

THEORIES OF DELEGATION

J. Bendor¹, A. Glazer², and T. Hammond³

¹*Graduate School of Business, Stanford University, Stanford, California 94305-5015,
e-mail: bendor_jonathan@gsb.stanford.edu*

²*Department of Economics, University of California, Irvine, California 92697;
e-mail: aglazer@uci.edu*

³*Department of Political Science, Michigan State University, East Lansing,
Michigan 48824; e-mail: thammond@msu.edu*

Key Words information, ally principle, hierarchy, credible commitment, noncooperative game theory

■ **Abstract** We survey modern models of delegation that assume a boss and a subordinate pursue their own goals. Among the major themes covered are the following: the conditions under which the boss will prefer to delegate versus those in which she will prefer to retain authority; how a boss can induce a subordinate to truthfully reveal information; when rational principals will use the ally principle (i.e. delegate to agents with similar goals); delegation in repeated interactions; and how delegation can overcome commitment problems. These themes are relevant to a wide variety of institutions, affecting intralegislative organization, executive-legislative relations, and central banks.

INTRODUCTION

Even God delegates, having told the Israelites after they left Egypt: “Behold, I send an Angel before thee, to keep thee in the way, and to bring thee into the place which I have prepared. Beware of him, and obey his voice, provoke him not; for he will not pardon your transgressions: for my name is in him” (*Exodus* 23:20–21). Delegation, then, may be done even by an all-powerful and an all-knowing entity. Yet the limitations of leaders have long been recognized as a reason for delegation. When Moses complained, “O my Lord, I am not eloquent, neither heretofore, nor since thou hast spoken unto thy servant: but I am slow of speech, and of a slow tongue” (*Exodus* 4:100), God’s solution was that Aaron should be Moses’s spokesman.

Recent years have seen heightened academic interest in the problem of delegation, as well as the development of sophisticated and insightful theories and models that explain various aspects of delegation. This essay surveys theories that have appeared since the last review of this topic (Aranson et al 1982). It also provides

a framework for thinking about delegation, both for legislative-executive and for intra-institutional relations.

Many of the most influential works surveyed here use formal models, particularly noncooperative game theory, as a central method of analysis. That too is a focus of our survey. We hope our survey will help the reader organize the literature and serve as an informal tutorial for political scientists who study delegation but who are not steeped in game theory. And although we consider mostly formal models of delegation, we do consider informal theories. Several of the best-known works are purely verbal, although most are inspired by rational choice theory. Further, the substantive differences between informal theories and formal models are rather small. Hence we treat them together, separating by content rather than form. In discussing results, we emphasize the ally principle. When will rational principals delegate to agents whose preferences are like theirs? That rational bosses prefer agents who are allies is part of an ancient wisdom about political hierarchies; the principle was old before Rome. But as we shall see, in some empirically significant situations, smart superiors will not use it.

The paper is organized as follows. The next section introduces the game-theoretic framework. It analyzes how problematic a delegation situation is by examining the decision makers' choices and payoffs (i.e. the normal form of delegation games). The third section analyzes the extensive form of delegation games (who does what when). The framework is then used to categorize the literature and serves as a springboard for examining it. The fourth section covers delegation in repeated interactions. Theories of delegation as credible commitment are discussed in the fifth section.

A TYPOLOGY OF DELEGATION SITUATIONS

Some elementary noncooperative game theory is useful in constructing a typology of delegation situations. Our typology is deliberately spare; it is constructed out of the minimal set of elements required for an analysis of delegation. Hence, we consider only two decision makers: a boss (or principal) and a subordinate (or agent). Any fewer than two makes delegation impossible; any more is excess baggage. Similarly, each player has two options: The boss can either delegate or not; the subordinate can either work or shirk. (There are many different interpretations of the work-or-shirk choice: specialize or remain uninformed, etc. Essentially, the subordinate can either act in a way that is good for the boss or not do so; see assumption 1, below.) So there are four possible outcomes, and we assume that each person has a strict preference ordering over the four. (Indifference—a knife-edge condition—is an unnecessary complication.) Figure 1 shows the normal form of the game. We assume that each player knows what game he or she is playing (i.e. the feasible actions and payoffs), each knows the other knows, and so on.

Different preference orderings identify different kinds of delegation situations, thus creating a typology. We do not discuss all possible orderings; there are too

		Agent	
		Work	Shirk
Boss	Delegate	a, w	b, x
	Control	c, y	d, z

Assumption 1: Agent's effort is valuable for the boss: $a > b$
and $c > d$.

Assumption 2: Both players prefer (delegate, work) to (control,
shirk): $a > d$ and $w > z$.

Figure 1 A typology of delegation situations.

many,¹ and some are irrelevant for theories of delegation. Instead, we examine a subset that represents a spectrum of delegation situations, ranging from difficult (delegation is valuable but problematic) to easy. We create this subset by restricting the preference orderings in two ways.

Assumption 1 (A1): The subordinate's effort is valuable for the boss.

We assume that, all else being equal, the boss prefers that her subordinate work rather than shirk. Thus, she is better off, for a given move of hers, if the agent works: $a > b$ and $c > d$. For example, in many of the delegation models analyzed here, the subordinate "works" by acquiring information about a random variable that is unobserved by the boss. Since the boss is typically assumed to be risk averse, she prefers an informed, autonomous subordinate to an ignorant, autonomous one; similarly, she does not prefer an ignorant, controlled agent to an informed, controlled one.

A parallel assumption that we do not make is that the subordinate must prefer having discretion rather than being controlled (holding fixed his move). We are agnostic about this property; there probably are important situations in which it does not hold, e.g. because the subordinate wants the boss held responsible for the outcome, which is more likely if she controlled the reins.

Assumption 2 (A2): Deals between the boss and the subordinate are possible because both prefer the outcome of (delegate, work) to that of (control, shirk).

¹Each player has 4! preference orderings, creating $(24)^2 = 576$ possible games.

This assumption is at the heart of the normative theory of institutions embraced, explicitly or implicitly, by most of the delegation literature. Although the two actors might conflict over effort—the boss wants the subordinate to work hard whereas the latter might prefer to take it easy—complete conflict does not prevail because there is at least one thing that both can agree on: (control, shirk) is collectively—i.e. Pareto—suboptimal. Thus, there is room for a compromise that would prevent the Pareto-inferior outcome. Simultaneously, A2 creates the possibility that the institution might be caught in a bad equilibrium.²

Given A1 and A2, the boss has 5 possible orderings; the subordinate, 12. This leaves a much more manageable subset of 60 games, which fall into 5 major categories. We start with the one in which the delegation “problem” is least pressing and move on to tougher cases.³

Category 1: The game has a unique, Pareto-optimal Nash equilibrium.

Consider a busy boss facing many important issues. The task at hand, however, is relatively minor. Given her crowded agenda, the opportunity costs of giving detailed instructions to her subordinate are high. So she has a dominant strategy of delegating: $a > c$ and $b > d$. If the subordinate works hard, great; if not, so be it. Now suppose that the subordinate has no dominant strategy. Instead, his best move depends on what the boss does. If she exerts control, then he optimizes by shirking; if she delegates, then he should exert himself. Since the boss will inevitably delegate, the predicted outcome is (delegate, work).⁴ This is benign; one can easily show (given A1 and A2 and the particular assumptions defining this case) that this outcome is efficient. This must often occur in organizations, but students of delegation have mostly ignored it, and for good reason—it is totally unproblematic. Being busy, the boss delegates and takes the consequences. The outcome is both stable and efficient. End of story.

Alternatively, suppose that the agent has a dominant strategy of working hard, possibly because the project at hand is personally significant, whereas the boss has a dominant strategy of exerting control. Hence the outcome is (control, work). Here delegation does not occur, but the important point is that this situation is just as unproblematic as the first example. This is an efficient Nash equilibrium, just as (delegate, work) was. And because it is supported by two dominant strategies, it is very stable.

²Given games with exclusively Pareto-optimal outcomes, one could not address the issue of stable yet inefficient outcomes.

³One should note, however, that some delegation games are easy only because another game—e.g. of personnel selection—has been played, and played well. For example, if the subordinate has a dominant strategy of working hard, that says something about the organization’s personnel policies.

⁴Because the boss has a strictly dominant strategy, this is the predicted outcome regardless of the timing of moves. So a normal form analysis suffices here. Timing matters in other delegation games, as we will see shortly.

Category 2: There are multiple equilibria, but they are Pareto-ranked.

In all category 1 games, at least one player has a dominant strategy. So these games are strategically simple (given the informational assumptions); for example, the sequence of moves does not matter.

But now consider what happens if the boss's best move depends on what her subordinate does and vice versa. For example, suppose that both have "deal-making" preferences, i.e. if the boss is willing to delegate then the subordinate will work hard and vice versa, but if the boss wants to exert control then the subordinate wants to shirk, and if the subordinate shirks then the boss wants to run things. Then there are two Nash equilibria, which by A2 are Pareto-ranked. (The better equilibrium is not only Pareto-superior to the bad one, it must also be Pareto-optimal.) Thus, delegation here amounts to a coordination problem, handled easily by ordinary communication or signaling via moves. Suppose, for instance, the boss moves first. If she delegates, the subordinate will work hard; if she controls, he will shirk. Anticipating this and preferring the former to the latter, she delegates.

Category 3: There are multiple equilibria that are not Pareto-ranked.

Out of the 60 games we are studying, only 2 have multiple equilibria that involve conflict. The boss prefers to play one equilibrium and the subordinate the other. In one of these, the boss prefers an outcome in which good work is done to one in which it is not ($c > b$), but if the subordinate shirks, then the boss would rather not be closely associated with the unsatisfactory outcome ($b > d$). The subordinate, however, prefers the laid-back equilibrium of (delegate, shirk) to the intense one of (control, work). Hence, conflict over equilibria.

Here the sequence of moves is important, as it is in Chicken and other games with multiple, conflict-laden equilibria. There is a first-mover advantage. If the boss moves first and takes charge by issuing detailed orders, then the subordinate is compelled to work hard. But if the latter can preempt the former by a passive form of worker sabotage, then the boss will distance herself from the issue by delegating. Hence, delegation in this situation is genuinely problematic.

Category 4: The game has a unique, Pareto-inferior Nash equilibrium.

This class of situations is studied intensively in the literature. In almost all of these games, the deficient equilibrium is (control, shirk). Whenever this holds, the culprit is a "bad" dominant strategy of either or both players. For example, suppose in an intralegislative game (e.g. Gilligan & Krehbiel 1987, Diermeier 1995) the floor has a dominant strategy of asserting control over the bill in the last stage. The committee lacks a dominant strategy; instead, it has the conventional "deal-making" conditional best responses (specialize if the floor delegates, otherwise remain uninformed). Given that the floor will control the process no matter what the committee does, the latter will not specialize. The result is a collectively bad equilibrium—the very stuff of much of the delegation literature.

All category 4 games resemble the Prisoner's Dilemma in that collective optimality and individual rationality diverge: Whatever is Pareto-optimal is not Nash.

The sequence of moves is irrelevant in these games, since one or both sides have dominant strategies.

Category 5: No pure Nash equilibria exist.

In almost all of the literature's models of delegation, the games have pure Nash equilibria. However, category 5 shows that this is not inevitable. Suppose that the boss and the subordinate are struggling over the allocation of credit (should the project go well) or blame (if things go badly). The boss can distance herself from the process and avoid some blame if she delegates. But if things work out well, then she would like to claim credit by putting her stamp on the process, i.e. by controlling the subordinate. For the subordinate, however, working hard is not worthwhile if the boss steals the credit, so in this event the subordinate shirks—whence the boss prefers to delegate. But then the credit is there for the taking, so the subordinate wants to work hard. And around they go. Thus, the only Nash equilibrium involves both sides randomizing. Although we have not shown that this outcome is Pareto-inferior, most real-world decision makers would probably consider it a delegation problem. If this pattern persists, then the institution is floundering, unable to settle on a stable arrangement.

WHO DOES WHAT WHEN? THE EXTENSIVE FORMS OF DELEGATION GAMES

Although a delegation game's normal form helps us to categorize delegation problems by their degree of difficulty, it does not represent the timing of moves—who does what when—which helps us understand a game's dynamics. Nor does the normal form tell us who knows what pieces of information when different actions are taken. In contrast, the extensive form shows both who knows what when and who does what when. This feature of delegation games is obviously important when there are informational asymmetries between principal and agent, and scholars in this field tend to regard such asymmetries as central to delegation. Hence, game-theoretic models increasingly use the extensive form. Indeed, one can use this type of analysis to categorize much of the literature by describing assumptions about who does what when and who knows what when. First, we briefly describe two major types of models in these terms; later we add details.

In one class of models, the boss has a choice of either delegating to an agent or not. If she delegates, then the agent makes the choice; if she does not delegate, then she makes the choice herself. Typically what drives these models is that the agent can become more informed than the boss. If the agent is delegated authority, then he has an opportunity to acquire information before he chooses an action, whereas the boss must decide in the face of uncertainty (Figure 2a).

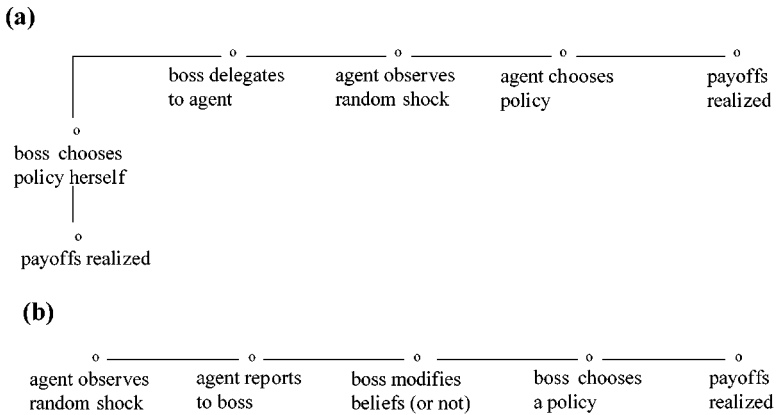


Figure 2 Two types of delegations games: (a) Delegation-of-authority game, (b) signaling game.

The other main class of models typically reverses the sequence; the agent moves before the boss (for an exception, see Gilligan & Krehbiel 1987). Again the agent may choose to become informed. If he acquires private information, he then reports to the boss. The report may either describe (all or some of) the information gathered by the agent, or it may recommend an action.⁵ The boss then decodes the message or the information embedded in the recommended action, and given this information, she chooses an action (Figure 2*b*). These are called signaling models. Models in which the agent can send a message to the boss for free are called cheap-talk models (as in “talk is cheap”).

These two types of models represent substantively different kinds of delegation. In the first type, the boss explicitly decides whether or not to give a subordinate the authority to take actions, without requiring that the agent first report back to the boss (e.g. about the information that the agent has gleaned). Presumably time matters here. In the second type, the agent has been given different kinds of authority: either to gather certain kinds of information or to make proposals to the boss.⁶ Here the agent lacks the authority to take actions; only the boss has that right.

We now take a closer look at these two major types of models.

⁵In some variants, there is little formal difference between reporting information and recommending an action, since the boss may be able to infer the agent’s private information from his recommended action. In other variants, however, there is a substantive difference between reporting information and recommending an option; the latter may include constraints on the boss’s choice set (e.g. she can choose only between the recommended alternative and the status quo; see Gilligan & Krehbiel 1987). We return to this distinction shortly.

⁶These models typically presume—sensibly—that agents authorized to make proposals are also authorized to gather information.

Delegating the Authority to Act

In the first type of model, the boss moves first by deciding whether or not to delegate the authority to pick policy to one of several possible agents. If the boss controls the issue, then she picks a policy, a random shock is realized, and the outcome and the payoffs are generated by the combination of the selected policy and the shock. If she chooses to delegate, she picks an agent from the feasible set and authorizes him to select a policy. After being picked, the agent learns the value of the disturbance and then chooses a policy. Consequently, he makes a more informed choice than the boss would.⁷

Probably no paper in the literature has used only this stripped-down model, but even this simple formulation illustrates what is generally regarded as the major trade-off facing the boss: Is the gain produced by delegating the decision to a more informed party worth the loss produced by having someone with different preferences make the choice? So before we turn to the inevitably more complicated variants, let us see what this unadorned model can tell us about the fundamental trade-off between expertise and control.

Consider a boss and a set of agents. Each has single-peaked preferences over outcomes in a unidimensional space. Preferences are quadratic, so agent i 's utility equals $-(x - x_i)^2$, where x_i is i 's ideal point. Quadratic preferences imply that agents are risk-averse; for every feasible outcome, they prefer that outcome (for sure) to any lottery whose mean is that outcome.

Actors choose policies rather than outcomes directly. Policies are perturbed by random shocks (ϵ) to yield outcomes. The conventional assumption is that outcomes equal policies plus the random error: $X = P + \epsilon$. Most of the models assume, for the sake of tractability, that the shock is distributed uniformly (for more general formulations, see Bawn 1995 and Bendor et al 2000). It is convenient to assume that ϵ is centered around zero.

The basic model uses a substantively important (though common) assumption about the relation between extreme values of shocks and policies: If someone chooses a policy after he observes the disturbance, then he can always find a policy that, when added to the shock, yields his ideal point in the outcome space. Thus, the agent is a perfect shock absorber.

All of the above is common knowledge. Given that the agent moves last—there is no ex post auditing in the basic model—it does not matter whether the boss observes the random disturbance or not. It is natural, however, to assume that asymmetric information prevails; the agent knows both the shock and his chosen policy, but the boss observes only the outcome.

Unlike signaling models, which typically have multiple equilibria (see footnote 10 below), this kind of delegation game has a unique equilibrium (setting aside knife-edge possibilities, e.g. two agents who are equidistant from the boss's ideal point). The equilibrium has the following properties:

⁷In many of the models, the agent learns the exact value of the shock; for exceptions, see Bawn (1995) and Bendor et al (2000).

1. The ally principle holds here. If the boss delegates, then she picks the agent whose ideal point is the closest to hers. The reason is simple. Since the agent knows the value of the shock before he chooses a policy and can exploit compensatory flexibility to the hilt, he picks his induced ideal point in policy space, i.e. the one that yields his ideal point in outcome space. The boss naturally wants this outcome to be as close as possible to her own ideal.
2. If the boss does not delegate, then she chooses the policy that maximizes her expected utility, given her knowledge of the shock's distribution and the policy "technology" (outcomes = policy + noise). If ϵ is distributed uniformly around zero, then the expected outcome is her ideal point.
3. The boss has a pair of certainty equivalents in outcome space: She is indifferent between getting such an outcome (for sure) and choosing not to delegate. Hence, these two certainty equivalents define a delegation cutoff rule for the boss: If there is an agent who is ideologically closer to the boss than this threshold, then she delegates; if not, she does not. Everything inside this boundary we call the boss's delegation set.
4. The equilibrium outcome is efficient if and only if the boss delegates. Delegation must be efficient because the agent who has been granted authority gets his ideal outcome, which he strictly prefers to anything else. If the boss controls the policy, then the outcome must be inefficient because everyone is risk-averse and so would prefer that the boss get her ideal outcome for sure (instead of in expectation). This could be achieved if the boss delegated to an agent who promised to implement that outcome and delivered on the promise. That arrangement is technically feasible, but such promises are not credible and so the boss disregards them. Thus, the inability of agents to credibly commit, plus risk aversion, leads to inefficient outcomes (when no agent is in the delegation set).

Comparative Statics of the Basic Model

Let us now see how variations in several of the model's key parameters affect the equilibrium. [All these results are proven for the general N -dimensional setting in Bendor et al (2000).]

Increased Conflict Suppose conflict increases; the ideal points of all agents move away from the boss's. What happens? There are three possibilities. First, if previously no agent was in the delegation set, then nothing changes; the boss continues to exert control. Second, if previously the best agent had been in the delegation set but the shift makes him unacceptable, then the boss switches from delegating to controlling, and the expected outcome shifts toward the boss's ideal

point. Third, if the change in preferences is small enough and the delegation set had not been empty, then the boss would still want to delegate. In this case, the expected outcome moves away from her ideal point.

As the best agent's ideal point moves further from the boss's, the latter's expected utility falls or stays constant. Thus, the intuition that the boss is better off with more like-minded agents is upheld in this setting.

Increased Policy Uncertainty Suppose there is an increase in policy uncertainty. Because the boss is risk-averse, this change reduces her expected utility of not delegating. But the value of delegating is unchanged, since the best agent will still be picked and he will attain his ideal outcome. Hence, the boss becomes less choosy; her delegation cutoff moves away from her ideal point. Thus, the value of delegating, relative to not delegating, increases in policy uncertainty.

Increased Aversion to Risk Now hold policy uncertainty constant and make the players more averse to risk. Then the expected utility of controlling the issue falls for the boss, since the (unchanged) uncertainty hurts more. Thus, the boss's tolerance for delegating increases; the delegation set grows larger.

Variants of the Basic Model

We now consider some important variants of the basic delegation-of-action model.

Personnel Uncertainty Suppose the boss can make an educated guess about subordinates' ideal points but does not know them precisely. Consequently, she does not know precisely what a subordinate would do if given authority over the issue. It is fairly straightforward to show that this reduces the value of delegating, if the boss is risk-averse (e.g. Bawn 1995, Bendor et al 2000).

Multiple Principals For most democratic political systems, grants of delegation are made by multiple principals. Even in unicameral parliamentary systems, as in England, grants of delegation are made by the entire Parliament, perhaps by some subset of the Parliament (such as the majority party), and perhaps even by some subset of the majority party (such as the Cabinet), but only rarely by the Prime Minister alone. Similarly, grants of delegation in the United States are made not by some unicameral "Congress" but by the House and Senate, each relying extensively on its own set of committees, and by a President with a veto (which of course may be overridden).

We know of only a few models of delegation in which multiple principals with vetoes play a major role (Calvert et al 1989, Volden 2000; we take up Volden's contribution below). Calvert et al model the appointment of the agent—a bureau chief—as a Nash bargaining game,⁸ set in an N -dimensional policy space, between

⁸The appendix discusses the choice of modeling appointments by cooperative game theory (the Nash bargaining solution) or noncooperative theory (the Rubinstein alternating offer game). Today's intellectual fashions favor the latter.

Congress (a unitary actor) and the President. They show (1989:595) that if the ideal point of either principal changes, then the ideal point of the selected bureau chief will move in the same direction. (However, this result presumes that the set of feasible agents corresponds to the entire set of feasible policies.) Their model of appointments also suggests an interesting reconceptualization of agency discretion. That an agency sets policy without direct interference from politicians “does not constitute agency discretion. Rather, agency discretion occurs when the agency succeeds in choosing a policy in line with agency goals, when these goals *differ* from what the executive and legislature expected *at the appointment stage*” (Calvert et al 1989:605, original emphasis).

Hammond & Knott (1996) explicitly discuss problems that multiple principals (e.g. House, Senate, and President) might have in controlling an agency. Their model also has implications for a theory of delegation. In particular, if principals with conflicting objectives have difficulty agreeing on a policy for the agency to adopt (e.g. because of coordination costs), they could at least agree on limiting the agency’s discretion so that it cannot adopt a policy that is worse for any principal than the initial status quo. However, because these coordination costs are not explicitly represented, the model can only characterize the upper bound of acceptable discretion. An explicit model of coordination costs might well show that the agency’s actual range of discretion would fall inside this upper limit.

Multiple Agents Although as a modeling strategy it makes sense to start off with dyadic theories—a boss delegating authority to just one agent—few interesting cases in the real world of politics involve only a single agent. Most of the time, authority is delegated to a substantial group of people. This leads to two types of interesting and important problems.

One type pertains to shirking and free-riding. When just one person is given authority, that individual can certainly shirk. But when authority is delegated to several people, they can not only shirk but also free-ride on each other’s efforts. This can greatly reduce the total effort supplied by the subordinates. There is a substantial literature in economics on the design of incentive systems for motivating groups of individuals [see Miller (1992) for an excellent and accessible introduction to this literature]. However, a problem can plague these systems: The boss has her own kind of shirking, which entails manipulating the subordinates’ incentive system for her own benefit (Miller 1992:Ch. 7, Miller & Hammond 1994). Thus, even if problems of the subordinates shirking and free-riding can be overcome, properly motivating the boss herself remains a serious problem. Miller (1992:Ch. 20–21) argues that this problem can be ameliorated if the boss can credibly commit to leaving the subordinates’ incentive system alone. The topic of credible commitment is of major importance in the study of delegation, and we return to it below.

The second class of problems arises from a common mode of delegation, wherein a large problem is broken down into several pieces that are doled out

to different subunits of the organization (as in, e.g., Shepsle 1979). Each subunit makes what it perceives as the best choice on its own part of the problem, and then the subunits' decisions are aggregated by higher officials into an overall choice. Because each piece of the problem is assigned to just one subunit, shirking and free-riding may be avoided. However, this aggregation process has other undesirable properties. For example, Hammond & Miller (1985), reinterpreting a social choice theorem by Sen (1970), show that this delegation-and-aggregation can produce a Pareto-suboptimal outcome. This possibility seems to be virtually inherent in this kind of multi-person, multi-jurisdiction delegation.

Ex Ante Controls on Agents McCubbins et al (1987, 1989; see Bawn 1995 for a formalization) argue that ex ante controls on agencies sometimes work better than ex post controls. The latter, they assert, can be hindered because a bureau's policy deviation may benefit some of the many principals (e.g. the President, senators, representatives) who passed the legislation in question, and those coalition leaders will block efforts to bring the agency back into line. [This is the explanation for bureaucratic autonomy in Hammond & Knott (1996).] Thus, any one principal's threat to impose sanctions ex post would not work, and the agency's policy choice would stand. McCubbins et al suggest that Congress and the President respond to this weakness of ex post controls by imposing administrative procedures as ex ante controls. [Bawn (1997) and Huber & Shipan (1999) present formal models of delegation in which ex ante and ex post controls can be either complements or substitutes.]

McCubbins et al (1987, 1989) argue that two classes of administrative procedures are important. One class mandates information revelation by the agency. For example, the agency may be required to follow notice-and-comment procedures and provide reasons for its regulatory proposals. These procedures ensure that the agency's political environment mirrors that faced by the principals; the principals' constituents can then challenge a new agency policy before it is adopted. The other class of procedures "stacks the deck" on behalf of the principals' constituents by requiring the agency to heed the demands of particular actors and rely on particular kinds of decision-making standards. For example, statutes can specify that only certain actors have legal standing in agency proceedings, and burden-of-proof requirements can protect the economic and legal positions of favored constituents. These procedures together serve as an autopilot for the agency; as the principals' political environment changes in a particular way, so will the agency's. Overall, McCubbins et al suggest, these procedural tools allow principals to get what they want from the agency, even in a changing environment, without having to pass new legislation (which might be impossible, given disagreements among the multiple principals).

Arnold (1987), Robinson (1989), Horn & Shepsle (1989), Hill & Brazier (1991), and Mashaw (1994) argue that such procedures will be less effective than McCubbins et al suggest. For example, information revelation cannot overcome the impact of multiple principals. Even if an agency must announce its intended

decision before final adoption, objections by some principals may be negated by support from others. In such cases, the agency could proceed with its decision without fear of ex post retribution. Moreover, information-revelation procedures can cut two ways: Sometimes the publicity forced by the procedures can help the old coalition maintain its policies, but they might impede these efforts by alerting a broader public. And even if the agency's new political environment mirrors that of the original coalition, enmeshing agency decision making in a complex web of procedures might hinder agency adaptation to this new environment. Most generally, the critics argue that precisely because administrative procedures have political consequences, conflicting principals may find it hard to agree on procedures that will induce the agency to select a particular policy. The less the principals can agree on specific policies in the first place, the less they will be able to agree on administrative procedures that would later lead to a specific policy.

Choosing Institutions Fiorina (1986; see also 1982) presumed that delegation will occur; the question he studied was, to what kind of institution will authority be delegated? Substantively, he focused on the debate in the late 1800s over controlling the railroads' pricing practices. Proponents of regulation had to choose how regulatory authority would be delegated. Should it be delegated by creating a regulatory commission with broad powers? Or should Congress pass a law creating legal standards for railroad behavior, which plaintiffs could then use in court? Congress chose the former course when it created the Interstate Commerce Commission. Why was this choice made?

To answer this question, Fiorina sketched a simple decision-theoretic model. Each legislator compared the expected utility of agency (commission) enforcement of railroad regulation with that of court enforcement, preferring the enforcement method with the higher expected utility. More specifically, each legislator had beliefs, represented by a subjective probability distribution, about the level of regulation that either institution would actually enforce. Regarding preferences, Fiorina assumed that legislators had single-peaked and concave utility functions and so were risk-averse. However, he dropped the conventional assumption of symmetric utility functions. (Indeed, he gave reasons for believing that in this domain they would usually be asymmetric.) The preferences of the median legislator would then determine the legislature's choice. (Presidents played little role in Fiorina's model.)

Because Fiorina was interested in the congressional concerns in the late 1800s about the long-run behavior of the agencies and the courts, he sketched how time would affect the thinking of legislators. In his model, however, time periods entered only via the probability distributions of anticipated effects; short-run versus long-run concerns were not explicitly analyzed. Nonetheless, his model—one of the first formal models of delegation in political science—was an important precursor of the many models that followed over the next 15 years.

Delegation, Information, and Risk Aversion

Almost all of the preceding models (and most of the signaling models as well) assume a risk-averse boss. They either presume or explicitly argue that delegation and risk aversion are tightly connected. The idea is straightforward; delegating to an informed agent can reduce risk, which must benefit a risk-averse principal.

The connection certainly exists, but the literature overstates it. More precisely, an uninformed principal can benefit from delegating to an informed agent even if the boss is risk-seeking or risk-neutral. A simple example will show how. To sharpen the point, we consider a case with nonspatial preferences, so that the principal can in a natural way be risk-neutral over all possible outcomes. Suppose that in an urban machine a (real) boss is thinking about accepting a bribe for a large public works contract. If he takes the bribe he gets $b > 0$; otherwise he gets nothing. However, there is a chance ($p > 0$) that a federal prosecutor will pay attention to the contract; if he does, then he will catch the boss, recover the bribe, and fine him an amount $c > 0$. So the expected amount of money from taking the bribe equals $(1 - p)b - pc$. Suppose this exceeds zero, and the boss is risk-neutral in money. Then he will take the bribe (and the consequences).

But now suppose that he has someone in the district attorney's office to whom he can delegate the matter, for a fee. As an insider, the agent will know for sure whether the DA is watching the public works contract. If the DA is monitoring it, then the agent turns down the bribe, on the boss's behalf. But if not, then the agent is authorized to go through with the deal. (We ignore the possibility that the agent might doublecross the boss.) The expected payoff to the boss is $(1 - p)b + p \cdot 0$, minus the agent's fee. Clearly, if the fee is less than pc , then the boss is better off delegating to the informed agent than he is making the choice by himself, ignorant of the prosecutor's actions.

What drives this example is that the optimal strategy is state-contingent: Accept the bribe if the law isn't looking; otherwise, don't. The optimality of state-contingent plans does not depend on risk aversion. Thus, the value of delegating to an informed agent can turn on the value of information alone; risk aversion is unnecessary.

It is true that the more risk-averse the boss is the more valuable the agent's information is, and the more the boss would be willing to pay him.⁹ But this comparative static effect does not undermine the general point. Delegating to an informed agent is optimal for an uninformed principal, but this does not imply that the principal must be risk-averse. Hence we conclude that although uncertainty (or risk) may be central to delegation, risk aversion is not.

⁹In this simple example, we can even quantify the two different benefits from delegation. A risk-neutral boss will pay the agent up to pc , so this is the value of getting the optimal state-contingent plan. Now suppose the boss is risk-averse, so that he is willing to pay up to a fee of f , which exceeds pc . The difference, $f - pc$, is the amount attributable to risk aversion.

Signaling Models of Delegation

Signaling models tend to be complex, but much of their logic can be explained rather simply. The signaling model presented here is stripped down in order to convey key properties as clearly as possible.

Consider a legislature that is deciding how much money to spend on defense. There are three possible levels: low, medium, and high, denoted $\{L, M, H\}$. The issue is one-dimensional, so the floor median (FM) is effectively the boss. The legislature has delegated information-gathering authority to an agent, which may be a legislative committee or an executive agency. Here we consider the agent to be a committee (C). (This model could also be used to analyze delegation between a president and an agency.) The agent is authorized to find out the true level of foreign threat and to report this to the legislature. The foreign threat can also be low, medium, or high ($\{l, m, h\}$).

The FM prefers to match spending to threat level, and it has single-peaked preferences. Thus, if we use $F_{J,k}$ to denote the payoff to the floor, given a spending level of J and a threat of k , we assume that (1) $F_{L,l} > F_{M,l} > F_{H,l}$, (2) $F_{L,m} < F_{M,m} > F_{H,m}$, and (3) $F_{L,h} < F_{M,h} < F_{H,h}$.

The game unfolds as follows. First the agent learns the level of threat and sends a message to the boss. Although the boss can require that the agent submit a report, she cannot enforce either truth-telling or precision. Thus, the agent may misreport his private information, e.g. report that the threat is high when it is medium, or he may send a true but imprecise message, e.g. that the threat is not low (when it is in fact medium). If the boss believes that the agent's message contains useful information, she will modify her beliefs about the threat. (For simplicity, assume that initially she believed that each state could occur with probability one third.)

Whether and how she modifies her beliefs is at the heart of the matter, and so we must analyze this step in detail. Because this is a game-theoretic analysis, revision of beliefs must satisfy several rationality-and-equilibrium requirements. First, her beliefs must be realized in equilibrium; she cannot be systematically fooled by the agent's report. (If the agent sends an imprecise message, then the boss may rationally guess wrong about the threat.) Thus, in equilibrium she must extract all that can be inferred from the message, by using her knowledge about the agent's preferences and the structure of the situation. (Suppose, for example, that the agent has a dominant strategy of saying that the threat is high, regardless of the real state of the world. Then in equilibrium the boss realizes that the report contains no information.) Furthermore, a fully rational boss must revise beliefs according to Bayes' law. Finally, since the agent knows that the boss is rational and hence can anticipate how she changes her beliefs (and ultimately uses them), it must be optimal for him to send the postulated report.

Let us see how all this works out by considering three possible types of agents, identified by how close their preferences are to the FM's.

1. Ally. An agent that is an ally of the FM shares preference properties 1–3. Note that this does not mean that the two have identical preferences.

For example, the FM may have no defense bases in her district (and so $F_{L,l} > F_{M,m} > F_{H,h}$), whereas the committee does (so $C_{M,m} > C_{L,l} > C_{H,h}$). What is crucial is that both want to match spending to threat level.

2. Moderate outlier. A moderate outlier prefers medium spending if the threat is medium or low, so $C_{M,l} > C_{L,l}$. Otherwise, he shares the rest of 1–3 with the FM.
3. Extreme outlier. This agent is a “high demander” for defense spending and prefers H over M and M over L , in all states of the world.

What will happen if the agent is an ally? Clearly, there is an equilibrium in which (a) he reports exactly what he observed, (b) the FM modifies her beliefs by taking the report at face value, and then (c) she picks the spending level that matches the reported threat. Given the boss’s response, it is clearly optimal for the agent to tell the precise truth. He too wants the spending to match the threat, and (given the boss’s strategy) he can ensure this by reporting accurately.¹⁰ [This is called a (fully) separating equilibrium, because different messages correspond to different states of nature.]

What if the agent is a moderate outlier? The fully separating outcome cannot be sustained as an equilibrium with this type of agent. Here’s why. Suppose that the boss responded, as postulated, by completely believing the agent’s report. Anticipating this, the agent would never report a low threat; instead, he would say “medium” if the true threat was either medium or low. But then the boss would not rationally take a report of “medium threat” at face value; instead, she would infer from such a message only that the threat is not high. Thus, if the agent is a moderate outlier, then reports cannot be completely informative; hence, information can be lost if an agent who is an ally is replaced by a moderate outlier.

However, there is also a partially separating equilibrium, in which (a) the agent tells part of the truth, reporting high threats accurately and in the other situations saying that the threat is “not high,” (b) the boss revises her beliefs in accord with the message,¹¹ and (c) chooses H if she believes that the threat is h and picks the budget that maximizes her expected utility if she believes that the threat is either medium or low, with equal likelihood. Each player’s strategy is optimal given the other’s, so this outcome is an equilibrium. Hence, with a moderately outlying agent, some information can be transmitted.

But it is pointless to delegate information gathering to an agent who is an extreme high demander for defense spending. Such a subordinate will never send

¹⁰Careful readers will note that we said that there is *an* equilibrium in which these good things happen. Signaling games always have multiple equilibria, and in particular always have “uninformative” equilibria in which the agent sends a message that contains no information and the boss ignores the report. Our exposition follows the literature’s (underjustified) convention of focusing on the best—most informative—equilibria.

¹¹If the report says “not high,” then she believes that the chance of a low threat is one half, as is the chance of a medium threat, while the chance of a high threat is zero.

a meaningful message. Realizing this, the FM ignores the report. (If she did not disregard it fully, then the agent would exploit that gullibility by invariably reporting “high.” But then the message would, in fact, contain no information, and a rational decision maker does not change beliefs in response to worthless signals.) Thus, the only outcomes involve “pooling equilibria,” so-called because in equilibrium the report collapses (pools) all possible states of the world into one signal.

Thus, in this simple signaling model, a rational boss would always prefer allies over moderate outliers and moderate outliers over extreme ones. With allies she can (in some equilibria) obtain the whole truth. With moderate outliers she can (in some equilibria) get some of the facts. But from extreme outliers she learns nothing. (This is a common finding in signaling models. Roughly speaking, the closer the preferences of agent and principal, the more information communicated in the “best” equilibrium.)

Extensions of the Basic Model

This basic model (or more precisely, its already more complicated counterpart in the literature) has been extended in several directions. We consider these extensions serially, as one-at-a-time modifications of the basic model.

Selecting Agents Models of cheap talk suggest that principals do better with agents whose preferences are similar to their own. This naturally led scholars to think about models that endogenized the agent. In a first stage, the boss faced a set of feasible agents and selected one who would be delegated authority to gather information. If principals can select agents, then we get the following prediction about legislative organization: “Legislative committees will *not*, as a matter of practice, be composed predominantly of high demanders or preference outliers” (Krehbiel 1991:95; original emphasis). See, however, Krishna & Morgan (2001) for a modification of that conclusion.

Costly Information Gathering Another natural extension of the basic model is to assume that the agent must pay a cost $c \geq 0$ in order to acquire information. In the simplest versions of this extension, the boss knows whether the agent pays the cost (and so becomes informed) or remains ignorant. In the latter case, the agent is as uninformed as the boss, and so she knows that the report is worthless and therefore ignores it. Hence, in these simple models, the boss listens only to agents who have paid “the price of admission” by incurring the cost of becoming specialists.

One can also turn this point around. Suppose the agent is an extreme outlier. (This could happen in the executive branch if civil service laws protected career officials, even ones with unusual preferences.) Because such an agent will be ignored, it is pointless for him to pay the information-gathering cost. Only allies and moderate outliers will do so (provided, of course, that c is not too high).

Multiple (and Heterogeneous) Agents It is obvious from the defense spending example that having multiple agents will not help the boss if they have similar

preferences which differ greatly from hers. A group of extreme outliers will all report that the threat is high, whatever they have observed. Thus, signaling models direct us to consider the wisdom of delegating to multiple agents with heterogeneous preferences (Gilligan & Krehbiel 1989a, Krishna & Morgan 2001). This comports with informal conventional wisdom—it is generally considered prudent to get information from diverse sources—but the signaling models explain in detail how and why this works.

Krishna & Morgan (2001) show that if the committee members' ideal points bracket the FM's (in the one-dimensional policy space), then the FM can devise a clever scheme for playing the two agents off against each other so that all information is revealed in equilibrium under the open rule. (This is desirable for the floor because the open rule gives the committee no proposal power; the floor can pass any bill it pleases.) Thus, they show that employing multiple diverse agents can get the boss just what she wants in a signaling game.

Proposing Options Thus far, the signaling models have concerned situations in which the agent was delegated only a limited authority, to gather information and to report the findings.¹² Proposal power was off-limits. However, as long as the boss's choice set remains unconstrained, adding proposal power to the list of what can be delegated changes little in the formal models. For instance, in several of the Gilligan-Krehbiel models, if the committee proposes a bill and the open rule obtains, this is equivalent to the committee sending the floor a report on its private information, since the bill allows the FM to infer what the committee's underlying private information must be.

However, once there are constraints on the FM's choice set, matters change. The literature has examined two types of constraints, both motivated by legislative procedures. In the closed rule, the boss must choose between the subordinate's proposal and the status quo. In the modified open rule, she gets proposals from two agents and can pick either.

Why would the FM voluntarily constrain her choice set? In an important series of papers, Gilligan & Krehbiel (1987, 1989a,b, 1990) give an intriguing answer: By limiting her options, she cedes influence over the final decision to her subordinate. This seems bizarre—why would a principal give up power to an agent?—but Gilligan & Krehbiel explain why it makes sense. Giving the subordinate a stake in the outcome gives him incentives to improve that outcome. Thus, committees that operate under restrictive rules have more incentive to uncover policy-relevant information. The ensuing reduction of the outcome's variance benefits everyone, since all actors are risk-averse. Formal analysis of homogeneous committees bears out these claims. [See Krishna & Morgan (2001) for some qualifications about heterogeneous committees.]

¹²Students of political institutions will recognize, however, that such authority is not trivial. For example, environmental analysts in the Corps of Engineers often had to lobby hard to obtain the authority to investigate suspected environmental impacts of proposed Army Corps projects (Taylor 1984).

From Signaling to Screening

Baron has criticized the assumption of signaling models that the principal will passively await the agent's message. "To the extent [the legislature] has instruments at its disposal to structure the work of the committee, the legislature gains by moving first rather than waiting for the committee to act" (2000:485). He constructs a model in which the floor precommits to providing resources to the committee; the level of resources depends on the nature of the report the agent will send to the legislature. Thus, the boss is committing to a schedule of rewards that depends on the agent's subsequent behavior. Because this incentive system screens the information the agent subsequently transmits, these are called screening models in game theory.

Although many of the results of signaling models stand up in Baron's screening model—for example, the more divergent the players' preferences, the less information transmitted by the agent—there are also many implications that differ (for a clear summary of the overlap and the differences, see Baron 2000:502). Perhaps most importantly, the parent body is strictly better off screening the agent's reports. This conclusion poses a problem for the signaling approach: Given this potential benefit, why would the parent body let the committee move first (Baron 2000:503)?

Screening by Requiring Lengthy Justifications

A natural way to segue from a signaling model into a screening one is to assume that the boss selectively imposes different costs on different signals (e.g. the higher the subordinate's budget request, the longer the required justification). Reconsider the defense spending example. Suppose that the boss is stuck with a moderate outlier for an agent, as defined above. She knows that if she responds appropriately then the agent will truthfully report a high threat, but his message will not distinguish between low and medium threats. What can she do? Under certain circumstances, she can selectively impose signaling costs on the subordinate so that it will be optimal for him to signal m if and only if the threat truly is medium. [The requirement is that the agent is more interested in getting medium appropriations when the threat really is medium than when the threat is low: $V(M, m) - V(L, m) > V(M, l) - V(L, l)$.] When this condition holds, the boss requires the agent to send a lengthy justification for a signal of m . (If the boss gets a message of l , without justification, she believes it.) Preparing such a report takes time, which is costly, and the boss's requirement is sufficiently laborious so as to drive a wedge between sending that message under the two possible states of the world: $V(M, m) - V(L, m) - c > 0$ but $V(M, l) - V(L, l) - c < 0$, where c is the time cost. Then it is optimal for the agent to write a lengthy document if and only if he observes a medium threat. Hence, by screening different types of agents (i.e. agents who observe low threats versus those who observe medium ones), the boss has ensured that delegating information gathering and communication to an agent can result in a fully informative equilibrium, despite the players' (moderate) difference in preferences.

Combining Signaling and the Delegation of Authority

Almost all the work on delegation is one of the above two types. But in a path-breaking work, Epstein & O'Halloran (1999) analyze the behavior of a boss who first receives a report from one agent and then decides whether to delegate authority to another agent—an empirically plausible sequence. (Epstein & O'Halloran's empirical setting is legislative-executive relations, so they interpret the boss as the FM, the reporting agent as a congressional committee, and the implementing agent as a bureau, but this setup applies to intra-executive structures as well. Imagine a boss, a staffer writing a report, and a "line" subordinate responsible for implementation.)

As usual, of course, realism comes at a price—analyzing the model is rather complicated—but Epstein & O'Halloran (1999) show that the presence of multiple agents with different responsibilities leads to interesting new conclusions. Their most interesting results, relative to what one can find using the simpler, dyadic principal-agent models, all have a coalitional flavor. They arise from how (for example) the presence of the committee affects the boss's stance toward the agency, or how the agency's presence affects the committee's posture toward the FM. Let us look at two of these results. First, the ally principle suggests that, in general, the further the agency is from the boss, the less likely the latter is to delegate. Epstein & O'Halloran derive a significant variant of this principle: The closer the committee is to the boss, the less likely the latter is to delegate authority to the agency [1999:68 (Figure 4.3)]. (Students of bureaucracy will recognize this as a source of conflict between staff and line.) Why? Recall that the value of delegating to an informed subordinate is that he can condition his action on information unavailable to the boss. But in the context of signaling games, the ally principle says that in general, the closer are principal and agent, the more information is communicated in the latter's report, in equilibrium. Hence, the boss becomes more informed via the signaling game, whence the value of delegating discretion to the agency falls (or remains constant).

Second, consider what can happen when the committee and the agency are ideologically close, as happens in some "iron triangles." The committee realizes that the more information it gives the floor, the less authority will be delegated to the agency. Although this is fine if the committee and floor are allies, it may be undesirable (for the committee) if the committee is closer to the agency than it is to the floor. Thus, Epstein & O'Halloran (1999:70) find that the presence of an ideologically compatible agency may induce a committee to transmit less information to its parent body. (This holds even if getting information is free; hence, the finding does not turn on a collective action problem.)

These findings reveal that specialized attention to legislative institutions, on the one hand, and executive structures, on the other, can blind us to important interactions across organizational boundaries: "committee decision making, and legislative organization more generally, should not be studied in a vacuum; they exist in the shadow of delegation to the executive, and the committee's anticipations about the

anticipated delegation regime will influence committee members' actions earlier on" (Epstein & O'Halloran 1999:71).

Volden (2000) raises two questions about the Epstein & O'Halloran model and crafts his own model in response. First, he points out that although their model is supposed to be a theory of the major institutions that make national policy, in fact it does not match the federal institutions very well. For example, the President cannot veto bills passed by Congress. (Hence, in some ways, the Epstein & O'Halloran model better represents a unicameral parliamentary system than one with separation of powers.) Second, Volden notes that although the model focuses on essentially new agencies, most major legislation involves bureaus that already exist and so enjoy some status quo level of discretion. [We would add that although their theory focuses on new agencies, with a clean delegation slate, most of their data (Appendix B) pertain to preexisting agencies that already have some discretion.]

Volden thus modifies the Epstein & O'Halloran model in several ways: The President has a veto, and agencies may have some preexisting level of discretion. (He also omits an integral part of their model, the legislative committee and its signaling to the floor.) These changes yield results that differ in some ways from those of Epstein & O'Halloran. For example, suppose the bureau chief's ideal point is in between the President's and the FM's. Under some conditions, the legislature will want to reduce the agency's discretion, but the President opposes this move. If he has a veto, he may be able to block the change; without a veto, he could not.¹³ Moreover, limiting the agency's discretion is possible only when the bureau chief's preferences are extreme relative to those of both the legislature and the President, and even in this case, such limitations are more likely when the political leaders' preferences are similar. Volden (2000:12) remarks about this last result, "This finding of a more-unified government as the most-likely condition for the limitation of discretion stands in sharp contrast to earlier work that suggested that such limitations are more likely to arise under divided government."

DELEGATION IN REPEATED CONTEXTS

In all of the preceding models and many of the verbal theories, the delegation problem was fundamentally one-shot. (In many cases, the game unfolded in stages, but the boss made only one delegation decision.) This is a sensible starting point for formal analysis, but in the real world, many delegation issues recur. Hence, it was a salutary development when scholars started to construct models of repeated delegation (Diermeier 1995).

¹³Epstein & O'Halloran (1999:59) justify the omission of presidential vetoes by arguing that although inclusion of a veto would lead to higher equilibrium levels of discretion, the model's comparative statics would be unchanged (1999:59).

Before we turn to specific models, a few overall remarks about repeated-game theory and delegation are in order. In general, repeated interaction yields the possibility of stick-and-carrot strategies of conditional cooperation. For delegation, this means that both the boss and the subordinate are more able to deter each other from cheating. For example, if the subordinate cheats (say by exploiting the discretion given to him), then the boss might retaliate by seizing control in the next period. Or if the boss cheats (say by grabbing credit), then the subordinate could retaliate by shirking in the next period.

Further, under certain parametric conditions, these stick-and-carrot strategies are individually rational and so form Nash equilibria. (Moreover, they could involve only credible punishment threats and so be subgame perfect.) Thus, repetition may reduce the harsh separation between what is individually rational and what is collectively optimal. To see how, consider the toughest delegation circumstance identified by our typology, when in the one-shot game both parties have dominant strategies to cheat on the benign arrangement of (delegate, work). Thus the Pareto-inferior outcome of (control, shirk) is supported by the rock-solid foundation of dominant strategies. It is, in short, the notorious Prisoner's Dilemma.

However, if the game is repeated, then the strategic repertoire expands greatly, which may permit cooperation to be stabilized. Suppose, for instance, that each side uses the role-appropriate definition of Tit for Tat. The subordinate will work hard in his first encounter with his boss. Thereafter, in any period $t > 1$, he will work hard if and only if in $t - 1$ the boss delegated. Similarly, if the boss plays Tit for Tat, then she would begin the relationship by delegating. Subsequently, she would delegate in period $t > 1$ if and only if the subordinate had worked hard in $t - 1$. As is well-known (e.g. Axelrod 1984), two Tit for Tat strategies form a Nash equilibrium in the Prisoner's Dilemma if and only if the future is sufficiently important for both players.¹⁴ This implies that an efficient institutional arrangement can be strategically stable. (Obviously, when this deal can stand up in the most difficult delegation case, then it also works in less trying circumstances, e.g. when only one party has a one-shot dominant strategy of cheating.)

An important caveat: Although the "cooperative" outcome of (delegate, work) becomes individually rational once repetition is added to the picture, many other outcomes—including the bad one of (control, shirk)—are also equilibria. This fact, known in game theory as the folk theorem, implies that the good outcome is not guaranteed. Repeat play merely makes it a live option, in contrast to its dead-on-arrival status in the one-shot game.

The preceding was a "free" application of the theory of repeated games to delegation issues. We used standard results from the literature, which, though not specialized to the study of delegation, apply directly. The models we now examine are tailored more specifically for delegation problems.

¹⁴This holds whether or not the game is symmetric, which matters because the delegation game is obviously asymmetric—the players have different roles and action sets.

Delegation with Overlapping Generations

In the real world, people retire or die, so no one plays an infinitely (or even indefinitely) repeated game. Thus, the assumption of the standard repeated-game models that if a player reaches period t , the probability that he will reach period $t + 1$ is a fixed probability is unwarranted. What remains, then, of the implications of the standard models, of (for example) the idea that delegation and decentralized effort may be maintained by longterm reciprocity between principal and agent? To answer the general question, game theorists have constructed models of organizations populated by overlapping generations of finitely lived members. The idea is elegantly simple.

For example, suppose that each generation or cohort in an institution lives seven periods. A new generation enters every period. Everyone in a particular cohort enters and leaves the institution at the same time. Every generation has n members. Thus, in the steady state, the organization has $7n$ members. Suppose that in every period everyone has a choice of either providing a collective good, to be shared equally by all organization members, or not. Providing the collective good imposes a private cost of $c > 0$. The game is an N -person Prisoner's Dilemma. In the one-shot game, each person has a strictly dominant strategy of shirking, so the only Nash outcome is that everyone shirks, but the all-work outcome is Pareto-superior to the all-shirk outcome. Further, if there were only one cohort of (finitely lived) players, or if there were many but they did not overlap, then again individual rationality would single out the dismal outcome as the unique prediction.

Even if the generations overlap, nothing can be done about players in their last period. Because they are in End Game, they are certain to shirk. Thus, full cooperation is unattainable. Consider, however, the following idea. Since "old" members are going to shirk no matter what, the proposed semicooperative equilibrium will allow for that, without punishment. The young and middle-aged (those in their first five periods in the organization) are expected to provide collective goods, and anyone shirking in these cohorts is punished. The key is that the punishment is delivered not only by all "workers" who are in the organization when the violation occurs; newcomers also punish the deviant by shirking after they arrive. For example, someone who shirks in his fifth period in the organization is punished by being denied retirement pay from the young and middle-aged cohorts who normally would have worked in his sixth and seventh periods. Assuming that the standard payoff conditions are met—the one-period temptation of switching to shirking is outweighed by the subsequent loss of cooperation—it is individually rational to work in periods one through five and retire in periods six and seven.

Such arrangements can work only if certain temporal and demographic conditions hold. In particular, people must live long enough, and cohorts must overlap sufficiently, so that retirement benefits supplied by current workers are attractive enough to induce effort from the young (Cremer 1986). But when these conditions do hold, finitely lived agents (in indefinitely enduring communities) can be induced

to cooperate much of the time. Thus, models of overlapping generations drive a wedge between the lifespan of people and the lifespan of their institutions, and so (in Cremer's words) reveal the significance—indeed, the necessity—of institutions in stabilizing good outcomes.

Diermeier (1995) puts this idea to work in the complex setting of a majority-rule legislature. The relations between the committee and the floor in the stage game are (by design) canonical, though for simplicity he does not explain how information is transmitted from committee to floor. Instead, if the committee specializes, then all legislators are automatically informed about the random shock's realization. (By finessing the complexities of signaling models, this conserves analytical resources for the repeated-game aspects of the theory.) There is a fixed set of "teams" whose members have identical preferences. People live for a known number of periods; after they die, they are replaced by someone from the same team.

Diermeier's main result is that, given the usual repeated-game provisos (the future must be sufficiently important, the payoff from reneging must not be too big, etc), there exists a benign Nash equilibrium wherein the committee specializes and the floor defers to the committee's bill. However, because the players have finite lifespans, the standard proviso about the future's importance goes beyond the usual requirement on discount factors. In addition, the institution's demography must have certain properties. For example, generational overlap cannot be too short, for this would preclude the required transfer of benefits across cohorts. (In the extreme, if the teams turned over completely in one period then floor-committee cooperation would collapse.) Further, random shocks to the demography—e.g. a sudden increase in new members from electorally insecure districts (Diermeier 1995:351)—could destabilize the delegation equilibrium. Thus, an overlapping-generations model of delegation generates new and interesting predictions.

More generally, models of repeat play show that the floor need not tie its own hands procedurally to induce the committee to specialize. A tacit understanding, upheld by strategies of conditional cooperation, can sustain the (delegate, work) outcome.

Delegation and the Search for Policies

In the standard formulation, policy uncertainty is represented by a random shock that perturbs programs. But this is only one type of relevant uncertainty that affects governmental policies in the real world. Often, all decision makers—principals and agents alike—are unsure which kind of program is preferable. Do charter schools outperform conventional public schools? Is missile system x better than system y ?

To represent such problems, search-theoretic models, typically set in a repeated-game framework, seem a natural choice. Delegation occurs in these contexts because the boss lacks the time or expertise to carry out search, so the boss's basic choice is similar to that of the standard formulation. She can either delegate and

later make a more informed choice—possibly biased by the subordinate’s goals—or choose under ignorance.

Ting (1999:Ch. 3) has constructed a search-theoretic model of delegation. In his setup, decision makers are initially uncertain about the distributions of outcomes produced by different policies. By experimenting with various options in different periods, they can learn how these distributions compare. In decision theory, these are called multi-armed bandit problems. (The name evokes the image of people playing slot machines.) Multi-armed bandit models are usually cast as single-person optimization problems. Ting’s setting is strategic; the boss can delegate search authority to more than one agent, and the agents’ preferences may differ from each other and from the boss’s.

Search-theoretic models lend themselves naturally to investigating issues that arise from the interactions of multiple agents searching in parallel. (In contrast, in the standard “outcomes = policies + shocks” model, the boss typically needs only one agent to observe the perturbation and then choose an appropriate policy.) Several of Ting’s results pertain to these interactions.¹⁵

1. Suppose the boss delegates. Then simply adding agents to engage in parallel search does not ensure that there will be more experimentation. The reason is that some agents may free-ride on others’ efforts. “[I]nformation is a public good, and rational bureaucrats may shirk by choosing conservative policies and letting others bear the risk of experimentation” (Ting 1999:63).
2. This informational free-riding is more likely if (holding beliefs constant) all agents who search are allies of the boss than if some of them are not allies. Thus, Ting’s results show that when search matters, the ally principle need not hold in its unvarnished form.

CREDIBLE COMMITMENT AND DELEGATION

Thus far, the ally principle—that a boss prefers subordinates who resemble herself ideologically—has stood up well. Although we noted a few caveats, they only weakly challenge the principle. Formal theorizing seems to track informal thinking here. However, the study of delegation has uncovered an important class of situations in which it is optimal for a boss to violate the ally principle. These situations

¹⁵It is well established in the literature on bureaucracy that having multiple, information-providing agents is more reliable than having only one, when subordinates are fallible but nonstrategic (Landau 1969, Bendor 1985, Heimann 1997). Under these conditions, engineering-reliability theory tells us that redundant systems are more reliable than nonredundant ones. But entirely new problems appear when agents can behave strategically, by e.g. engaging in biased search driven by “mission orientations” (Bendor et al 1987) such as the Air Force’s preference for planes over missiles.

are characterized by problems of credible commitment, often between a political leader and his constituents.

Recall the essence of the game in which the problem arises. In the first stage, a leader makes a promise, say by passing a law, to some or all of his constituents, e.g. to avoid expropriating wealth generated by investments (North & Weingast 1989). In the second stage, the constituents can take an action—for example, invest—based on their beliefs about the leader's promise. In the third stage, the leader moves again, for example by faithfully implementing the tax law or by acting outside the law. Everyone would be better off if the leader could commit to following through on his promise to obey the law. But in the game sketched above, it may be suboptimal for the leader to do so. Instead, if the citizens invest in stage two, then he will renege on his promise in stage three and expropriate their profits. Anticipating this doublecross, businessmen refuse to invest in stage two. (More generally, if the model allows for degrees of investment, then investment will be suboptimally low in that stage.)

In this game, constituents can be seen as the boss, and the political leader as the subordinate. Voters who recognize the commitment problem may choose to elect a leader who will avoid expropriation. As Persson & Tabellini (1994) suggest, the majority elects a policy maker who, after the investments are made, favors a lower tax rate than the majority does.

Related historical experience (North & Weingast 1989) bears out this analysis. Under the Stuart monarchy and until the Glorious Revolution in 1688, the British Crown raised money in ways that hurt investment. It reneged on loans, seized private assets, undermined the common law courts by creating its own set of courts to uphold the legality of its actions, and tried to rig parliamentary representation in its favor. Ultimately, private lenders became unwilling to make loans to the Crown at all. Following the Glorious Revolution, courts closely tied to the Crown were abolished, and all cases involving private property had to be tried by common law courts. Common law judges could now be removed only for violations of good behavior. A law requiring the regular meeting of Parliament was passed. Parliament's exclusive authority to raise new taxes was reestablished. Parliament gained the right to audit how the government spent its funds. Royal prerogative powers were limited.

In short order, the new King borrowed far more money and at far lower interest rates than had his predecessors. North & Weingast argue that the central reason is that the Crown, because of the institutional changes made during and after the Glorious Revolution, became credibly committed to honoring its financial obligations. It could no longer renege on these obligations because it had lost its authority to do so—such decisions were now Parliament's. And the diversity of interests in the legislature, as well as the strong presence of the commercially minded Whigs, made Parliament unlikely to renege on the government's debts.

The idea that a boss might delegate in order to make a policy credible has been especially prominent in the literature on central banks and monetary policy.

Delegation and Commitment Problems in Monetary Policy

To examine delegation of monetary policy, we must first briefly describe some of the effects of monetary policy. Although in the short run an increase in the money supply can increase employment and output, workers and firms who realize that demand for labor and for goods is higher may demand higher wages and higher prices, generating inflation. The problem may be even more severe. Firms and workers who anticipate an increase in the money supply may increase wages and prices, thereby reducing demand for goods and labor. Government may then want to stimulate the economy to avoid a recession and may therefore increase the money supply. The anticipated, and realized, increased money supply would then leave the economy with the same output and employment, but with higher prices. If, instead, firms and workers expected little increase in the money supply, they would not increase prices and wages, and the economy would enjoy the same level of output with lower prices.

In the following, we think of Congress or the Executive as the boss delegating monetary policy to a central bank. If the central bank has the same objectives as the boss, then the delegation has no real effects. The interesting question is, can the performance of the economy be improved by delegating to a central bank whose preferences differ from the boss's? The surprising answer is yes.

The credibility problem can be formulated in terms of the objective function of the boss, which is often called the social welfare function. Let social welfare directly decline with inflation (p) and decline with the gap between actual output, y , and an ideal output y^* . (Higher output may sometimes be undesirable, perhaps because it increases the trade deficit, reduces prospects for future growth, or increases inequality.) That is, let the social welfare function be

$$U = -(y - y^*)^2 - \theta p^2, \quad 1.$$

where θ represents the weight on inflation.

Let output increase with the size of unanticipated inflation, which is the difference between actual inflation (p) and expected inflation (p_e).

$$y = p - p_e \quad 2.$$

Substituting Equation 2 into Equation 1 gives

$$U = -(p - p_e - y^*)^2 - \theta p^2. \quad 3.$$

Alternatively, we can view the objective as minimizing the social loss,

$$L = (p - p_e - y^*)^2 + \theta p^2. \quad 4.$$

Suppose first that the boss could credibly commit to an inflation rate, p . From Equation 4 it is clear that the optimal choice is $p = 0$. Since the boss committed to $p = 0$, that is the inflation rate expected by firms, workers, and consumers (in

short, by economic agents). So $p - p_e = 0$, and the social loss under commitment, L^c , is just

$$L^c = (y^*)^2.$$

Suppose next that the boss cannot commit. She instead takes expectations of inflation as given, minimizing Equation 4 given those expectations. The first-order condition for minimizing L is

$$\frac{\partial((p - p_e - y^*)^2 + \theta p^2)}{\partial p} = 0,$$

with the solution

$$p = \frac{p_e - y^*}{1 + \theta}.$$

If we assume that people have rational expectations, so that $p_e = p$, we find that the solution under discretion, called p^d , is

$$p^d = \frac{y^*}{\theta}. \quad 5.$$

Substituting Equation 5 into Equation 4 yields

$$L^d = \frac{(y^*)^2(1 + \theta)}{\theta}.$$

Note that for all positive finite values of θ , $(y^*)^2 < \frac{(y^*)^2(1+\theta)}{\theta}$. Thus, if the boss aims to maximize social welfare but cannot commit to fixing inflation at zero, she will be tempted to stimulate the economy through inflation; firms, workers, and consumers will anticipate this inflation, and social welfare will be lower than it would be with commitment. [This idea was formulated by Barro & Gordon (1983) and Backus & Driffill (1985); Persson (1988) offers a more complete discussion.]

Three proposals have been offered to ameliorate the problem. All involve delegating decisions to a central banker but making his objective differ from the boss's. The first is to delegate control of monetary policy to a "conservative" central banker, one who strongly dislikes inflation, as indicated by a high value of θ (Rogoff 1985). Such reasoning played a role in President Carter's appointment of Paul Volcker as Chairman of the Federal Reserve Bank in 1979 (Greider 1987). The nomination of Volcker, described as "rigidly conservative, very right-wing, and not a team player," was applauded by business and banking leaders.

As we saw, however, no finite value of θ can attain the solution attainable under commitment. Other tools can do better. Consider setting an inflation target, p_b . [The seminal work on this topic is by Svensson (1997). For a discussion of the credibility of inflation targeting, see Cukierman (2000).] The central bank is somehow induced to minimize not Equation 4 but rather to minimize

$$(p - p_e - y^*)^2 + \theta(p - p_b)^2.$$

Assuming again that the boss cannot commit to policy, but that economic agents anticipate the central bank's actions when they forecast inflation, realized inflation will be

$$p^t = \frac{y^* + \theta p_b}{\theta}.$$

To induce an outcome with zero inflation, the boss should set p_b to satisfy

$$\frac{y^* + \theta p_b}{\theta} = 0,$$

obtaining

$$p_b = -\frac{y^*}{\theta}.$$

Realized inflation will be zero, and social loss the same as under commitment. So delegation combined with inflation targeting solves the commitment problem.

The third tool is for the boss to delegate and then to impose a fine on the central banker of f per unit of inflation. This is called an inflation contract. The central banker minimizes

$$L^r = (p - p_e - y^*)^2 + \theta p^2 - fp.$$

With rational expectations by economic agents, realized inflation is

$$p^r = \frac{2y^* + f}{2\theta}.$$

The value of f that would induce zero inflation is $f = -2y^*$. So an inflation contract, like an inflation target, can induce the central banker to choose the level of inflation that maximizes social welfare.

These results are reasonably robust. One important exception is that the model described above assumes that the boss fully knows the central banker's preferences. If there is uncertainty (say about θ), then gaining the benefits from delegation may require the boss to both set an inflation target (or use an inflation contract) and try to choose a central banker who dislikes inflation (Muscatelli 1999).

A serious criticism of these approaches is that they assume that at some level the boss can commit, whereas at others she cannot (McCallum 1995, Jensen 1997). Consider inflation targeting, where the boss assigns the central banker an inflation target, p_b , that enters into the banker's loss function. If a surprise oil shock hurts the economy, why wouldn't the boss who announced an inflation target surreptitiously change it, inducing the banker to increase inflation and thus output? Or why not surprise the public by appointing a central banker who cares little about inflation?

Summary Our discussion of central banking shows that the government may increase welfare by delegating authority to an official whose preferences differ from the government's. We also see that such benefits from "strategic delegation" can apply in many more situations.

These results are interesting mainly because they show that it can be optimal to ignore the ally principle, and they explain why and when this is so. After all, common sense supports the finding that it is often rational for an uninformed principal to delegate authority to an informed agent; it is rather obvious that tapping the agent's greater expertise can provide gains. (Further, when delegation is beneficial for this reason, then the ally principle is typically operative.) But the value of deliberately handing authority over to an agent whose goals differ markedly from one's own—more, that this preference difference is essential to delegation being useful—is much less intuitive. Hence, it is in such contexts that noncooperative game theory and the analysis of subgame perfection reveal their value.

The Difficulties of Credible Commitment

In the provocative opening of one of his papers on “the politics of structural choice,” Moe (1989:267) asserts that “American public bureaucracy is not designed to be effective.” This is an intriguing claim. To argue that an institution turns out to be inept is one thing; to maintain that it was not designed to be effective is quite another. Why would this happen?

In Moe's (informal) theory (Moe 1989, 1990a,b, 1991; Moe & Caldwell 1994; Moe & Wilson 1994), the central problems facing politicians creating a new agency in a democracy are political turnover and political uncertainty. The party in power today can give a new agency statutory authority, but in a democracy, power is ephemeral. Eventually the opponents will run things, and then they might gut the agency or turn it in an unpalatable direction. What to do?

Moe argues that the party currently in power (say, the Democrats) faces two distinct choices. On the one hand, the Democrats could heed the classical reasons for discretion and give the agency much leeway so that it can adapt to changes in its task environment. Doing that, however, simultaneously makes it possible for their opponents to redirect the agency. On the other hand, the Democrats could enmesh the agency in structural and procedural constraints. Doing so would confer a political advantage—these constraints would make it hard for new political superiors to influence the agency in unwanted ways—but it would simultaneously make the agency rigid and thus technically inefficient. Essentially, then, Moe is suggesting that the Democrats may prefer to have a somewhat inefficient, but still environmentally oriented, Environmental Protection Agency today *and tomorrow* over the bundle of an effective, environmentally oriented EPA today and an (effectively) gutted EPA tomorrow.

There is a puzzle, however: Nobody likes inefficient agencies. Thus, the Republicans could suggest the following deal: “Don't burden the agency today with constraints that degrade performance. Let it remain flexible. In return, we promise that when we come into power we won't gut the agency but will only moderate its environmental enthusiasm.” Such a deal may well be mutually beneficial (under plausible assumptions politicians prefer “policy-smoothing” over a volatile path with the same average policy), but there is a hitch: The party in power must do

its part in the present while the out-party is agreeing to deliver in the future. Will it? If the Democrats doubt the word of their opponents, then they may do the relatively safe thing, namely insulate and rigidify the agency. Hence, the Pareto-inferior outcome of inefficient agencies arises precisely because future principals cannot make credible commitments today about how they will run agencies tomorrow. Consequently, Moe's theory can be seen as a pessimistic counterpoint to the studies of central banking. Whereas in those studies, quasi-independent agencies were the solution to problems of credible commitment, in Moe's account they reflect the inability of politicians to credibly commit to stable and efficient policy compromises.

This theory posits that the turnover of political principals is a central problem for the design of bureaucracies, but one could challenge this key premise and its presumed effects. So what if there will be turnover of principals? Today's governing party should rationally create an organization that is best suited to carry out the party's optimal policy; if the out-party takes over, it will do the same. Such is life in a competitive polity. And if the system at hand is a parliamentary one—particularly a parliamentary system with only two major parties, as in Britain—this point has merit (Moe & Caldwell 1994). In such systems, if a single party takes control of Parliament it will be in a position to sweep away the procedural and structural constraints imposed on agencies by previous governments. Thus, agencies in these systems cannot be insulated from the influence of future principals. And since by hypothesis this insulation comes at a price—inefficient rigidity—rational parties in power would not do it in (two-party) parliamentary democracies.

In contrast, in systems with separation of powers, it is unlikely that any one party will have the kind of commanding position enjoyed by (single-party) winners in parliamentary systems; American-type polities have too many veto points. Hence, Moe argues, if the Democrats enmesh the EPA today in constraining laws, tomorrow the Republicans may be unable to overturn this statutory status quo. (This was exactly what President Reagan and EPA head Gorsuch found out in the early 1980s.) Separation of power systems are designed to favor retaining the status quo, which tempts the incumbent party to surround an important agency with stabilizing (yet debilitating) constraints.

Although Moe's theory is plausible, his specific predictions depend on specific assumptions concerning the actors' expectations about future political trends. Moe argues that when the members of the governing coalition believe the future will become politically less favorable for their program, they will be inclined to constrain the agency, so that neither the agency nor future politicians can easily modify the program. Sometimes, however, ruling politicians believe that the tides of political fortunes are in their favor. In this case, they may give the agency discretion so that it can unilaterally adapt to new contingencies without having to engage in time-consuming negotiations with the future Congress and President. This modification, which hinges on the politicians' forecasts about political trends and on their degree of risk aversion, might allow Moe's theory to account for the fact that Congress has at times given agencies great legal discretion.

CONCLUSIONS

The study of delegation is flourishing in political science. It displays unusual degrees of coherence and cumulation. These two properties are related. Scholars working in this area have shown an admirable degree of self-restraint; resisting the temptation to engage in ceaseless but undisciplined innovation, they have built on each other's work to an extent that is rare in the discipline.

This incremental research strategy has facilitated theoretical progress. It greatly enhances comparability across models, making it easier to figure out when new assumptions (e.g. superiors can monitor agents *ex post*) matter—that is, when they affect predicted outcomes. This cumulative work has helped us understand an important piece of conventional wisdom, the ally principle, that is central to the politics of delegation. Now we not only have rigorous models explaining why rational superiors would use the principle but we also better understand when they would not. Perhaps most striking here are the models that analyze problems of credible commitment, in which superiors deliberately delegate to subordinates who have different goals in order to demonstrate that they are serious about certain policies. Explicit theorizing, by generating surprising yet not implausible results, shows a clear “value added” here.

Of course, because all research strategies involve trade-offs, this cumulation and coherence has come at a price. In particular, although the subfield has benefited from using a common set of building blocks in the different models—even down to the quadratic utility functions, uniform noise, and additive policy-shock technology—it is time to stop relying on such a narrow set of assumptions and to start generalizing results. After all, where is it written that policy shocks are uniformly distributed, or preferences quadratic? The field is mature enough to work on theoretical sensitivity testing. What would happen, for example, if signaling models relied only on an assumption that superiors are risk-averse, rather than assuming a particular functional form? Would the qualitative pattern of results about, say, the ally principle stand up if we relaxed the assumption about uniformly distributed noise and instead assumed a much larger class of disturbances? These are not mere technical issues; they go to the heart of what we believe matters substantively. Typically, informal theorizing uses broad notions—e.g. decision makers are risk-averse—rather than specific postulates, such as quadratic preferences. Thus, if we wish to faithfully represent our informal hunches, we should work as much as possible with the broad notions. If we then discover that certain conjectures do not hold under the general premises, we will have learned something.

A second problem is that too little attention is paid to a central empirical issue: What real-world institution or process is being modeled by a particular formulation? We have no quarrel at all with a research strategy of starting with the simplest of models in the simplest of settings and then incrementally making the models more complex (and, one hopes, more realistic). Indeed, as we argued earlier, this incremental model-building strategy is highly desirable. Yet many modelers (present

company included!) have been rather casual about the prima facie plausibility of their models.

A third common problem is that even when the real-world institutions are adequately represented in a model, the decisional context may not be. In particular, models often presume that principals are designing on a blank slate—that they delegate without an ongoing agency or an ongoing program in the background. Yet this background—e.g. whether an agency already exists and already has at least some de facto authority—probably matters. Changing an existing level of discretion is politically harder than giving discretion to an entirely new agency because, for example, multiple actors may have vetos over changes in an agency's status quo level of discretion (Volden 2000:2).

Finally, consider delegation in complex projects, e.g. the D-Day invasion. Given well-known cognitive constraints (Simon 1990), it was impossible for General Eisenhower to explicitly delegate all requisite authority to all the relevant field commanders. Instead, by virtue of the complexity and uncertainty of their tasks, those officers had a great deal of discretion “on the ground” that Eisenhower could not have anticipated. [See Ambrose (1994) for a fascinating description of the improvisation on the beaches; see Lindblom (1959) for a relevant theory.] Yet current models of delegation in political science blithely assume that the principal is perfectly rational, regardless of the complexity of the task facing her.

All of these issues stand as challenges to the ongoing study of delegation in political science. Given the strides that have been made in the past 15 years, we are confident that these challenges will be taken up and overcome.

Visit the Annual Reviews home page at www.AnnualReviews.org

LITERATURE CITED

- Ambrose S. 1994. *D-Day, June 6, 1944: The Climactic Battle of World War II*. New York: Simon & Schuster
- Aranson PH, Gellhorn E, Robinson GO. 1982. A theory of legislative delegation. *Cornell Law Rev.* 68:1–67
- Arnold RD. 1987. Political control of administrative officials. *J. Law Econ. Org.* 3:279–86
- Axelrod R. 1984. *The Evolution of Cooperation*. New York: Basic Books
- Backus D, Driffill J. 1985. Rational expectations and policy credibility following a change in regime. *Rev. Econ. Stud.* 52:211–21
- Baron D. 2000. Legislative organization with informational committees. *Am. J. Polit. Sci.* 44:485–505
- Barro RJ, Gordon DB. 1983. Rules, discretion and reputation in a model of monetary policy. *J. Monetary Econ.* 12:101–21
- Bawn K. 1995. Political control versus expertise: congressional choices about administrative procedures. *Am. Polit. Sci. Rev.* 9:62–73
- Bawn K. 1997. Choosing strategies to control the bureaucracy: statutory constraints, oversight, and the committee system. *J. Law Econ. Org.* 13:101–26
- Bendor J. 1985. *Parallel Systems: Redundancy in Government*. Berkeley: Univ. Calif. Press
- Bendor J, Glazer A, Hammond TH, Meirowitz A. 2000. *Simple spatial models of delegation*. Unpubl. ms., Grad. School Bus., Stanford Univ., Stanford, Calif.

- Bendor J, Taylor S, Van Gaalen R. 1987. Stacking the deck: bureaucratic missions and policy design. *Am. Polit. Sci. Rev.* 81:873–96
- Calvert R, McCubbins M, Weingast BR. 1989. A theory of political control and agency discretion. *Am. J. Polit. Sci.* 33:588–611
- Cremer J. 1986. Cooperation in ongoing organizations. *Q. J. Econ.* 101:33–49
- Cukierman A. 2000. Establishing a reputation for dependability by means of inflation targets. *Econ. Gov.* 1:53–76
- Diermeier D. 1995. Commitment, deference, and legislative institutions. *Am. Polit. Sci. Rev.* 89:344–55
- Epstein D, O'Halloran S. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers*. Cambridge: Cambridge Univ. Press
- Fiorina M. 1982. Legislative choice of regulatory forms: legal process or administrative process. *Pub. Choice* 39:33–66
- Fiorina M. 1986. Legislative uncertainty, legislative control, and the delegation of legislative power. *J. Law Econ. Org.* 3:287–335
- Gilligan TW, Krehbiel K. 1987. Collective decision-making and standing committees: an informational rationale for restrictive amendment procedures. *J. Law Econ. Org.* 3:287–335
- Gilligan TW, Krehbiel K. 1989a. Asymmetric information and legislative rules with a heterogeneous committee. *Am. J. Polit. Sci.* 33:460–90
- Gilligan TW, Krehbiel K. 1989b. Collective choice without procedural commitment. *Models of Strategic Choice in Politics*, ed. P Ordeshook, pp. 295–314. Ann Arbor: Univ. Mich. Press
- Gilligan TW, Krehbiel K. 1990. The organization of informative committees by a rational legislature. *Am. J. Polit. Sci.* 34:531–64
- Greider W. 1987. *Secrets of the Temple: How the Federal Reserve Runs the Country*. New York: Simon & Schuster
- Hammond TH, Knott JH. 1996. Who controls the bureaucracy? Presidential power, congressional dominance, legal constraints, and bureaucratic autonomy in a model of multi-institutional policymaking. *J. Law Econ. Org.* 12:121–68
- Hammond TH, Miller GJ. 1985. A social choice perspective on authority and expertise in bureaucracy. *Am. J. Polit. Sci.* 29:611–38
- Heimann CFL. 1997. *Acceptable Risks: Politics, Policy, and Risky Technologies*. Ann Arbor: Univ. Mich. Press
- Hill JS, Brazier JE. 1991. Constraining administrative decisions: a critical examination of the structure and process hypothesis. *J. Law Econ. Org.* 7:373–400
- Horn MJ, Shepsle KA. 1989. Commentary on 'Administrative arrangements and the political control of agencies': administrative process and organizational form as legislative responses to agency costs. *Virg. Law Rev.* 75:499–508
- Huber JD, Shipan CR. 1999. A comparative theory of statutory control of bureaucrats. Presented at Annu. Meet. Am. Polit. Sci. Assoc., 95th, Atlanta
- Jensen H. 1997. Credibility of optimal monetary delegation. *Am. Econ. Rev.* 87:911–20
- Krehbiel K. 1991. *Information and Legislative Organization*. Ann Arbor: Univ. Mich. Press
- Krishna V, Morgan J. 2001. Asymmetric information and legislative rules: amendments. *Am. Polit. Sci. Rev.* In press
- Landau M. 1969. Redundancy, rationality, and the problem of duplication and overlap. *Pub. Admin. Rev.* 29:346–58
- Lindblom CE. 1959. The science of 'muddling through'. *Publ. Admin. Rev.* 19:79–88
- Mashaw JL. 1994. Improving the environment of agency rulemaking: an essay on management, games, and accountability. *Law Contemp. Prob.* 57:185–257
- McCallum BT. 1995. Two fallacies concerning central bank independence. *Am. Econ. Rev.* 85:207–11

- McCubbins M, Noll RG, Weingast BR. 1987. Administrative procedures as instruments of political control. *J. Law Econ. Org.* 3:243–77
- McCubbins M, Noll RG, Weingast BR. 1989. Structure and process, politics and policy: administrative arrangements and the political control of agencies. *Virg. Law Rev.* 75:431–82
- Miller GJ. 1992. *Managerial Dilemmas: The Political Economy of Hierarchy*. Cambridge: Cambridge Univ. Press
- Miller GJ, Hammond TH. 1994. Why politics is more fundamental than economics: incentive-compatible mechanisms are not credible. *Pub. Choice* 6:5–26
- Moe TM. 1989. The politics of structural choice. *Can the Government Govern?* ed. J Chubb, P Peterson, pp. 267–329. Washington, DC: Brookings
- Moe TM. 1990a. Political institutions: the neglected side of the story. *J. Law Econ. Org.* 6:213–54
- Moe TM. 1990b. The politics of structural choice: toward a theory of public bureaucracy. *Organization Theory: From Chester Barnard to the Present and Beyond*, ed. O Williamson, pp. 116–53. New York: Oxford Univ. Press
- Moe TM. 1991. Politics and the theory of organization. *J. Law Econ. Org.* 7:106–29
- Moe TM, Caldwell M. 1994. The institutional foundations of democratic government: a comparison of presidential and parliamentary systems. *J. Inst. Theor. Econ.* 150:171–95
- Moe TM, Wilson SA. 1994. Presidents and the politics of structure. *Law Contemp. Prob.* 57:1–44
- Muscattelli VA. 1999. Inflation contracts and inflation targets under uncertainty: Why we might need conservative central bankers. *Economica* 66:241–54
- North DC, Weingast BR. 1989. Constitutions and commitment: the evolution of institutions governing public choice in seventeenth-century England. *J. Econ. Hist.* 44:803–32
- Persson T. 1988. Credibility of macroeconomic policy: an introduction and a broad survey. *Eur. Econ. Rev.* 32:519–32
- Persson T, Tabellini G. 1994. Representative democracy and capital taxation. *J. Pub. Econ.* 55:53–70
- Robinson GO. 1989. Commentary on ‘Administrative arrangements and the political control of agencies’: political uses of structure and process. *Virg. Law Rev.* 75:483–98
- Rogoff K. 1985. The optimal degree of commitment to an intermediate monetary target. *Q. J. Econ.* 110:1169–90
- Sen A. 1970. The impossibility of a Paretian liberal. *J. Polit. Econ.* 78:152–57
- Shepsle K. 1979. Institutional arrangements and equilibrium in multidimensional voting models. *Am. J. Polit. Sci.* 23:27–59
- Simon HA. 1990. Invariants of human behavior. *Annu. Rev. Psychol.* 41:1–19
- Svensson LEO. 1997. Optimal inflation targets, ‘conservative’ central banks, and linear inflation contracts. *Am. Econ. Rev.* 87:98–114
- Taylor S. 1984. *Making Bureaucracies Think: The Environmental Impact Statement Strategy of Administrative Reform*. Stanford, CA: Stanford Univ. Press
- Ting M. 1999. *Essays on the positive theory of bureaucratic politics*. PhD thesis. Stanford Univ., Stanford, CA
- Volden C. 2000. The political economy of restrictions on bureaucratic discretion. Unpubl. ms., School Politics & Econ., Claremont Grad. Univ., Claremont, CA



CONTENTS

TRANSNATIONAL POLITICS: Contention and Institutions in International Politics, <i>Sidney Tarrow</i>	1
MECHANISMS IN POLITICAL PROCESSES, <i>Charles Tilly</i>	21
DEMOCRATIZATION AND ECONOMIC REFORM, <i>V. Bunce</i>	43
PSYCHOLOGY AND INTERNATIONAL RELATIONS THEORY, <i>J. M. Goldgeier, P. E. Tetlock</i>	67
POLITICAL TRADITIONS AND POLITICAL CHANGE: The Significance of Postwar Japanese Politics for Political Science, <i>Bradley Richardson, Dennis Patterson</i>	93
RELIGION AND COMPARATIVE POLITICS, <i>Anthony Gill</i>	117
TOWARD A FOURTH GENERATION OF REVOLUTIONARY THEORY, <i>Jack A. Goldstone</i>	139
POLITICAL CONSEQUENCES OF MINORITY GROUP FORMATION, <i>M. Hechter, D. Okamoto</i>	189
POLITICAL KNOWLEDGE, POLITICAL ENGAGEMENT, AND CIVIC EDUCATION, <i>William A. Galston</i>	217
THEORIES OF DELEGATION, <i>J. Bendor, A. Glazer, T. Hammond</i>	235
TIME-SERIES-CROSS-SECTION DATA: What Have We Learned in the Past Few Years?, <i>Nathaniel Beck</i>	271
THE ORIGINS AND WAR PRONENESS OF INTERSTATE RIVALRIES, <i>John Vasquez, Christopher S. Leskiw</i>	295
THE POLITICAL ECONOMY OF INTERNATIONAL MONETARY RELATIONS, <i>J. Lawrence Broz, Jeffrey A. Frieden</i>	317
BIOLOGY AND POLITICS: Linking Nature and Nurture, <i>R.D. Masters</i>	345
WOMEN'S MOVEMENTS AT CENTURY'S END: Excavation and Advances in Political Science, <i>Karen Beckwith</i>	371
TAKING STOCK: The Constructivist Research Program in International Relations and Comparative Politics, <i>Martha Finnemore, Kathryn Sikkink</i>	391
VIACRATIC AMERICA: Plessy on Foot v. Brown on Wheels, <i>Douglas W. Rae</i>	417