \ldots, \underline{\phi}_n(\tilde{\Phi}))\}$ inducing success probability $\tilde{\Phi}$ that satisfies NC. Since each agent's best response to any agent strategies is single-valued and unique, no equilibrium in strictly mixed strategies exists. Thus, ϕ^* is the unique equilibrium strategy profile.

(ii) To show $\phi_1^*(n)$ is decreasing in n, consider any n'' and n', where $n'' > n'$. Clearly, $\gamma^{n''}(\Phi) > \gamma^{n'}(\Phi)$, and thus $\phi_1^{n''*}(\Phi) < \phi_1^{n'*}(\Phi)$ for all Φ. Let $\tilde{\phi}_1(\Phi, r_1) = \{\underline{\phi}_1(\Phi), \bar{\phi}_1(\Phi)\}$ represent the correspondence of possible values of $\phi_1^*(n)$ for each Φ, and let $f(\tilde{\phi}_1)$ represent its inverse with respect to Φ. Note that $f(\cdot)$ is a continuous function over $[0,1]$ and $f(0) = f(1) = 1$. By (i) the unique equilibrium is characterized by $\phi_1^{n*}(\Phi) \in \tilde{\phi}_1(\Phi, r_1)$. By the facts that $\phi_1^{n*}(\Phi)$ is increasing and $\phi_1^{n*}(1) \leq r_1$, f satisfies $f(\tilde{\phi}_1) \leq \tilde{\Phi}$ for $\tilde{\phi}_1 \geq \phi_1^*$, where $\phi_1^{n*}(\tilde{\Phi}) = \tilde{\phi}_1$. Then $\phi_1^{n''*}(\Phi) < \phi_1^{n'*}(\Phi)$ implies there exists no $\tilde{\phi}_1 \geq \phi_1^*(n')$ such that $\tilde{\phi}_1 = \phi_1^{n''*}(\Phi) \in \tilde{\phi}_1(\Phi, r_1)$. Therefore, $\phi_1^*(n') > \phi_1^*(n'')$.

To show $\phi_1^*(n)$ is increasing in r_1, let $\phi_1^{*''}$ and $\phi_1^{*'}$ denote A1's equilibrium policies under

r_1'' and r_1', respectively, where $r_1'' > r_1'$. Also, let $\phi_1^{n*}(\Phi, r_1)$ denote A1's best response to $\gamma^n(\Phi)$ under r_1. Clearly, $\phi_1^{n*}(\Phi, r_1'') > \phi_1^{n*}(\Phi, r_1')$ for all Φ. Let $\tilde{z}$ denote the value of ϕ_1 such that $\phi_1^{n*}(\Phi, r_1'') \in \tilde{\phi}_1(\Phi, r_1')$. By a symmetric argument with the first part of (ii), $z > \phi_1^{*'}$. Now note that f is increasing in r_1 over $\tilde{\phi}_1 \in (0,1)$. Since $\phi_1^{n*}(\Phi, r_1'')$ is increasing in Φ, this implies that there exists no $\tilde{\phi}_1 \leq z$ such that $\tilde{\phi}_1 = \phi_1^{n*}(\Phi, r_1'') \in \tilde{\phi}_1(\Phi, r_1'')$. Thus, $\phi_1^{*''} > z$, implying $\phi_1^{*''} > \phi_1^{*'}$.

(iii) The result holds trivially for $n = 1$; thus, assume $n > 1$. Suppose otherwise. Let $\Phi^{*''}$ and $\Phi^{*'}$ denote the equilibrium success probabilities and $\phi_i^{*''}$ and $\phi_i^{*'}$ Ai's equilibrium policies under r_1'' and r_1', respectively, where $r_1'' > r_1'$. By (8) and the fact that $\phi_i^*(n) = \underline{\phi}_i(\Phi^*)$ for $i > 1$, $\Phi^{*''} < \Phi^{*'}$ implies $\phi_i^{*''} > \phi_i^{*'}$ for $i > 1$. But by part (ii), $\phi_1^{*''} > \phi_1^{*'}$, so by (5) $\phi_i^{*''} < \phi_i^{*'}$ for some $i > 1$: contradiction. ■

Proof of Proposition 2. (This proof uses the notation developed in the proof of Comment 3.) By the proof of Comment 3, in equilibrium $\bar{\phi}_1(\Phi) = \phi_1^{n*}(\Phi) = r_1(1-\gamma^n(\Phi))$ or $\underline{\phi}_1(\Phi) = \phi_1^{n*}(\Phi)$. There are two cases. First, suppose that $\phi_1^{n'*}(1-\frac{1}{4r_1}) \leq \frac{1}{2}$ for some n' (equivalently, $\gamma^{n'}(1-\frac{1}{4r_1}) \geq 1-\frac{1}{2r_1}$). Then, by the proof of Comment 3(i), $\bar{\phi}_1(\Phi) \neq \phi_1^{n*}(\Phi)$ for any Φ, and hence is characterized by $\underline{\phi}_1(\Phi) = \phi_1^{n*}(\Phi)$ (and so $\phi_1^*(n') < \frac{1}{2}$). Recall that $\phi_1^{n*}(\cdot)$ is increasing and $\underline{\phi}_1(\cdot)$ decreasing in Φ. Also, since $\gamma^n(\Phi)$ is weakly increasing in n for all Φ, $\phi_1^{n*}(\Phi) = r_1(1-\gamma^n(\Phi))$ is decreasing in n for all Φ. These facts imply $\underline{\phi}_1(\Phi^*(n')) \geq \underline{\phi}_1(\Phi^*(n'+1))$. Then $\underline{\phi}_1(\cdot)$ decreasing in Φ implies $\Phi^*(n'+1) \geq \Phi^*(n')$. Thus $\Phi^*(n)$ is increasing in n for all $n \geq n'$.

Second, suppose $\phi_1^{n'*}(1-\frac{1}{4r_1}) > \frac{1}{2}$ for some n' (equivalently, $\gamma^{n'}(1-\frac{1}{4r_1}) < 1-\frac{1}{2r_1}$). If $\gamma^{n'+1}(1-\frac{1}{4r_1}) < 1-\frac{1}{2r_1}$, then $\phi_1^{n*}(\Phi) > \frac{1}{2}$ for all Φ and A1's equilibrium policy must lie along $\bar{\phi}_1(\cdot)$ (*i.e.*, $\phi_1^*(n'+1) > \frac{1}{2}$). Recall that $\bar{\phi}_1(\cdot)$ is increasing in Φ. Then by a symmetric argument with the first case, $\Phi^*(n'+1) < \Phi^*(n')$. Thus $\Phi^*(n)$ is decreasing in n for $n < n'+1$ such that $\gamma^{n'+1}(1-\frac{1}{4r_1}) < 1-\frac{1}{2r_1}$. If $\gamma^{n'+1}(1-\frac{1}{4r_1}) \geq 1-\frac{1}{2r_1}$, then from the first case $\Phi^*(n)$ is increasing in n for all $n \geq n'+1$. Note that either $\Phi^*(n') > \Phi^*(n'+1)$ or $\Phi^*(n') \leq \Phi^*(n'+1)$ may obtain, depending on the value of $r_{n'+1}$.

Combining results, if there exists n' such that $\gamma^{n'}(1-\frac{1}{4r_1}) < 1-\frac{1}{2r_1}$ and $\gamma^{n'+1}(1-\frac{1}{4r_1}) \geq 1-\frac{1}{2r_1}$, then $\tilde{n} = n'+1$. Since $\gamma^n(\cdot)$ is increasing in n, this condition is equivalent to $\tilde{n} = \min\{n \mid \phi_1^*(n) \leq \frac{1}{2}\}$. If no such n' exists, then $\tilde{n} = N+1$. ■

Proof of Proposition 3. (i) If $\Phi^*(n') - n'k' > \Phi^*(n'') - n''k'$ for some k' and $n'' > n'$, then clearly $\Phi^*(n') - n'k > \Phi^*(n'') - n''k$ for any $k > k'$. Thus, n^* is non-increasing in k.

(ii) Manipulating (6), $\Phi^*(n) \leq \Phi^*(1)$ for all n if $\phi_1^*(n) > 1-r_1$ for all n. As the result clearly obtains for $r_1 = 1$ I consider only the interior case: by (5) $\phi_1^*(n) > 1-r_1$ iff $\prod_{i=2}^N(1-\phi_i^*(n)) > \frac{1}{r_1}-1$. An upper bound on $\phi_i^*(n)$ is r_i, so a sufficient condition is $\prod_{i=2}^N(1-r_i) > \frac{1}{r_1}-1$. Since $1-\sum_{i=2}^N r_i < \prod_{i=2}^N(1-r_i)$, this is also satisfied if $\sum_{i=2}^N r_i < 2-\frac{1}{r_1}$.

(iii) If $r_1 \leq \frac{1}{2}$, then clearly $\tilde{n} = 1$. By Proposition 2 $\Phi^*(2) \geq \Phi^*(1)$, with the inequality strict if $r_2 > 0$. Thus for any $k < \Phi^*(2) - \Phi^*(1)$, $n^* > 1$. ∎

Proof of Comment 5. I first derive the optimal termination rule for P. First note that if $r_i = 1$ for some Ai, then Ai's dominant strategy is a policy of 1. Any termination rule is then optimal if the replacement rule is $\rho_\tau = \text{A}i\ \forall\tau$. Second, if $r_i = 0$ for some Ai, then Ai's dominant strategy is a policy of 0. In this case, a termination rule with $M_i = 1$ for all such Ai is trivially optimal.

Now consider agents such that $r_i \in (0,1)$. I show that Ai's policy choice, $\phi_i^{M_i*}$, is decreasing in M_i. For any $M_i > 1$, the following recursion defines Ai's value function upon its choice of some arbitrary ϕ_i:

$$v_i^{M_i}(\phi_i) = w_i\phi_i - m_i\phi_i^2 + \delta[\phi_i v_i^{M_i}(\phi_i) + (1-\phi_i)v_i^{M_i-1}(\phi_i^{M_i-1*})].$$

Thus, $v_i^{M_i}(\phi_i) = \frac{w_i\phi_i - m_i\phi_i^2 + \delta(1-\phi_i)v_i^{M_i-1}(\phi_i^{M_i-1*})}{1-\delta\phi_i}$, and for $M_i = 1$, $v_i^1(\phi_i) = \frac{w_i\phi_i - m_i\phi_i^2}{1-\delta\phi_i}$. It is obvious that $v_i^2(\phi_i^{2*}) > v_i^1(\phi_i^{1*})$. Manipulating $v_i^{M_i}(\phi_i)$, for any $M_i > 2$, $v_i^{M_i}(\phi_i^{M_i*}) > v_i^{M_i-1}(\phi_i^{M_i-1*})$ if $\exists\phi_i'$ such that $w_i\phi_i' - m_i\phi_i'^2 > (1-\delta)v_i^{M_i-1}(\phi_i^{M_i-1*})$. Let $\phi_i' = r_i = \phi_i^*(1)$. Since $r_i < 1$ (by assumption), termination occurs with probability one, and so $v_i^{M_i-1}(\phi_i^{M_i-1*}) < \frac{w_i r_i - m_i r_i^2}{1-\delta}$, establishing the inequality. Thus by induction $v_i^{M_i}(\phi_i^{M_i*})$ is increasing in M_i.

The optimal policy is derived by differentiating $v_i^{M_i}(\phi_i)$, which yields: $\frac{\mathrm{d}v_i^{M_i}}{\mathrm{d}\phi_i} = \frac{w_i - 2m_i\phi_i - \delta v_i^{M_i-1}(\phi_i^{M_i-1*})}{1-\delta\phi_i} + \delta\frac{w_i\phi_i - m_i\phi_i^2 + \delta(1-\phi_i)v_i^{M_i-1}(\phi_i^{M_i-1*})}{(1-\delta\phi_i)^2}$. To show that first-order conditions are sufficient, note $\frac{\mathrm{d}v_i^{M_i}}{\mathrm{d}\phi_i} < 0$ if and only if: $(1-\delta\phi_i)[w_i - 2m_i\phi_i - \delta v_i^{M_i-1}(\phi_i^{M_i-1*})] + \delta[w_i\phi_i - m_i\phi_i^2 + \delta(1-\phi_i)v_i^{M_i-1}(\phi_i^{M_i-1*})] < 0$. Rearranging terms, this condition is equivalent to:

$$\nu(\phi_i) = \delta m_i\phi_i^2 - \phi_i + w_i - \delta(1-\delta)v_i^{M_i-1}(\phi_i^{M_i-1*}) < 0.$$

Now note that $\frac{\mathrm{d}\nu}{\mathrm{d}\phi_i} = 2\delta m_i\phi - 2m_i$, so $\frac{\mathrm{d}\nu}{\mathrm{d}\phi_i} < 0$ if and only if $\phi_i < \frac{1}{\delta}$, which holds for all $\phi_i \leq 1$. Since $\nu(\cdot)$ is strictly decreasing on $[0,1]$, $v_i^{M_i}(\cdot)$ is pseudoconcave and first-order conditions are sufficient for characterizing a maximum. At an interior solution:

$$\phi_i^{M_i*} = \frac{1 - \sqrt{1 - 2\delta r_i + \delta^2(1-\delta)v_i^{M_i-1}(\phi_i^{M_i-1*})/m_i}}{\delta}. \quad (11)$$

Otherwise, at a corner solution or if (11) has no real roots, $\phi_i^{M_i*} = 0$ or 1. Since $r_i > 0$ (by assumption), Ai can achieve strictly positive utility by choosing some $\phi_i > 0$, and therefore 0 cannot be optimal. Thus, if (11) does not characterize $\phi_i^{M_i*}$, then $\phi_i^{M_i*} = 1$.

Since $v_i^{M_i}(\phi_i^{M_i*})$ is increasing in M_i, it is clear from (11) and the corner case that $\phi_i^{M_i*}$ is weakly decreasing in M_i. Thus a type-Ai agent chooses the highest policy when $M_i = 1$. For any replacement rule, P's optimal termination rule is therefore $\{M_i = 1\}$ for all Ai.

I now derive the optimal replacement rule. Given any rule $\{\rho_\tau\}$, P's expected payoff in each period is the expectation of some distribution over $\{\phi_i^{1*}\}$. This payoff is maximized if the policy chosen is $\max_i\{\phi_i^{1*}\}$ with certainty. By inspection of (11), the replacement rule $\{\rho_\tau = \text{A}1\}$ $\forall\tau$ must be optimal, since $r_1 \geq r_j$ for $j \neq 1$.

Finally, since not all termination and replacement rules are Nash, it will be necessary to verify that $\{M_i = 1\}$ $\forall i$ and $\{\rho_\tau = \text{A}1\}$ $\forall\tau$ is a best response. Substituting into (11), since $\phi_1^{1*} = \min\{1, \frac{1-\sqrt{1-2\delta r_i}}{\delta}\} > r_1 > r_j$ for $j \neq 1$, P achieves the highest policy by following these rules. ■

Proof of Proposition 4. If $r_1 \leq 0$, then $r_i \leq 0$ for all Ai, and 0 is a dominant strategy for all agents. Otherwise, substituting into (11), $\phi_1^{1*} = \min\{\frac{1-\sqrt{1-2\delta r_1}}{\delta}, 1\}$. ■

Proof of Proposition 5. Suppose $n = s$ and $r_i < 1$ for all Ai ($i \leq N$). Since $\mu_{-i}^q(\mathbf{0}) = 0$, by (10) each A$i$'s best response is $\phi_i^{s*}(n) = 0$. Thus $\phi^{s*}(s) = \mathbf{0}$ is an equilibrium of the subgame. To show uniqueness of this equilibrium, suppose $\phi_j^{s*}(n) > 0$ for some Aj. By (10), $\phi_i^{s*}(n) = \frac{w_i \mu_{-i}^{s-1}(\phi)}{2m_i} \in [0, 1)$ for each Ai. If $\phi_i^{s*}(n) = 0$ for any Ai, then $\mu_{-i}^{s-1}(\phi) = 0$ and thus $\phi_j^{s*}(n) = 0$: contradiction. Otherwise, if $\phi_i^{s*}(n) > 0$ for all Ai, then since $n = s$, $\mu_{-i}^{s-1}(\phi) < \phi_k$ for any $k \neq i$, and therefore by (10) $\phi_i^{s*}(n) < \min_{k\neq i}\{\phi_k\}$ for all Ai: contradiction.

Clearly, $\phi^{s*}(n) = \mathbf{0}$ for all $n < s$. Thus, P prefers 0 agents to j agents for $1 \leq j \leq s$. Thus, $n^{s*} = 0$ or $n^{s*} > s$. ■

REFERENCES

Armacost, Michael. 1969. *The Politics of Weapons Innovation: The Thor-Jupiter Controversy.* New York: Columbia University Press.

Austen-Smith, David, and Jeffrey S. Banks. 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review* 90(1): 34-45.

Banks, Jeffrey S., and Rangarajan K. Sundaram. 1998. "Optimal Retention in Agency Problems." *Journal of Economic Theory* 82(2): 293-323.

Barlow, Richard E., and Frank Proschan. 1965. *Mathematical Theory of Reliability.* New York: John Wiley and Sons, Inc.

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

Bendor, Jonathan B. 1985. *Parallel Systems: Redundancy in Government.* Berkeley, CA: University of California Press.

Brehm, John, and Scott Gates. 1997. *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public.* Ann Arbor: The University of Michigan Press.

Calvert, Randall L., Mathew D. McCubbins, and Barry R. Weingast. 1989. "A Theory of Political Control and Agency Discretion." *American Journal of Political Science* 33(3): 588-611.

Carpenter, Daniel P. 2000. "Stochastic Prediction and Estimation of Nonlinear Political Durations: An Application to the Lifetime of Bureaus." In *Political Complexity: Nonlinear Models of Politics*, ed. D. Richards. Ann Arbor: The University of Michigan Press, 209-238.

Carpenter, Daniel P. 2001. *The Forging of Bureaucratic Autonomy: Reputations, Networks, and Policy Innovation in Executive Agencies, 1862-1928.* Princeton: Princeton University Press.

Carpenter, Daniel P., and David E. Lewis. 2002. "Political Learning, Fiscal Constraints, and the Lifetime of Bureaus." Unpublished manuscript, Harvard University.

Chang, Kelly H., David E. Lewis, and Nolan M. McCarty. 2001. "The Tenure of Political Appointees." Unpublished manuscript, Princeton University.

Chisholm, Donald W. 1989. *Coordination Without Hierarchy: Informal Structures in Multiorganizational Systems.* Berkeley, CA: University of California Press.

Dixit, Avinash K. 1995. *The Making of Economic Policy: A Transaction Cost Politics Perspective.* Cambridge: MIT Press.

Donahue, John. 1991. *The Privatization Decision: Public Ends, Private Means.* Reprint ed. New York: Basic Books.

Epstein, David, and Sharyn O'Halloran. 1994. "Administrative Procedures, Information, and Agency Discretion." *American Journal of Political Science* 39(3): 697-722.

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

Esteban, Joan, and Debraj Ray. 2001. "Collective Action and the Group Size Paradox." *American Political Science Review* 95(3): 663-672.

Gailmard, Sean. 2002a. "Expertise, Subversion, and Bureaucratic Discretion." *Journal of Law, Economics, and Organization* 18(2): 536-555.

Gailmard, Sean. 2002b. "Multiple Principals and Outside Information in Bureaucratic Policy Making." Unpublished manuscript, University of Chicago.

Heimann, C. F. Larry. 1993. "Understanding the Challenger Disaster: Organizational Structure and the Design of Reliable Systems." *American Political Science Review* 87(2): 421-435.

Heimann, C. F. Larry. 1997. *Acceptable Risks: Politics, Policy, and Risky Technologies.* Ann Arbor: The University of Michigan Press.

Hirshleifer, Jack. 1983. "From Weakest-Link to Best-Shot: The Voluntary Provision of Public Goods." *Public Choice* 41(3): 371-386.

Holmstrom, Bengt. 1982. "Moral Hazard in Teams." *The Bell Journal of Economics* 13(2): 324-340.

Kaufman, Herbert. 1976. *Are Government Organizations Immortal?* Washington, DC: The Brookings Institution.

Kiewiet, D. Roderick, and Mathew D. McCubbins. 1991. *The Logic of Delegation: Congressional Parties and the Appropriations Process.* Chicago: University of Chicago Press.

Kremer, Michael. 1993. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics* 108(3): 551-575.

Landau, Martin. 1969. "Redundancy, Rationality, and the Problem of Duplication and Overlap." *Public Administration Review* 29(4): 346-358.

Landau, Martin. 1991. "On Multiorganizational Systems in Public Administration." *Journal of Public Administration Research and Theory* 1(1): 5-18.

LaPorte, Todd R., and Paula M. Consolini. 1991. "Working in Practice But Not in Theory: Theoretical Challenges of High-Reliability Organizations." *Journal of Public Administration Research and Theory* 1(1): 19-47.

Lerner, Allan W. 1986. "There is More Than One Way to be Redundant." *Administration and Society* 18(3): 334-359.

Lewis, David E. 2002. "The Politics of Agency Termination: Confronting the Myth of Agency Immortality." *Journal of Politics* 64(1): 89-107.

Lohmann, Susanne, and Hugo Hopenhayn. 1998. "Delegation and the Regulation of Risk." *Games and Economic Behavior* 23(2): 222-246.

Maass, Arthur. 1951. *Muddy Waters.* Cambridge: Harvard University Press.

Miller, Gary, and Terry M. Moe. 1983. "Bureaucrats, Legislators, and the Size of Government." *American Political Science Review* 77(2): 297-322.

Moe, Terry M. 1989. "The Politics of Bureaucratic Structure." In *Can the Government Govern?*, ed. J. E. Chubb and P. E. Peterson. Washington, DC: The Brookings Institution, 267-329.

Montinola, Gabriella. 2002. "The Efficient Secret Revisited: The Emergence of 'Clean Government' in Chile." Unpublished manuscript, University of California, Davis.

Niskanen, William. 1971. *Bureaucracy and Representative Government.* Chicago: Aldine-Atherton.

Osborne, David E., and Ted Gaebler. 1992. *Reinventing Government: How the Entrepreneurial Spirit is Transforming the Public Sector.* Reading, MA: Addison-Wesley.

Perrow, Charles. 1984. *Normal Accidents: Living with High-Risk Technologies.* New York: Basic Books.

Rochlin, Gene I., Todd R. LaPorte, and Karlene H. Roberts. 1987. "The Self-Designing High-Reliability Organization: Aircraft Carrier Flight Operations at Sea." *Naval War College Review* (Autumn): 76-90.

Sagan, Scott D. 1993. *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons.* Princeton: Princeton University Press.

Sah, Raaj K., and Joseph E. Stiglitz. 1986. "The Architecture of Economic Systems: Hierarchies and Polyarchies." *The American Economic Review* 76(4): 716-727.

Salamon, Lester. 1978. *Welfare: The Elusive Consensus.* New York: Praeger.

Sapolsky, Harvey. 1972. *The Polaris System Development.* Cambridge: Harvard University Press.

Silver, Jay. 1996. "Influencing Agency Policies Through Congressional Appropriations." Unpublished manuscript, Stanford Graduate School of Business.

Sobel, Joel. 1992. "How to Count to One Thousand." *The Economic Journal* 102(410): 1-8.

Ting, Michael M. 2001. "The 'Power of the Purse' and its Implications for Bureaucratic Policy-Making." *Public Choice* 106(3-4): 243-274.

Wilson, James Q. 2000. *Bureaucracy: What Government Agencies Do and Why They Do It.* 2nd ed. New York: Basic Books.

Figure 2: Number of Agencies as a Function of k and r_1 (N = 2)

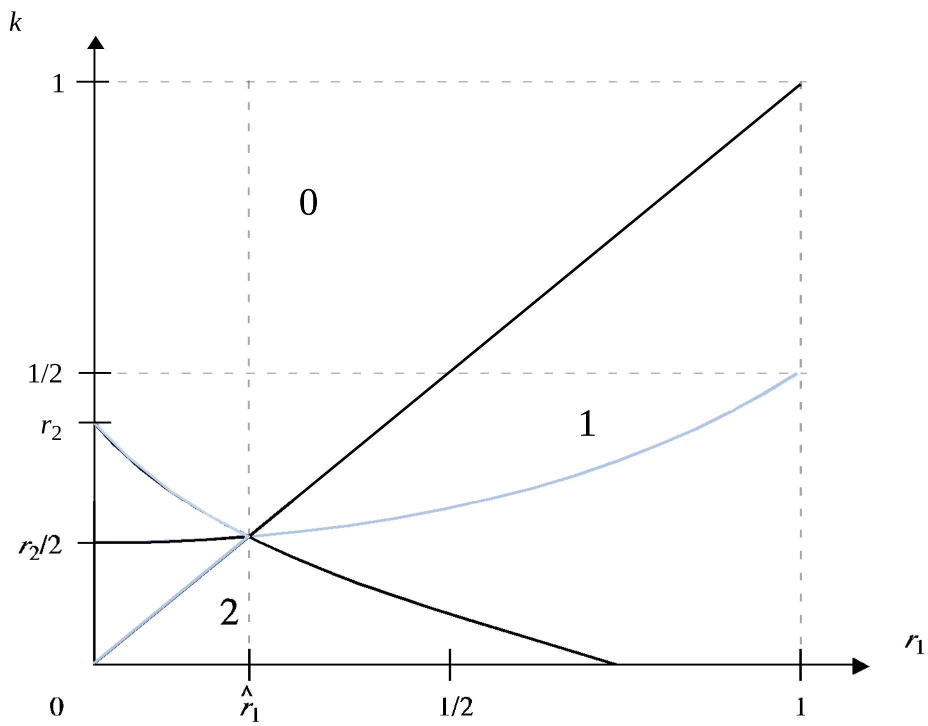

Figure 3: Game Theoretic and Classical Redundancy