

Spatial Models of Delegation

JONATHAN BENDOR *Stanford University*

ADAM MEIROWITZ *Princeton University*

Although a large literature on delegation exists, few models have pushed beyond a core set of canonical assumptions. This approach may be justified on grounds of tractability, but the failure to grasp the significance of different assumptions and push beyond specific models has limited our understanding of the incentives for delegation. Consequently, the justifications for delegation that have received recent scrutiny and testing differ from some of the more plausible justifications offered by informal studies of delegation. We show that surprisingly few results in the literature hinge on risk aversion, and surprisingly many turn on the ignored, though equally canonical, technological assumption that uncertainty is fixed (relative to policies). Relaxing the key assumptions about dimensionality and functional forms provides a clearer intuition about delegation—one that is closer to classical treatments. The theory allows us to relate different institutional features (commitment, specialization costs, monitoring, multiple principals) to delegation's observable properties.

Although the practice of delegation dates back to antiquity—the Book of Exodus refers to it—only in the last century has much scholarship examined its causes and effects, and only recently has research focused on developing and testing formal theories that explore the logic of delegation.¹ Despite this recent surge in formal work on delegation, scholars have failed to generalize several important properties of the early pathbreaking models. Instead of revisiting assumptions that were initially imposed for the sake of convenience, recent papers tend to maintain and reinforce the canonical assumptions, applying these limited models to ever-richer institutional settings. Although the breadth of understanding about delegation across various settings is impressive, this approach has come at a price. It has produced an understanding that only partly captures delegation's informational rationale in even the simplest contexts. The literature is rich in predictions about delegation in complicated institutional settings when stylized foundational assumptions are satisfied but poor in explanations for and intuitions about why a poorly informed boss chooses to delegate to subordinates with different preferences. Moreover, we lack an understanding of what basic features drive the boss's choice of which agent is given control of policy.

For example, scholars have provided detailed explanations about how separation of power systems affect

delegation by legislatures to agencies, but the basic question “Does the informational rationale for delegation hinge on the boss's risk preferences?” has been ignored. Despite the analysis of complicated signaling and screening games between bosses with ex post controls and their subordinates, we do not yet understand when or why the ally principle (the boss picks the most ideologically similar agent as delegatee) fails in *simple* settings.

We try to put models of delegation on more solid footing by examining a larger (more general) set of environments. We eschew the canonical assumptions of quadratic utility (or even risk aversion), unidimensional outcome spaces, and additive and homogeneous random shocks to policy. Our analysis focuses first on two basic questions: When does delegation occur? and To whom does the boss delegate? We then consider richer, institutionally driven extensions, e.g., the role of commitment and monitoring devices, the impact of multiple bosses, and the institution in which their preferences are aggregated. Through a sequence of nested models with varying levels of abstraction, the analysis isolates the relations between various assumptions and different findings.

To give readers a sense of the theory's fertility, we summarize four implications now.² (1) As the uncertainty between policies and outcomes rises a principal becomes more prone to delegate, *regardless* of her risk preferences. This holds for multiple principals (e.g., in a separation of powers system) if uncertainty increases “sufficiently.” (2) In political hierarchies with a single principal and multiple (informed) agents, if the latter can precommit to delivering particular outcomes, then they will compete away their informational advantages: In equilibrium the boss will get her ideal point (regardless of the policy space's dimensionality). Hence in these circumstances the agents' preferences do not matter: As in certain kinds of Wittmanesque electoral competition, conflict is resolved completely in favor of the principal. (3) A classical feature of delegation, the

Jonathan Bendor is Walter and Elise Haas Professor of Political Economics and Organizations, Graduate School of Business, Stanford University, Stanford CA 94305 (bendor.jonathan@gsb.stanford.edu). Adam Meirowitz is Assistant Professor, Department of Politics, Corwin Hall, Princeton University, Princeton, NJ 08544 (ameirowi@princeton.edu).

Ami Glazer and Tom Hammond provided invaluable guidance and assistance. Their contributions far exceeded the collegial norm. In addition, we have benefited from seminar participants at Stanford University, the 2002 APSA and 2001 MPSA meetings, and a reading group at Princeton University. We appreciate comments from Brandice Caines-Wrone, Alex Cukierman, Sven Feldmann, Mark Fey, James Fowler, John Huber, David Lewis, Nolan McCarty, Charles Shipan, and Joel Sobel, as well as three reviewers and the editor of the *Review*.

¹ See Bendor, Glazer, and Hammond 2001 for a review of models of delegation.

² These summaries omit important details. The implications are stated more precisely later.

ally principle, will not hold if obtaining information is costly and agents cannot precommit to doing so. In such situations the boss and her best ally face a collective action problem, which the principal solves by delegating authority to an agent whose policy preferences are an “intermediate” distance from hers. (4) If the boss can select a discretion level and monitor compliance, the optimal level of discretion may not depend monotonically on the agent’s ideological distance from the boss.³

Before turning to these implications, however, we must first examine the “big picture”—the informational rationale for delegation. Though the canonical model’s diffusion seems to indicate a scholarly consensus about this rationale, this is not really so: The traditional perspective on delegation provides a different view.

The classical analysis of delegation offers a simple and persuasive perspective. This view was elegantly expressed by Alexander Hamilton ([1788]. 1961, 199–200) in *The Federalist* No. 23’s examination of defense policy and management:

The authorities essential to the common defense are these: to raise armies; to build and equip fleets; to prescribe rules for the government of both; to direct their operations; to provide for their support. These powers ought to exist without limitation, *because it is impossible to foresee or to define the extent and variety of national exigencies, and the correspondent extent and variety of the means which may be necessary to satisfy them.* The circumstances that endanger the safety of nations are infinite, and for this reason no constitutional shackles can wisely be imposed on the power to which the care of it is committed. (italics in the original).

Hamilton thus advanced a logic for transferring control from less informed to more informed officials, anticipating that the more knowledgeable ones would make better choices. His view has persisted to the present day and is particularly prominent in the study of bureaucracy. Indeed, that delegating to experts or specialists can yield superior outcomes has been conventional wisdom in public administration for most of the last century.⁴ Translated into the language of modern economics, the traditional perspective says that optimal plans are usually *state-contingent*. Hence, if a principal is ignorant of the state of the world, then it may be optimal to delegate to an agent who *does* know the true state, provided that the costs of delegating (financial, political) are not too high.

In contrast, in modern formal models of delegation, making outcomes stochastically better is *not* the

key to delegation. Instead, the focus is on how delegation can make outcomes *less risky* (e.g., Gilligan and Krehbiel 1989, 462; 1990, 536–37; Epstein 1997, 277; 1999, 6; Epstein and O’Halloran 1999, 235).⁵

This focus on risk reduction has had significant implications for other features of current models of delegation, thereby affecting our current understanding of the phenomenon. In particular, because lowering risk is unequivocally desirable if *and only if* a principal is risk-averse, most current formal models assume that political principals are risk averse. (For an exception see Huber and Shipan 2002.) Indeed, they typically assume a specific *degree* of risk aversion, by positing a quadratic utility function. For example, in discussing the assumptions of the canonical model of delegation, Epstein and O’Halloran (1999) write:

This quadratic loss function has two important implications. First, actors will prefer outcomes that are closer to their ideal point than those further away; this is the distributive component of their utility. But actors will also be risk-averse, meaning that they will dislike uncertainty over policy outcomes; *this is the informational component.* Even though actors want outcomes close to their preferred policy, then, they may nonetheless be willing at the margin to accept a system that biases policy away from their preferences if they can simultaneously reduce the uncertainty associated with these outcomes. (54; emphasis added)

The contrasts between the traditional and the modern perspectives are clear and striking. The former asserts that the informational basis for delegation is very simple: When optimal policies are state-contingent and the costs of delegating not too high, *any* rational, uninformed principal prefers to transfer authority to an informed agent. The latter asserts that the informational basis for delegation is based on risk reduction. Recall Epstein and O’Halloran’s (1999) explanation: “Actors will also be risk-averse, meaning that they will dislike uncertainty over policy outcomes” (54). Under this perspective, delegation depends intimately on the risk preferences of principals.

Substantively, this analysis comes down on the traditional side. We show that what drives the decision to delegate in spatial models is not aversion to risky outcomes but simply an aversion to bad—i.e., distant—ones. (More precisely, the quasi-concavity inherent in Euclidean preferences by itself creates an informational rationale for delegation; concavity of the loss function is inessential.)⁶

As noted, to ensure tractability the canonical delegation model typically uses three other specialized

³ To avoid confusing pronouns, throughout this article the principal or boss is assumed to be female and the agents to be male.

⁴ Consider these remarks from two eminent texts: “A second limitation of legislative control results from the increasing complexity and volume of governmental activity. Legislators . . . are highly dependent on the very administrative agencies they seek to control for *the information necessary for rational choice*, for recommended solutions to public problems, and even for the identification and description of the problems themselves” (Simon, Smithburg, and Thompson 1950, 524; emphasis added). “In general, authority should be placed at the lowest level at which all essential elements of information are available” (Wilson 1989, 372).

⁵ Distribution *W* is less risky than *Z* if the latter has “fatter tails.” To keep the comparison clean, most scholars assume a *mean-preserving spread*. Loosely speaking, one produces a mean-preserving spread of *W* by “squashing” it down—shifting probability mass from the distribution’s center to its tails. For two-parameter distributions (e.g., the normal or uniform), *Z* is a mean-preserving spread of *W* if it has the same mean and more variance.

⁶ Other rationales for delegation include comparative advantage (regarding, e.g., time). An extension of our basic model captures this logic.

premises: (1) Both policies and outcomes are unidimensional; (2) the policy technology is stark—outcomes = policies + random shocks; and (3) shocks are from a specific distribution, usually the uniform. Whether these assumptions are substantively important or only analytically useful has received scant attention in the literature.⁷ (It is argued, however, that assuming uniformly distributed shocks is just for computational convenience [e.g., Gilligan and Krehbiel 1990, 536].)

Our focus on more general models of delegation allows us to show that *none* of the canonical model's specialized premises drive the decision to delegate in political institutions, and only the technological assumption strongly affects decisions about which agents are given authority. Though we find that current scholarship overstates the importance of risk aversion, the assumption that outcomes = policies + random shock (especially the premise that the shock's distribution is unaffected by the policy) turns out to be very important—a fact that has gone unappreciated by the literature. Relaxing this convenient but empirically suspect assumption can dramatically affect predictions about which agents are given authority.

To show that we only need very general assumptions about preferences, policy and outcome spaces, technologies, and uncertainty to understand the core of delegation, we construct a nested sequence of models. Model **A** is the most general. It assumes only that agents have Euclidean preferences over outcomes (which lie in N -dimensional Euclidean space) and that agents make policy choices (which lie in K -dimensional Euclidean space, where K need not equal N). The structure of uncertainty and the relation between policies and outcomes are unspecified. Yet even at this level of generality we can show the validity of Hamilton's insight that delegation is driven by the search for state-contingent plans. Hence, by using a more general formulation, we can isolate the importance of specific features: For example, we can identify exactly which results truly hinge on risk aversion (in the classical sense)—and which do not. Thus, we depart from and generalize the canonical model not merely to check the robustness of its results but, more fundamentally, to locate and correct errors in the received wisdom about delegation.

Model **A'** imposes an additional technological assumption: The mapping from policy to outcomes has the common additive form, outcome = policy + random shock, where the shock is independent of the policy choice. This simple structure allows for stronger results. Indeed, the richness of **A'**'s results indicates that the literature has greatly underestimated the technological assumption's significance: We think that this part of the canonical model is more powerful than the

premises of risk aversion or unidimensionality, whose importance has been *overestimated*.⁸

We then consider model **A''**, which additionally assumes risk-averse decision makers. Moving from **A'** to **A''** allows us to see that only some features of delegation hinge on attempts to reduce risk rather than on the more fundamental desire to avoid bad outcomes and secure good ones.

Even the most specific model, **A''**, is considerably more general than the canonical one: It dispenses with specific functional forms for the utility functions (only concavity is presumed) and distributions of the shocks (only symmetry is required), and it is set in an arbitrary-dimensional space. Though the move to more general models often complicates analysis and obscures intuition, we find that for this subject the move *aids* understanding. The more general analysis, by avoiding algebraic derivations, helps us see the relation between results and major assumptions.

To demonstrate that a “spatial” theory of delegation (i.e., one with Euclidean preferences) not only is powerful for the “big picture”—delegation's central features—but also can yield detailed institutional insights, we consider some extensions to the basic models. These focus on three key aspects of delegation examined by many variants of the canonical model: (1) how a boss induces subordinates to acquire information, (2) her ability to make policies contingent on the world's true state, and (3) how closely outcomes reflect her preferences.⁹

At each level of generality we discuss not only that level's results but also the omissions: what a given model does *not* imply. This comparison of results and assumptions clarifies which premises drive which results. To facilitate these comparisons we derive results at the level with the weakest set of sufficient assumptions. (For example, if a result can be established in model **A'** but *not* in **A**, then it is presented in **A'**.) An overview of the three models and their results is given in Tables 1 and 2.

Before turning to the analysis, it is helpful to formalize some key decision theoretic issues associated with both the traditional and the contemporary perspectives. Although these issues turn out to be vital in delegation decisions, the literature has mostly ignored them.

⁸ This explains our sequencing of nested models: In going from model **A** to **A'**, one generates more new results by assuming the additive policy technology (while keeping the general preferences of model **A**) than by positing risk aversion (while keeping the general policy technology of **A**).

⁹ We always allow for multiple agents but not always multiple principals. Two reasons explain the latter. First, if issue spaces are unidimensional, then collective choice procedures (e.g., majority rule) can ensure that one boss is pivotal. Second, though recent literature has focused on how Congress delegates authority to the executive branch, delegation often occurs in standard tree hierarchies (e.g., from the president to the armed services), where multiple subordinates report to one superior. Moreover, the issue of delegation confronts leaders in *all* political systems, not just democracies: A dictator—the quintessential solo boss—faces hard delegation choices. Our theories of delegation should be able to accommodate different kinds of political institutions.

⁷ Battaglini (2002), an exception to this neglect, shows that some of the canonical model's premises do matter substantively. He proves that when the outcome and policy spaces are at least two-dimensional, equilibria exist in which informed subordinates reveal their private information to the boss. (See also Baron and Meirowitz 2002.) This finding differs sharply from the conclusions in Gilligan and Krehbiel 1989 and from those in Krishna and Morgan 2001.

TABLE 1. Comparing the Models


	Model A	Model A'	Model A''
Assumption			
Policy, shocks, outcomes	General	$x = p + \omega$	$x = p + \omega$
Preferences	Euclidean over outcomes	Euclidean over outcomes	Euclidean and strictly concave over outcomes
Result			
Ally principle in basic version	Ally principle holds.	Same as A	Same as A'
Efficiency of delegation	—	—	Not delegating is ex ante inefficient.
Increased preference conflict	Boss is worse off. Effect on expected outcome is ambiguous.	Boss is worse off. Effect on expected outcome can be pinned down in some cases.	Same as A'
Increased policy uncertainty	—	Delegation set expands.	Same as A'
Increased risk aversion	—	—	Delegation set expands.

Note: Entries of — denote results or effects that cannot be established at the level of generality of the model in question.

TABLE 2. Extensions

Extension	Model A	Model A'	Model A''
Busy boss	As cost of control increases, delegation set expands.	Same as A	Same as A'
Informationally imperfect agents	As imperfection increases, delegation set contracts. Notion of “ally” is ambiguous but dominated agent will not be chosen.	Same as A	Same as A'
Costly information gathering	Ally principle may fail even if costs are homogeneous.	Ally principle holds if costs are homogeneous.	Same as A'
Ex ante control and ex post auditing	For some costs the system is used when no agents are in the delegation set. Compliance may be state-contingent. Discretion decreases in detection probability and allies may be harder to control.	Same as A	Same as A'
Reasserting control	(1) With heterogeneous competence, ally principle may fail. (2) If delegation occurs, it is to an undominated agent.	Same as A	Same as A'
Agents can commit	With two or more agents the boss gets her ideal outcome (generically).	With costly information-gathering, delegation set is a doughnut.	Same as A'
Multiple bosses	With reversion to control by a boss, delegation is more likely if the voting rule is less restrictive or bosses have more similar preferences.	Same as A	Same as A'
Political uncertainty	—	—	Delegation set is nonempty and expands as bosses become more locally risk averse.
			As uncertainty about preferences increases, delegation set decreases.

FIGURE 1. Spatial Preferences and Choice under Uncertainty: Distance Aversion versus Risk-Aversion



DECISION-THEORETIC FOUNDATIONS OF DELEGATION

Mapping Hamilton’s thesis—delegating can obtain superior outcomes—into modern theories of choice is more than translation; it also helps us to see the claim’s generality. Let us translate the intuitive but fuzzy notion of “superior outcomes” into its modern equivalent, *first-order stochastic dominance*. The gist of this concept, a standard one in decision theory, can be conveyed by a simple example. Suppose that a person with an insatiable desire for a good (e.g., money) chooses between two lotteries, X and Y , where $X = Y + 10$ dollars. Then X is shifted to Y ’s right, toward better (higher) values of the good; hence X first-order stochastically dominates Y . Hence the colloquial phrase, “It’s a no-brainer,” captures the situation: *Any* rational person with these monotonic preferences will take X over Y . *This conclusion requires no further assumptions*. In particular, the decision maker’s attitudes toward risk are irrelevant: Even if the stochastically better distribution were more variable than the dominated one, (e.g., $X = Y + Z$, where Z is five or 50 dollars with equal probability), it would not matter because all the “risk” is benign.

We can now reinterpret the traditional view as arguing that delegation produces outcome distributions that first-order stochastically dominate those produced without delegation: Delegating to qualified specialists will shift those distributions toward better values. Of course, given “political” (spatial) preferences, where people have bliss points and want outcomes to be close to ideal, the idea of first-order stochastic dominance must be used carefully. Because such preferences are noneconomic (satiating), “better” means “closer to the decision maker’s bliss point.” Hence, given spatial preferences it is easy to confuse first-order stochastic dominance with risk reduction.

A simple example reveals the decision theoretic features of our analysis. Figure 1 depicts a symmetric single-peaked utility function with a bliss point of zero.

Consider a lottery with two equally likely outcomes: x and $-x$. Though *outcomes* are random, in an important sense nothing of consequence is stochastic: The agent’s *utility* is the same for either realization of the “lottery.” This example demonstrates that theories of choice under uncertainty with strictly monotone preferences must be adapted when applied to Euclidean (or, more generally, quasi-concave but nonmonotone) preferences. Although variance in outcomes produces variance in utility when preferences are insatiable (and outcomes are one-dimensional), in the spatial setting *this relation need not hold*.

Figure 1 illustrates a second important point: With Euclidean preferences the utility function’s curvature is only weakly related to choice under uncertainty. The figure’s utility function is not concave. Indeed, for all neighborhoods to the left or right of the bliss point the function is strictly *convex*. Hence it does not exhibit risk aversion in the classical sense. To see this, consider a lottery between outcome $-x$ and outcome $-y$. This lottery’s expected utility exceeds the utility of its expected value. Despite this, the utility function exhibits features of risk aversion given lotteries that are sufficiently symmetric and centered close to the ideal point. For example, the lottery with two equally likely outcomes, y and $-y$ (where $y > x$) is a mean-preserving spread of the lottery of x and $-x$ with equal probability, and the agent strictly prefers the latter. This apparent preference for risk reduction is equivalent to something more fundamental: In this setting, more risk means that outcomes are *farther from the ideal point*, hence less desirable. It is therefore instructive to think about lotteries over *outcomes* in terms of the corresponding induced lotteries over *utility*. Our analysis shows that for some distributions, second-order stochastic dominance in the former is the same as first-order stochastic dominance in the latter.¹⁰ Moreover, when decision makers can choose policy to center the outcome distribution

¹⁰ Focusing on first-order stochastic dominance regarding lotteries over utilities highlights risk aversion’s second-order character.

at their ideal point, they typically confront outcome distributions in which this equivalence holds.

The presentation of the following text is informal: It focuses on conveying the intuition behind the results. Due to space constraints, only a few major results are proven in this paper (see the Appendix). For proofs of all results, see Bendor and Meirowitz 2003a and 2003b. Anything identified in boldface as a “conclusion” in the text is proven either in the Appendix or in the working paper.

MODEL A

We consider a boss and a finite set of subordinates, $\{1, \dots, m\}$, $m \geq 1$. The set of outcomes is called X , a subset of N -dimensional Euclidean space; a particular outcome is a vector x . Each agent, say i , has a unique ideal point y_i in X . The utility of outcome x to agent i is given by a common, continuous, and decreasing function of the distance between x and the ideal point y_i . Thus, preferences are Euclidean: Player i strictly prefers outcome x_1 over outcome x_2 if and only if x_1 is closer in Euclidean distance to his ideal point than is x_2 . (Hence in two dimensions indifference curves are circles.) To avoid the annoyance of an indifferent boss, we assume that all ideal points are strictly rank-ordered by their Euclidean distance from the boss’s ideal point.

Policies are located in a K -dimensional space (K need not equal N). The set of policies is P ; a particular policy is a vector p . Policies are perturbed by random shocks (ϵ) to yield outcomes. In the basic model we assume only that the outcome distribution, given an uninformed decision maker’s choice of a specific policy, is not degenerate: More than one outcome can arise. Thus, in contrast to most delegation models, the shock need not be uniformly distributed; indeed, it need not be continuous, symmetric, or even independent of the policy. We require only that together the policy and shock create *some* uncertainty in outcomes.

Because outcomes depend on both policy and disturbance and because weak premises about the disturbance’s distribution suffice, outcomes can be seen simply as random variables that depend on policies: x ’s distribution function given policy p is $F(x|p)$. In contrast to the canonical technology (outcomes = policies + shock), our basic model stipulates neither a technology nor a parametric family of shocks. Instead, uninformed agents have beliefs over outcomes,

via the policy-conditioned distribution of $F(x|p)$. We sometimes say that informed agents have learned the mapping from policy to outcomes. Explicitly, an informed agent has concentrated beliefs over outcomes (given policies). Since the canonical model presumes shocks, we sometimes describe informed agents as having learned the shock’s value. Explicitly, $F(x|p)$ can be modeled as a function $x = g(p, e)$ where e is some high-dimensional random variable. This allows for much flexibility in how policies and shocks generate outcomes. For example, the variance of outcomes may depend on which policy is chosen: In the canonical model this variance is constant.

We do, however, impose some structure on the policy technology. (The following properties—and much more—hold in the canonical model.) First, we assume that given the uncertainty about the mapping from policies to outcomes, each agent has an optimal policy (it need not be unique). Second, knowing the perturbation permits complete “shock absorption”: If a decision maker chooses a policy after he observes the disturbance, then he can always find a policy that, when combined with the shock, yields any desired (feasible) outcome. We call this *perfect shock absorption*. Simply put, the shock maps a policy (in P) into an outcome (in X). Perfect shock absorption means that agents who observe the shock can get any specific outcome, x , by selecting the appropriate policy, p . Thus, an agent who selects policy while knowing the shock will choose one that results in her ideal outcome. (Multiple policies may produce this optimal outcome.)

All of the above is common knowledge (as is the game itself, below). The extensive form of the basic game is simple. The boss can either delegate or not. If she doesn’t, then she chooses a policy, which combines with the random shock to yield an outcome. Payoffs are then distributed. If the boss delegates, she selects an agent from the feasible set. (Recall, however, that this set may have only one member.) Once the agent is selected and given authority, he observes the shock’s realization and *then* picks a policy. Payoffs are then distributed.¹¹

Throughout, we consider subgame-perfect Nash equilibria, which resolve the boss’s indifference (to control or delegate) in favor of delegation. (If an agent is indifferent between policies at any decision node, we assume that he selects, from his set of optimal policies, one that is best for the boss, in expectation.) Subgame perfection is required to rule out equilibria that rely on noncredible play in some of the game forms we consider. Because none of our game forms involve signaling, subgame perfection is sufficiently

(Recall that Von Neumann–Morgenstern utility functions are linear and so exhibit risk neutrality over utility lotteries.) Alternatively, one could define a notion of risk aversion relative to certain lotteries. (For example, all decision makers with Euclidean utility functions exhibit risk aversion when we examine only distributions that are centered at the agent’s bliss point.) However, we have chosen not to use such a notion here. Instead, it seems more illuminating to highlight the fact that what *seems* to be classical risk aversion is equivalent to a more fundamental property of Euclidean utility functions: Bad outcomes are necessarily far from one’s ideal, and when outcome lotteries are centered at one’s bliss point, greater dispersion (“risk”) moves outcomes farther from one’s global maximum.

¹¹ There are two major types of delegation models in political science, defined by game forms: Either the principal moves first, typically by choosing to delegate or not (as in this paper), or the agent does, typically by sending his superior a message about the state of nature. The latter (often called signaling models) are harder to analyze; hence to figure out what drives delegation we have chosen the former. We believe, however, that our claim that first-order stochastic dominance (over utilities) is more important than risk reduction for understanding delegation holds in signaling models; the sequence of moves should not affect this point.

restrictive: Appeals to stronger refinements are unnecessary.

Results: Properties of the Equilibrium

1. Generically, a unique equilibrium exists for each set of parameters.¹²
2. The equilibrium has the following properties.
 - a. If the boss does not delegate, then she chooses a policy that maximizes her expected utility, given her knowledge of $F(x | p)$. Because her preferences are strictly single-peaked (no flat spots) and because $F(x | p)$ is nondegenerate for all p , the expected utility of such a policy must be strictly less than the utility garnered by getting her ideal outcome with certainty. Because utility varies continuously in outcomes, this implies that she has *certainty equivalents* in outcome space: a set of outcomes such that she would be indifferent between getting any such outcome (for sure) and taking a policy that maximized her expected utility. Given preferences that are monotonic with Euclidean distance, all these certainty equivalents are equidistant from her ideal point. Since $F(x | p)$ is not degenerate, this distance must be strictly positive.

This set of certainty equivalents defines a *delegation cutoff rule* or ideological distance threshold for the boss: She delegates if and only if some agent is at least as ideologically close to her as this threshold. Call this threshold d^* ; anything within this boundary is the boss's *delegation set*.

Because delegating is optimal for the boss if her delegation set is not empty and because this set's radius is strictly positive, even model **A**'s very general assumptions imply the possibility of rational delegation. *None of the canonical model's four main premises (e.g., risk aversion) is necessary.* The key to delegation lies elsewhere: It is based on the boss's ignorance of the state of the world, captured by $F(x | p)$'s nondegeneracy, and the agent's ability to condition his policy choice on the true state. This combination is precisely the optimality of state-contingent plans emphasized by the traditional perspective.

- b. If the boss delegates, then she picks the agent whose ideal outcome is the closest to hers, in Euclidean distance. (We sometimes call this person the "best agent.") Since the agent knows the shock's value before choosing a policy and is a perfect shock absorber, he picks his induced ideal point in policy space, thus obtaining his ideal outcome. This is clearly optimal for the agent.

Thus this ancient political principle—the boss picks the (ideologically) closest agent—is valid

given very general assumptions about preferences, shocks, and policy technology. In particular, her risk preferences are immaterial. The ally principle is driven instead by the principal's desire to minimize the costs of delegating by appointing someone with preferences as similar to hers as possible. This satisfies *first-order* (not *second-order*) stochastic dominance regarding the random distance between the boss's ideal point and the final outcome; hence, risk preferences are irrelevant.

- c. If the boss delegates, then the equilibrium is efficient, both *ex ante* and *ex post*. (It must be efficient: The delegatee gets his ideal outcome, which he strictly prefers to all others.) If, however, the boss controls the policy, then the outcome may be inefficient. For example, as model **A**' will show, if the players are risk-averse, then *ex ante* everyone would prefer that the boss get her ideal outcome for sure, which could be achieved if the boss delegated to an agent who promised to implement that outcome (and delivered on the promise), over the lottery produced by the boss making an uninformed choice. Because agents are perfect shock absorbers, if one promises to implement the boss's ideal outcome, he is technically able to do so. But the agent's promise to carry out this plan is not credible, so the boss disregards it.

Turn now to the comparative statics of preference conflict. Suppose that conflict increases: The ideal points of all agents move an equal distance away from the boss's. What happens? Three possibilities arise. First, if no agent had been in the delegation set previously, then nothing changes: The boss keeps the reins in her hands and the expected outcome is unchanged. Suppose, however, that before the exogenous shift delegation had prevailed. Then we could get the second possibility—the situation changes qualitatively—if the best agent becomes unacceptable to the boss, so she no longer delegates. Third, the preference change could be sufficiently small so that delegation remains attractive. In the second and third cases the expected outcomes shift, *but they can move in opposed directions*. In case three it must move *away* from the boss's ideal; in case two it can move *toward* it. Hence an increase in ideological conflict between the boss and her agents does not always push policy in the same direction. That greater preference alignment (in case two) can move expected outcomes *away* from the boss's ideal might contradict conventional wisdom.

An exogenous increase in conflict does have *one* consistent (and unsurprising) effect: The boss is (weakly) worse off the more her subordinates are outliers. More precisely, as the best agent's ideal point moves away from the boss's, the latter's expected utility falls or stays constant.

Conclusion: *When preferences between agents and the boss conflict more, (1) the movement of the expected*

¹² If the boss is indifferent between delegating (to the closest agent) and retaining control, the equilibrium is not unique. Due to her indifference these equilibria are payoff-equivalent to the boss (but not to the agents).

outcome is ambiguous, but (2) the boss's expected utility unambiguously falls.

Extensions of the Basic Model

The following extensions are to be taken as one-at-a-time modifications. (We will see how this serial approach matters when we examine, e.g., the boss's cost of effort—point 1—and, later, the subordinates' cost of information-gathering, in point 3.)

1. A Busy Boss. Delegation often arises when a boss confronts many issues. That can be easily represented here by a cost of effort of controlling, reflecting the opportunity cost of attending to the matter at hand. If the boss's expected utility of controlling is additive in policy outcomes and effort, then the busier she is the more she wants to delegate.

Conclusion: *When the boss becomes busier the delegation set expands.*

2. Informationally Imperfect Agents. In the basic model agents always observe the state of nature. What if agents may fail to gather information? Consider a simple type of fallibility: After the boss delegates, with probability q_i agent i observes the state of nature; with probability $1 - q_i$ he does not. An agent who does not observe ϵ must make an uninformed decision. This puts the boss in the worst of all possible worlds: There is policy bias but no compensating adaptation of the policy to the state of the world. Hence, the expected value of delegating to i is increasing in i 's competence (q_i). Consequently, if agents are homogeneously imperfect ($q_i = q$ for all i), as q falls the delegation set contracts: The less delegation is worth to the boss via increased expertise, the more ideological affinity she demands.¹³

Note that for $q < 1$ the delegation set need not be convex or connected. In fact, the *meaning* of “ally” is unclear in this setting. If all agents are perfectly competent, then the boss can ignore the location of their induced ideal policies: She knows that if she delegates, the delegatee will set policy to attain his ideal *outcome*, so only outcome proximity matters. But if shock absorption is imperfect, then the boss must also worry about the location of the agents' induced ideal *policies*, and given the generality of model **A**, her closest ally in outcome space may not be the closest in policy space. Consider an example where the outcome and policy spaces are the real line. The boss's ideal point is 0, agent 1's ideal is $-1 + \epsilon$ (with $\epsilon > 0$ and small) and agent 2's is 1, so agent 1's ideal is closer to the boss's than is agent 2's. Suppose that $F(x | p)$ is as follows. If $p > 0$, then $x = p + 0.1$ or $x = p - 0.1$ with equal probability. If $p < 0$, then $x = p + 0.2$ and $x = p - 0.2$ are equally likely. If agent 1 or 2 is chosen but does not observe the shock, he will select policies $-1 + \epsilon$ and 1, respectively.

¹³ This implies the following cutoff rule for expertise: For any agent i with ideal point y_i^* (distinct from the boss's ideal) who would be in the delegation set for $q = 1$, there is a threshold $\bar{q}_{y_i^*} > 0$ such that i is acceptable to the boss if and only if $q \geq \bar{q}_{y_i^*}$.

Hence agent 1's policy entails more outcome risk than agent 2's. If the boss is risk-averse and ϵ is sufficiently small, then the boss prefers delegating to agent 2. Yet this ordering is the opposite of her ordering of the agents' ideal points in outcome space. This reveals that, because the meaning of ally is ambiguous in this setting, the ally *principle* is also unclear.

Now consider a general increase in competence: q_i rises for each agent. Obviously this cannot hurt the boss. More interesting is the effect on the expected outcome in equilibrium: By inducing delegation, higher competence (like reduced preference-conflict) can move the expected outcome *away* from the boss's ideal point. This point has substantive bite and differs from conventional wisdom. Agencies with more competent subordinates may on average attain outcomes that are farther from the boss's ideal. This points to the subtle yet important distinction, noted earlier, between a policy lottery's expected outcome and the expected *distance* of its outcomes from the boss's ideal point. The boss need not prefer a lottery whose expected value is close to her bliss point over one whose mean is farther away. This may seem perverse (Why prefer a lottery with a more biased mean?) or risk-averse (she accepts more bias in return for less variance), yet in fact it is never the former and need not be the latter.¹⁴

Since agents' competence may vary, ability and ideology may trade off. Thus the ally principle must be understood as a “*ceteris paribus*” rule: *All else equal*, a rational boss should choose her closest ally as an agent. If that ally is a dolt, delegating to him is silly.

Though the meaning of ally is ambiguous in this setting, one could use a *weaker* sense of the concept—a partial ordering over three dimensions: quality, proximity of ideal outcomes, and proximity of ideal policies (when choosing under ignorance). Then a weaker version of the ally principle holds: The Boss never delegates to a dominated agent.

However, she never delegates to outliers—agents outside her delegation set given $q_i = 1$ —even if they were perfectly competent and allies were fools. Thus, although some competence is necessary for receiving authority, it is *not* sufficient. (The same holds for preference similarity.) As in the old communist phrase, to receive authority an agent must be *both* “red” and “expert.”

Conclusion: (1) *As agent competence decreases the delegation set contracts.* (2) *With homogeneous competence the ally principle may be violated as an ally's ignorant choice may be worse than a more distant agent's.* (3) *With heterogeneous competence the ally principle may also be violated as the boss trades off competence for preference similarity.* (4) *A dominated agent (regarding competence, proximity of ideal outcomes, and proximity of ideal policies) is never picked.*

¹⁴ Recall the unidimensional example in which ignorant choice produces y or $-y$ with equal likelihood. The expected *outcome* is zero—exactly the boss's ideal point—but this “lottery” yields, *with certainty*, an outcome of distance y from her ideal. Hence, she prefers delegating to a sufficiently competent agent whose ideal is x , where $0 < x < y$, even though that yields a biased lottery.

3. Costly Information-Gathering. In the basic model agents freely observe the state of nature. But becoming informed is usually costly. So consider a model similar to that of Gilligan and Krehbiel (1987): An agent, i , who pays a cost $c_i > 0$, observes the true state; otherwise, i stays uninformed (unspecialized). Utility is additive in policy and the cost of getting informed.

All agents have cutoff rules: “Specialize if and only if $c_i \leq c_i^*$,” where the threshold c_i^* solves an indifference condition. Hence the boss may not select the closest agent: That one may have high specialization costs and may elect to remain uninformed, and an uninformed agent is useless. Instead, if she delegates, the boss will pick the closest agent of those who would specialize.

Even if c_i is a constant, agents may have different cutoff rules on the specialization decision because the benefits of acquiring information may vary. That is, given the generality of the $F(x|p)$ formulation, different delegates might face different amounts of uncertainty associated with their ignorantly selected policies.

Obviously, if everyone’s c_i is sufficiently large, then all agents prefer to stay ignorant. Anticipating this, the boss won’t delegate. But low specialization costs cannot make a preference outlier acceptable to the boss. Suppose that initially c_i is so high that i would stay ignorant if given authority. Then the boss will not delegate to i , regardless of his policy preferences. Now let c_i fall, so that i would specialize if appointed. Nevertheless, if i is such an outlier that he is not in the delegation set, then this parametric shift cannot bring him into the boss’s favor. As with competence, “sufficiently low” information-gathering costs are necessary but not sufficient for an agent to be picked. The boss wants subordinates who are both spatially proximate and informationally efficient. The latter ensures that inducing specialization is feasible; the former, that it is valuable (to the boss).

Conclusion: (1) As information-gathering costs rise the delegation set contracts. (2) Even with homogeneous costs the ally principle need not hold because the gains of acquiring information may vary across agents, making proximate agents less willing than distant ones to specialize.

Control Systems

4a. Ex Ante Controls and Ex Post Auditing. The boss may be able to exert some control over agents, by limiting their discretion ex ante and enforcing that with ex post auditing and punishments (e.g., Huber and Shipan 2002 and McNollGast 1987).¹⁵ We follow Huber and Shipan (H-S) in assuming that (1) reducing an agent’s discretion is costly for the boss and (2) this cost increases strictly and continuously the more discretion is reduced. The cost affects the boss’s utility additively.

The sequence of moves is as follows. The boss moves first, either delegating or not. If she sets policy, everything is as it was in the basic model. If she delegates, she can limit the agent’s discretion: An outcome is acceptable if and only if it is within the “discretion set,” i.e., the outcomes that are within d of her ideal point. (H-S assume that discretion pertains to policies, not outcomes. Both assumptions are reasonable; they describe different programmatic environments. Wilson [1989, 158–71] examines how policy and outcome observability varies empirically.) The bigger d is, the more discretion the agent has.

The boss can choose to pay a cost $c \geq 0$; if she does, then the final outcome will be inspected with probability $q \geq 0$. (For $c = 0$ auditing is a free “fire alarm” pulled by constituents [McCubbins and Schwartz 1984].) The delegatee learns the mapping from policy into outcomes and sets a policy. The game ends here if the boss did not pay the auditing fee. If she did pay it, then with probability q she learns whether the outcome is in the discretion set. If the outcome is acceptable, the game ends; if not, the agent pays a fine f . If the boss catches the agent deviating, then she must ignorantly choose a policy. (This departs from the H-S model, in which an unspecified nonstatutory mechanism forces a detected shirking agent to implement the boss’s ideal point.¹⁶)

If she decides to use a control scheme, the boss must select an agent, a discretion level, and a detection probability. For a fixed control scheme and random shock, the agent’s choice is simple: Either he can select, of the acceptable outcomes, the one closest to his ideal or he can get his own ideal outcome and risk detection.

Despite their differences, in both the H-S and our model the subordinate’s optimal compliance decision may be state contingent: He may comply for some, but not all, of the shock’s values. In the H-S model, if the agent is caught disobeying, then the final outcome is the boss’s ideal. Hence for the agent the policy consequences of being caught do not depend on the shock’s value. However, the consequences of disobeying do depend on the shock if he goes free.

In our model the boss chooses ignorantly, so a detected deviating agent *does* incur state-contingent impacts. Because the agent finds out the disturbance’s value *before* deciding whether to comply, the compliance decision is also state-contingent: Complying is optimal for some shocks but not others. (For example, disobeying is unattractive for an extreme perturbation: If the agent is caught, then the boss’s uninformed choice—insensitive to the shock—will produce an outcome that is awful for the agent. But if the shock would distort the boss’s policy selection toward the agent’s

¹⁵ Although these considerations affect agents directly only if the boss delegates authority, they also influence her decision to delegate in the first place.

¹⁶ We drop their premise for three reasons. First, the nonstatutory mechanism in H-S is a vital part of the model. Hence, it should be modeled explicitly, as a set of rational actors with specified monitoring capacities. Second, if catching a deviating agent means that the boss gets her ideal outcome, then she has a perverse incentive to induce agents to shirk. Third, if an auditor (part of H-S’s nonstatutory mechanism) sees the unacceptable outcome but not the policy, then agents may have an incentive to distort policies, to confuse the auditor.

bliss point, then being caught is not so bad.) Hence in equilibrium the boss, having observed that the agent deviated, must update her beliefs about the shock.¹⁷ Thus, although the agent's and Boss's decision problems differ in these two models, in either case compliance may depend on the shock. This helps explain why even optimal control schemes may only sometimes induce compliance: The agent's private information affects the relative costs of submitting to a control system (and systems that always induce compliance may cost too much). In contrast, in other models compliance is not state-contingent and so is binary: The agent always complies or never does. (If the agent won't comply, then installing a control system is stupid, so such models imply that if controls are used, then the agent always complies.)

Conclusion: *Even under optimal control schemes an agent may, depending on the shock's realization, choose to comply but, at other times, intentionally disobey.*

Although selecting a system that induces compliance for all possible shocks may not be *optimal*, a tightest control system—that the agent will invariably obey—always exists. This system is characterized by a discretion level d_m for which agent m is willing to enact an outcome exactly d_m from the boss's ideal, for every shock, rather than disobey and risk detection.¹⁸ Though the boss may not want to bind the agent so tightly, it is useful to analyze how this minimum discretion level is affected by the control variables q and f and the agent's ideal point, y_m . Intuitively, as q or f increases the agent becomes more willing to comply: d_m falls. Surprisingly, the ally principle may be violated here: Even if y_1 is closer to the boss's ideal, d_1 may exceed d_2 . This can happen because being caught deviating might hurt an outlier more than an ally. Suppose, e.g., that the fine is small: A shirking agent is punished mostly by losing control to the boss. The (common) loss function is piecewise concave: Almost-flat near the bliss point, it plunges steeply once a specific distance from the ideal. (For simplicity assume that the shock barely affects utility.) Thus, if y_1 is close to the boss's bliss point, then losing authority won't hurt agent 1 much, whereas if y_2 is far away, then agent 2's losses would be catastrophic. Then agent 2 will be more compliant than agent 1.

Conclusion: (1) *The discretion level, d_m , is decreasing in the detection probability and the fine.* (2) *Allies may be less controllable than outliers: The discretion level need not be increasing in the distance between y_m and the boss's ideal.*

¹⁷ Because compliance is shock dependent, the solution concept should be sequential or perfect Bayesian equilibria. A boss who chooses policy after learning that the agent shirked will optimize given updated beliefs about the shock; these beliefs must reflect the information implied by the agent's actions.

¹⁸ Recall that we assume that if the agent complies, then he chooses the final policy. But in H-S, nonstatutory mechanisms make policy revert to the boss's ideal with positive probability even if the agent complies.

H-S (2002) address similar issues in their proposition 1: "Increases in policy conflict . . . and decreases in the reliability of nonstatutory factors [analogous to our detection probability] always make it easier to satisfy the conditions under which the Politician adopts a low discretion law as opposed to a high discretion law" (248).¹⁹

Whereas H-S find that increases in policy conflict induce the boss to constrain agents more, we find that the boss *may* want to give more discretion to distant agents than to closer ones. The counterexamples to H-S's intuitive finding (i.e., more conflict results in less discretion) hinge on assumptions about the agents' preferences: These influence the discretion levels that agents will accept. As the above example shows, nonlinear loss functions and a richer set of shock realizations can yield different conclusions.

The two models differ still more in their analysis of detection and its impact. In ours, the detection probability and discretion level vary inversely: If the chance of detection falls, to induce compliance the boss must give the agent more discretion. H-S obtain the opposite result because they assume that even if the agent complies, nonstatutory factors affect the chance that the boss will get her ideal outcome. This cannot happen in our model, given optimal compliance (and $d_m > 0$). Thus the models differ crucially in their assumptions about when the boss gets her ideal outcome. Both assumptions seem to be reasonable descriptions of different institutions.

Plausibly, in the short run the boss can select d and q but not f . In Bendor and Meirowitz 2003a and 2003b we show that, first, if f is sufficiently big and $q > 0$, then any (d, q) pair will produce full compliance and, second, for any f, q where $f q > 0$ there is a discretion level d that the agent will obey. The boss's optimal control scheme can be designed as follows: (1) For each agent m select the optimal (d_m, q_m) ; (2) given these optimal (d, q) pairs for every agent, pick the optimal agent m ; (3) compare controlling this agent versus ignorantly choosing policy. Setting d sufficiently large or $q = 0$ is, in effect, delegating without control. Given our models' level of generality, many types of institutions can be optimal. For example, if the cost of monitoring and the marginal cost of reducing d are sufficiently small, then the boss will install the control system and use it to constrain an agent; if f is also sufficiently big, she will get her ideal outcome with certainty. But when control is very costly or fines are small the boss may prefer delegating without imposing constraints.

These findings suggest that when reputations matter (big f) and monitoring is cheap we should observe control schemes. But tight systems (small d) are not necessarily optimal even when some monitoring is cheap. For tight controls to be worthwhile, the *marginal* cost of reducing discretion must be small and/or fines for shirking must be large. Finally, assume that the optimal control scheme induces constant obedience. Then we

¹⁹ The topics are similar but the approaches differ: H-S examine the decision to select low discretion; we focus on the optimal level of discretion that induces full compliance (even if inducing full compliance is itself suboptimal).

assert that if f is increased, then d or q or both must fall. Why? The intuition hinges on this: If f goes up by a little, there is a correspondingly small decrease in q and/or d that will keep compliance optimal for the agent. So when f rises the boss can do strictly better: She can get the same outcome for less (lower q) or a closer one at the same cost (lower d) or both and split the gain. Thus, subordinates are more useful to the boss when she has a rich set of controls than when her only options are ignorant control or open-ended delegation.

Conclusion: (1) *The optimal control scheme varies with the costs of monitoring and punishment.* (2) *Tight systems (small d) are optimal when the marginal cost of monitoring is low and/or the fines for shirking are big.* (3) *In an optimal control system that induces full compliance, if f rises, then d or q or both fall.*

Part 2 of this conclusion is consistent with the second part of H-S's proposition 1: "A low discretion law is adopted in equilibrium only when legislative capacity is sufficiently high (a is sufficiently small)" (248). The rest of our conclusion makes predictions about the relations between institutional properties (the ability to punish disobedience) and observable control decisions (degree of discretion and monitoring effort).

Comparing our model to that of H-S is complicated for two reasons. First, the models make assumptions about preferences, dimensionality, policy technology, and shocks at very different levels of generality. Second, they differ on two key structural dimensions: how detecting noncompliance affects the final outcome and whether discretion pertains to policies or outcomes. The two types of intermediate models—those differing from either ours or H-S's in only one structural dimension—are worth investigating. Though doing so is beyond this paper's scope, we offer the following conjecture: The agent's compliance decision is independent of his private information (regarding how policies map into outcomes) in only one type of model—when discretion is over outcomes and detecting disobedience leads to the boss getting her ideal outcome.

4b. Reasserting Control. A simpler form of control is for the boss to reassert her authority if matters go awry. Of course, taking control is one thing; choosing wisely is another. If seizing the reins just puts the boss back in her original position—facing the tension between control and expertise—then she may have gained little.

Given, however, imperfectly competent agents (extension 2, above), the benefits of this simple control become clear. Consider the following time line. As always, initially the boss can either set policy (ignorantly) or delegate. If she delegates, she waits to see whether the agent manages to become informed. If he does, the boss remains passive. If, however, the agent fails—stays uninformed—then the boss reasserts control and sets policy herself. (One can easily extend this analysis by assuming that retaking control is costly.) True, she will do so ignorantly, but better than let an equally ignorant subordinate choose.

This ability to (freely) retake the reins implies that the delegation set is once more decisive: The boss will

delegate if and only if some agent's ideal outcome is in the delegation set. Because she can reassert control, the boss need not worry about whether agent i is "sufficiently" competent: If i does not observe the random shock, then the boss will take over.

This control process therefore restores clarity to the meaning of "ally." Since the boss will set policy if the agent is uninformed, that agent's induced ideal *policy* (when choosing in ignorance) is of no concern. This suggests the following prediction: The ally principle will hold more often in institutions that allow a boss, upon seeing that a subordinate is making uninformed decisions, to retake the policy reins.

Although the ability to reassert control implies that "ally" has only one meaning (proximity in outcome space), agents may still vary on the dimension of competence. Of course, *ceteris paribus* the boss prefers more able agents. Hence she will only select an agent from the *efficient set* of subordinates. This set is defined by a dominance relation. (For example, if $q_i = q_j$ and agent i is closer to the boss, then i dominates j .) Undominated agents form the efficient set, and the boss ignores any agent who is weakly dominated by another agent, regarding competence and outcome preferences. Because by assumption agents' ideal points vary in their distance from the boss's, if the efficient set contains multiple agents, then they must be strictly rank-ordered by competence: The nearest agent is the set's weakest member; the farthest one, its most able.²⁰

Conclusion: (1) *When the boss can reassert control, she prefers delegating to any agent in the delegation set (regardless of competence) over ignorantly choosing policy.* (2) *If the boss delegates and competence is homogeneous, then she will pick the closest agent.* (3) *If competence is heterogeneous, then the ally principle can fail but the chosen agent must be undominated.*

5. Agents Can Commit. In the basic model agents cannot commit to policies or outcomes. Hence the boss anticipates that any agent will, if given the opportunity—the authority—do what is in his best interest. Suppose, however, that agents can "join the boss's team" by committing to an outcome.²¹ (We are not relaxing the requirement of subgame perfection; we are examining subgame perfect equilibria in a new game form.) In this extension the agents first learn the shock's value and then announce (commit to) an outcome.²² Each agent's strategy space is the set of constant functions from the state space to the outcome space. The agents simultaneously announce this function, then the boss chooses to delegate or not. If she delegates, she picks an agent who then selects a policy that yields his announced outcome. Otherwise the boss selects policy ignorantly. As we will see, in equilibrium

²⁰ In the special case of homogeneous competence ($q_i = q$ for all i) the efficient set is a singleton, so we recover the ally principle.

²¹ Since we are modifying the basic model serially, here we continue to assume perfect shock absorption. Thus committing to outcomes is feasible.

²² Allowing agents to commit before they learn the state yields the same results.

the delegate commits to a policy that is suboptimal for him.^{23,24}

Commitment by itself can benefit the boss, even if there is only one agent ($m = 1$). Not always, of course: If this agent is in the delegation set, then commitment is useless—the agent will implement his ideal outcome, which he could credibly promise to do even if unable to precommit. But if the agent is outside the delegation set, then commitment changes the equilibrium dramatically. Without commitment the boss won't delegate—the distributional loss is too great—leading to a risky and possibly inefficient choice by the ignorant boss. But if the agent can commit to a policy before the boss's delegation decision, then the outcome is efficient: The agent will offer to deliver an outcome in the delegation set, and the boss will accept.²⁵

Now consider the additional benefits of having *multiple* agents who can commit. With $m > 1$, the situation resembles Wittmanesque competition among policy-oriented candidates. (Also related are agency models of competition [e.g., Ferejohn 1986].) Here, however, the boss is the only voter. Thus she is always pivotal, *regardless of the outcome space's dimensionality*. Two results follow. First, if the outcome space is unidimensional and at least two agents are ideologically heterogeneous (their ideal points bracket the boss's), then with commitment the boss gets her ideal outcome in equilibrium. Second, in multidimensional spaces the boss gets her bliss point in equilibrium for "almost every" configuration of ideal points. Intuition can be gained by studying just two agents, say 1 and 2, with ideal points that are not collinear with the boss's. Consider any outcome proposed by agent 1 that is not the boss's ideal. Agent 2 can offer an outcome (on a line linking the boss's bliss point to agent 2's) that is closer to the boss's ideal point than is agent 1's proposal and that benefits both the boss and agent 2. Similarly, responding to any promise by agent 2, agent 1 can propose a policy closer to the boss's ideal. Hence, the only pair of proposals that are best responses to each other is two commitments at the boss's ideal outcome.

If all ideal points are perfectly collinear, then in effect the outcome space is unidimensional. Here, if all subordinates are on the same side of the boss, then an equilibrium exists where the outcome is the ideal point of the boss's ally. Since agents are purely policy oriented, more distant ones would not "outbid" the favored agent; that would hurt them. But if they are even slightly noncollinear, then an outlying agent can

always find an outcome that he and the boss prefer to the ally's ideal point, and so on. If the policy space has at least two dimensions, then perfectly collinear bliss points have measure zero in the space of possible preference profiles, so this happens "almost nowhere." Significantly, these results hold even if no agent is in the boss's delegation set.

Thus, because the principal's ability to base personnel choice on credible promises about outcomes wipes out the agents' informational advantage, inducing competition among agents who can precommit is a very powerful tool for bosses facing delegation decisions.

Conclusion: (1) *If there is just one agent and he is outside the delegation set, then the ability to commit helps the boss.* (2) *With multiple agents and a unidimensional outcome space the final outcome is the boss's ideal as long as two agents have ideal points that bracket the boss's.* (3) *With multiple agents and a multidimensional outcome space the final outcome is the boss's ideal for almost all preference profiles.*

6. Multiple Bosses. Many political settings have multiple bosses and multiple agents. The bosses determine, by some procedure or institution, which subordinates (if any) are authorized to make policy. One simple institution has a *primus inter pares*: an empowered boss b_1 who proposes a delegation decision. A stipulated number of other bosses must approve this proposal. If no winning coalition supports it, then b_1 selects a policy in ignorance.

In a richer bargaining model, bosses sequentially propose delegation options. After each proposal the bosses vote. The first alternative supported by a winning coalition is enacted. If none is accepted, then again, b_1 must select a policy in ignorance.

No offer may be acceptable, in either extension, for either uni- or multidimensional outcome spaces. (However, if all bosses are strictly risk-averse over outcomes, as in model **A''**, a "collective delegation set" always exists. See **A''** for details.)

Both the simple and the rich models presume that delegation (not b_1 's policy choice) requires approval. Consider the set of winning (decisive) coalitions, given a voting rule: e.g., under majority rule, coalitions with over half the voters. Compare rules a and b , where a is more restrictive (fewer decisive coalitions). Clearly, under rule a weakly fewer proposals will be acceptable to a decisive coalition, so reaching a decision to delegate will be harder. If the empowered boss is in every decisive coalition—she has a veto—then delegation will be even less likely. On the other hand, fixing the rule (and so the set of winning coalitions), as the bosses' preferences become more alike it gets easier to find an agent backed by some winning coalition. Thus in this model delegation is more likely when rules are less restrictive or principals more similar ideologically.

While analyzing these models fully is beyond this paper's scope, we briefly compare two institutions, each with simple majority rule. In one, only boss b_1 can propose to delegate authority; in the other, any boss can. Hence, delegation will occur under institution 1 if and

²³ Thus, if one keeps the basic game and drops subgame perfection, this subsection's equilibria are not supportable, as here we require commitment on as well as off the equilibrium path.

²⁴ Because the agents have private information their announcements might be considered signals. If so, the right model is a signaling game, with its associated complexity. Happily, this complication does not arise here: The uninformed boss picks an agent *after* hearing the agents' (binding) offers; appealing to perfect Bayesian equilibria does not refine the set of equilibria. Because informed agents can precommit and the boss's action set is sparse, interesting signaling cannot occur.

²⁵ The delegation set has many outcomes that both prefer to the boss's ignorant policy choice—a spatial "surplus." Hence, a bargaining model is needed to pin down the equilibrium allocations.

only if a majority coalition *that includes* b_1 is better off delegating than having b_1 choose ignorantly. But in institution 2 no one monopolizes agenda-setting, so delegation will occur if *any* majority coalition is better off delegating. Therefore we predict that delegation will occur more often in institution 2.

Conclusion: *With multiple bosses and reversion to control by an uninformed boss, delegation is more likely when (a) the rule is less restrictive or (b) the bosses' preferences are more similar.*

Part a has implications for institutions such as the separation of powers. When multiple branches of government have veto powers, the rule is more restrictive: Fewer coalitions can decisively implement delegation; more can block it. Hence, all else equal we should observe less delegation in separation of power systems.

Part b may appear intuitively obvious. Yet it contrasts sharply with proposition 1 of Epstein and O'Halloran (1999, p 75). Their result concludes, "The closer the preferences of the committee to those of the median floor voter, the less likely Congress is to delegate authority to the executive," implying that the more ideologically similar are the political superiors (the floor median and committee), the *less* likely they are to delegate. Why the difference? In their model, the committee knows more than the floor median but less than the agency. The committee and floor play a cheap talk signaling game before the floor decides whether to delegate to the agency. So proposition 1 holds for two related reasons. First, the committee gives the floor less information when they are ideologically far apart (a robust result in signaling games.) Second, this induces the poorly informed floor to rely on the *other* informed decision maker, the agency, by giving it authority.

Each model focuses on properties ignored by the other. Our model does not examine situations in which a *subset* of bosses (a committee) is informed. Thus, it ignores the possibility that informed principals and agents might compete, regarding giving information to uninformed principals. By examining only unidimensional outcome spaces, however, their model ignores situations in which *many* bosses—not just the (unitary) floor median—make delegation decisions. Hence, it cannot analyze how preference conflict among multiple bosses impacts such decisions.

MODEL A'

Model A' is strictly more specialized than A: It uses all of the latter's assumptions and adds others. The main new features concern the policy technology. Despite this greater specificity, A' is more general than the canonical model regarding utility functions (it does not assume risk aversion), the distribution of shocks, and the issue space's dimensionality.

Policy technology here is conventional, i.e., additive: outcomes = policy + noise. The noise, ϵ , is symmet-

rically distributed around the zero vector.²⁶ We again allow for N -dimensional outcome spaces. (Given the technology, policies and shocks must have the same dimension as outcomes.)

Comparative Statics

1. Increased Preference Conflict. Here we can definitively conclude that if the agents' ideal points move away from the boss's, then the expected equilibrium outcome in two cases move in opposite directions. If conflict increases enough to empty a once-occupied delegation set, then the boss sets policy and the expected outcome moves toward her bliss point. But if the change is not this large, then the expected outcome moves *away* from her ideal.

Conclusion: (1) *Small increases in preference conflict tend to move the expected outcome away from the boss's ideal.* (2) *Big increases tend to move the expected outcome toward her ideal.*

2. More Variable Shocks. If the shocks become more variable, via a mean-preserving spread of ϵ , will the boss's behavior be affected here? Conventional wisdom does not give a clear answer: Since A' does not presume risk aversion, why would more risk matter? The following result may therefore be surprising: If the uninformed boss's optimal induced policy equals her ideal outcome, then a riskier shock reduces the expected utility of not delegating. But delegating's value is unchanged, for the best agent will still be picked and will still get his ideal outcome. (Because he picks a policy *after* seeing the shock, its distribution is irrelevant.) Hence the boss becomes less choosy: Her delegation cutoff moves away from her ideal point. Thus in this circumstance delegation's relative benefits rise as the shock becomes riskier.²⁷ So the chance of observing delegation must be (weakly) increasing in this technological uncertainty.

This comparative static reveals the significance of spatial preferences, *independently* of risk aversion. As a key point, this is worth explaining. Consider the following example. Suppose that the boss's utility function is tent-shaped (piecewise linear) in a unidimensional space;²⁸ her bliss point is 0. The shock is $-x$ or x with equal odds. Given model A' (ϵ distributed symmetrically around zero, etc.), if the boss does not delegate, then setting $p = 0$ is an optimal choice. Now make ϵ riskier by letting it equal $\pm y$, with $y > x$. This lowers

²⁶ Assuming that the shock's mean is the zero vector causes no loss of generality: The distribution is common knowledge, so a rational (uninformed) agent could compensate for a non zero mean by shifting his policy appropriately.

²⁷ Model A' also implies that if the shock is unimodal and continuous, then for all Euclidean utility functions the boss's best induced policy is her ideal outcome. (Proof on request.) So the comparative static holds for this large class of shocks, including the canonical case of uniform noise.

²⁸ We could abuse terminology by referring to tent-shaped utility functions as locally risk neutral since their second derivatives are zero "almost everywhere."

the value of not delegating: Whether the outcome is $-y$ or y it is farther from her bliss point. This harms the boss simply because she has spatial preferences; that her utility function is locally linear nearly everywhere is irrelevant.

Thus, what drives this effect is not the second-order property of risk aversion but the first-order one of preferring near outcomes to far ones. Indeed, in this example a riskier shock hurts the boss even though *fundamentally she faces no uncertainty at all*: If she (optimally) chooses her bliss point, then she knows, *with certainty*, that the outcome will be a distance of y from her ideal. That it could be left or right of her ideal is immaterial, by virtue of Euclideanism. So it is more accurate to say that the parametric change makes determining outcomes, *ex ante*, harder for the boss than to say that it increases uncertainty.

Conclusion: *If the uninformed boss's optimal induced policy equals her ideal outcome, then an increase in the shock's variance (risk) makes the Boss less choosy: Her delegation set expands.*

Extensions of A'

1. Imperfectly Competent Agents, Revisited. Given model A's simple technology, the boss orders the agents' proximity in policy space just as she does in outcome space. Thus, again, "ally" has a clear meaning, so if the boss delegates, then she will pick from the efficient set of subordinates (as this set was defined for model A, extension 4b).

Even some of the efficient agents may be unacceptable. Model A tells us that every agent who would be in the delegation set if perfectly competent has a threshold level of competence that is necessary and sufficient for that agent to be acceptable. Further, in model A' (but *not* A), this competence threshold has a simple monotonicity property: The closer an agent is to the boss, the lower the threshold. This captures a trade-off between ideological similarity and expertise.

Conclusion: *Greater preference similarity permits less competent agents to be acceptable.*

This trade-off often confronts, for example, a revolutionary regime's choice of an officer corps immediately after a revolution.

2. Reasserting Control and Costly Information-Gathering, Revisited. If the boss can reassert control after seeing whether an agent gathered information, then agents know that delegation is reversible: An agent keeps authority only if he specializes. Thus, if acquiring information is costly, the boss's ability to re-take control gives her another way to create incentives for agents.

Consider the simple case: Becoming informed is equally costly for all agents. Then for moderate costs of specialization the boss will *not* use the ally principle. The reason: For moderate c 's, agents who are allies will not specialize. (Ideologically allied agents don't mind the boss's induced ideal policy, so they free-ride on her

uninformed choice.) But because the boss's ignorant choice hurts moderate outliers more significantly, they preempt this outcome by specializing.

Of course, the boss never delegates to extreme outliers. Hence, agents who are either very similar to or very different from the boss will not be picked. (In two dimensions the delegation set is a doughnut—a ring around a hole.) In a further extension, the boss could specialize, at a cost c_0 . Bendor and Meirowitz (2003a, 2003b) show that the delegation set can have a hole (exclude allies) here too. But when the boss's cost of specializing is less than the agents', the delegation set is empty; the boss must get the information herself or go without.

Conclusion: *If the boss can reassert control and information-gathering is costly, then the ally principle fails: Nearby agents aren't selected because they won't specialize.*

3. Multiple Bosses, Revisited. Suppose, as before, that multiple bosses (e.g., legislators) must vote on whether to delegate and, if so, to whom. We assume that when voting on delegation versus control they use any weakly Paretian voting rule: The entire coalition of principals is decisive. We do not specify what agenda and voting system are used to select an agent (after delegating) or a policy (following control).

The following matched pair of results shows that what really matters is the *relative* magnitudes of two fundamental parameters: the size of the shock and the distances among players' ideal points.²⁹ The former measures how hard it is, *ex ante*, for principals to control the outcome by themselves; the latter, their ideological conflicts. These results also show that the intuition underlying comparative static (2)—how a mean-preserving spread of the noise can reduce the expected utility of not delegating—generalizes to collective choice settings.

Conclusion: *(1) If utility is unbounded from below and preferences are held constant, then for a sufficiently big mean-preserving spread of the shock, every boss strictly prefers delegating, to any agent, over any choice of not delegating. (2) If the shock's distribution is held constant and the bosses' and agents' ideal points are sufficiently close together, then every boss strictly prefers delegating, to any agent, over any choice of not delegating.*³⁰

Part 2 of the conclusion has implications for the analysis of divided government in separation-of-power systems. It states that delegation is more likely the more similar are decision makers' preferences (for a fixed

²⁹ Several scholars working on the canonical model have noted this intertwining of ideal point distances and the shock's variance (Epstein 1998; Gilligan and Krehbiel 1987). However, they stressed how a risk-averse principal gains from risk reduction. In contrast, Model A' makes no use of those properties; what counts is not uncertainty but distance from bliss points.

³⁰ [Bendor and Meirowitz 2003a, 2003b] shows that both results can be strengthened: If the bosses' strategies satisfy a simple weak dominance condition, then delegation occurs in every subgame-perfect equilibrium.

distribution of shocks). So on average one would expect unified governments to delegate more.

These results are much stronger than merely saying that the bosses' Pareto set consists only of choices involving delegation. Under either condition, *every* delegation choice—authorizing any agent to implement policy—is unanimously preferred to *any* choice in which the bosses retain control. (Typically, of course, the bosses will disagree about *which* agent should be picked. But Bendor and Meirowitz [2003a, 2003b] show that if a reasonable dominance criterion guides the bosses' voting behavior, then they will vote unanimously to delegate.)

As with a single boss, what is central here is how far—in *any* direction—outcomes are from ideal points. When shocks are big, compared to the ideological bias produced by delegating, holding the policy reins is worse for all bosses than is delegating to a biased expert who will know the state of the world before selecting a policy. This jibes with the traditional public administration view of delegation: When uninformed principals cannot control outcomes within “tolerable” limits, then giving authority to informed agents is best. Thus, that the classical view slighted risk aversion was appropriate: *All* rational bosses prefer to delegate in the above circumstances. Risk aversion is not driving this response.

The conclusion's two parts are closely related to Epstein and O'Halloran's (1999, 75) other two propositions. Part 1 is similar to their, proposition 3, which states, “The closer the preferences of the president to those of the median floor voter, the more likely Congress is to delegate authority to the executive.” Part 2 of our conclusion is similar to their proposition 2: “The more uncertainty associated with a policy area the more likely Congress is to delegate authority to the executive.”

However, our model does not merely reproduce their reasonable results; it generalizes them significantly. Because their model presumes unidimensional policy and outcome spaces, the delegation decision is reduced to the choice of a unitary actor, the floor median. In contrast, because our model allows for multidimensional spaces, it confronts the possibility of multiple bosses and ensuing collective choice issues.

MODEL A''

We have argued that the modern perspective on delegation has overestimated the importance of risk reduction and its partner, risk aversion. But we have *not* claimed that they do not matter. That would be silly. Just as most people want to reduce risk when preferences are insatiable (most humans prefer more money to less but they also buy insurance), so too it can matter for the satiable preferences represented by bliss point utility functions. To show how risk reduction and risk aversion affect delegation choices, we now assume that utility functions are strictly concave. To allow us to use the Arrow–Pratt measure of absolute local risk aversion on the loss function (and so to do comparative statics on

the *degree* of local risk aversion), we assume that they are twice-differentiable.³¹ The literature's standard—the quadratic—is just one example of this class of utility functions.

Now that all parties are risk-averse an efficiency result is clear.

Conclusion: *If everyone is risk averse, not delegating is inefficient.*

Everyone would be better off if the boss delegated to an agent who implemented the boss's ideal outcome, but promises to do that are not credible. The agents' inability to commit harms all.

Comparative Statics: A More Risk-Averse Boss

Conclusion: *As the boss becomes more locally risk averse everywhere (i.e., the Arrow–Pratt measure of absolute risk aversion increases), the delegation set expands.*

This is intuitive: The more averse one is to risk, the more one likes certain outcomes—which an informed agent can provide.

Extensions of Model A''

1. Multiple Bosses with a Status Quo. Suppose that a status quo policy exists, and if the bosses cannot agree on an agent, then they (ignorantly) implement the status quo policy. Define the *collective* delegation set as the set of possible agents who could be picked (under a fixed voting rule) by some winning coalition, given the status quo.

Conclusion: (1) *If the rule is weakly Paretian and the status quo policy is in the bosses' Pareto set, the collective delegation set always exists.* (2) *For any weakly Paretian voting rule the collective delegation set expands as the bosses become more locally risk-averse everywhere.*

These results extend a main finding of Gilligan and Krehbiel (1987)—the informational rationale for delegation, given risk-averse principals—to multidimensional policy spaces, arbitrary (symmetric, mean-zero) perturbations, all strictly concave Euclidean utility functions, and any number of bosses. Their intuition is simple. Delegating to an agent whose ideal point is the reversion policy's expected outcome yields that expected outcome with certainty. Thus, all risk-averse bosses prefer delegating to agents with bliss points close to the reversion policy. Reducing uncertainty is then a collective good, just as Gilligan and Krehbiel argue, *and*

³¹ Given a twice-differentiable loss function, $h(\bullet)$, the Arrow–Pratt coefficient is the negative of the ratio of the second and first derivatives of $-h(\bullet)$. For two utility functions on the same domain, one has a higher coefficient than the second if the first's is larger for every value in the domain. So agent 1 is locally more risk-averse than agent 2 if the negative of agent 1's loss function is more concave (Mas-Colell, Whinston, and Green 1995).

for exactly the reason they specify: Even if the delegation outcome is no closer to any principal's ideal point than is the expected outcome under no delegation, all the bosses *ex ante* prefer the former to the latter because they dislike risk.

However, a collective delegation set does *not* exist in general in model \mathbf{A}' . (It doesn't, e.g., if the bosses' utility functions are piecewise convex [local risk acceptance], the shock's support is sufficiently small, voting is by unanimity, and the status quo policy is in the Pareto set.) Hence, although risk reduction and risk aversion are not *central* to delegation, at times they are important—and our nested models' sequence enables us to pin down when and how they matter.

2. Increased Political Uncertainty. Principals rarely know agents' ideal points with certainty. Suppose that the boss can only *estimate* them (albeit in an unbiased way). Then delegating entails a lottery (Bawn 1995), and the expected utility of delegating falls as this uncertainty increases. If this political uncertainty is sufficiently great, then the delegation set is empty.

CONCLUSIONS

We try to give the political theory of delegation a foundation wider than the canonical model's, thereby linking modern theory to classical perspectives. Following the latter, we have argued that bosses delegate primarily to make good outcomes more likely and bad ones less so; reducing risk is a second-order concern. Hence, whenever possible we have analyzed delegation only through the lens of spatial preferences (models \mathbf{A} and \mathbf{A}'), assuming risk aversion only in the most specialized model (\mathbf{A}''). This nested approach lays bare which premises drive different results.

To take an important example, the ally principle—an ancient rule of political delegation—holds in the very general setting of model \mathbf{A} . Yet it does not always hold in extensions of the models. When subordinates are not perfectly competent the principle may fail at the level of model \mathbf{A} , and if the imperfection is heterogeneous, then it may fail in any of the models. Further, if information is costly, precommitment is possible, or *ex post* controls are present, the ally principle need not hold in any of the models.

Our nested approach also enables us to get a better grasp of the canonical model's major premises and how much they matter. In particular, nesting the models showed that risk aversion (more generally, the shape of loss functions) is of only secondary importance; in contrast, the policy technology—especially its relation to uncertainty—turned out to be critical.

Regarding extensions, the analysis of institutions with multiple bosses, various control schemes, and exogenous commitment generates insights about the choices of political superiors and the relevance of institutional devices. Surprisingly, these insights do not come at the price of generality: They stem from models that are *less* restrictive than the canonical delegation model.

At the end of the day, however, our analysis provides good news for modern (game theoretic) studies of delegation: Much of the canonical model's intuition survives in this paper's more general settings, and many types of extensions turn out to be tractable.

This project—trying to see the heart of delegation by stripping away inessentials—is not done; the political theory of delegation can rest on a still wider base. We have relied on Euclideanism (symmetric preferences) to keep the analysis simple and to make key notions (e.g., “ally”) intuitively meaningful. But this reliance springs not from substantive concerns but from modeling's craft demands. A model yet more general than \mathbf{A} could be built, discarding Euclideanism and assuming only that preferences are (strictly) single-peaked. Such a model would help us see more clearly which properties of delegation depend only on bliss points. Further, because such preferences, in contrast to “economic” (insatiable) ones, are quintessentially political, such a model would also lead us to a truly political theory of delegation.

APPENDIX

A full appendix with proofs of all the results is given by Bendor and Meirowitz (2003b).

Model A

Assume that preferences are representable by utility functions $u_i(x) := h(\|x - y_i\|)$, where y_i in the interior of X is agent i 's ideal point and $h: \mathbb{R}^1 \rightarrow \mathbb{R}^1$ is a strictly decreasing continuous function. Without loss of generality, assume that $h(0) = 0$. We also assume that no decision maker other than the boss has $y_m = 0$. **Perfect shock absorption** means that for all $\varepsilon' \in S$ and $x' \in X$, there exists a unique $p \in P$ s.t. $x' = g(p, \varepsilon')$. We use superscripts to denote coordinates of a vector. Conditional on a policy p the outcome x is a random variable with distribution $F(x | p)$. We assume that the integrability conditions (1) $\int |x^j| dF(x | p) < \infty$ for all j and all $p \in X$ and (2) $\int h(\|x - y_i\|) dF(x | p) > -\infty$ for all $y_i \in X$ and all $p \in X$ are satisfied. We also assume that for every $y_m \in X$ a solution to the problem $\max_{p \in P} \int h(\|x - y_m\|) dF(x | p)$ exists. Finally, we assume that $\|y_m\| < \|y_{m+1}\|$.

Extensions of Model A: (5) Agents Can Precommit.

Preferences satisfy **diversity** if either (i) there does not exist a vector $s \in X$ s.t. for all $m \in \mathbf{M}$, $x_m = \lambda_m s$ for some $\lambda_m \in \mathbb{R}^1$, or (ii) such a vector does exist then there must be two agents $m, m' \in \mathbf{M}$ s.t. for the λ_m and $\lambda_{m'}$ described above $\lambda_m < 0$ and $\lambda_{m'} > 0$.

Proposition 1. *If there are at least two agents in \mathbf{M} and preferences satisfy diversity, then in every subgame-perfect equilibrium at least two agents in \mathbf{M} offer the outcome 0 (the boss's ideal point) and the boss accepts one such offer. Moreover, all such equilibria yield the boss the same utility.*

Proof. By way of a proof by contradiction, assume that diversity is satisfied and the claim does not hold. This implies that either at most one agent promises $x_m = 0$ or at least two agents make this offer and the boss does not delegate to one of these agents. If the latter occurs, then with positive probability the final outcome is not 0, but since a lottery that assigns probability 1 to the final outcome 0 maximizes the boss's

expected utility and is available to the boss, this cannot occur in equilibrium. This implies that it must be the case that we have a subgame-perfect equilibrium in which at most one agent announces $x_m = 0$. There are two subcases to consider: (a) No agent offers $x_m = 0$ and (b) exactly one agent offers $x_m = 0$. We treat these cases in turn.

(a) Let $x' = \arg \min_M \|x_m\|$. There are three separate sub-sub cases that we consider to derive the contradiction for case a. (1) If the boss does not accept x' , which is the best available offer, then it must be the case that no offer is in D , delegation does not occur, and the final outcome is nondegenerate. But a deviation by any agent m to offer the policy $(h^{-1}(\underline{v})/\|y_m\|)y_m$, which is in the delegation set D and which he prefers to the lottery induced by the boss's uninformed best response, dominates m 's offer in the conjectured equilibrium, contradicting the assumption that no offer is in the delegation set in an equilibrium. (2) If the boss does accept an offer (which must be x' by subgame perfection) and part i of diversity is satisfied, then we know that the set $\{z \in X : \|z - y_m\| < \|x' - y_m\| \text{ and } \|z\| < \|x'\|\}$ is not empty. This is the set of points that both the boss and m prefer to the policy x' . Since this set is nonempty there exists a deviation by m that dominates her offer under the conjectured equilibrium. (3) If the boss does accept an offer (x') and part i of diversity is not satisfied but part ii is, then the result follows from the median voter theorem.

(b) If b attains, then we know that the boss will accept this offer. However, since $y_m \neq 0$ and $\min_{n \in M} \|x_n\| := \gamma > 0$, a deviation by agent m to, say, $x'_m = (\gamma/2 \|y_m\|)y_m$ results in delegation by the boss to m but a final outcome that m prefers to the outcome 0. This contradicts the fact that $x_m = 0$ is a best response for m when no other agent offers 0. ■

Model A'

Assume, additionally, that (1) $S \subset P = X = \mathbb{R}^n$, (2) $g(p, \varepsilon) = p + \varepsilon$, and (3) the distribution $F^j(\varepsilon^j)$ of each coordinate of ε is symmetric about 0.

Comparative Statics for A': (2) More Variable Shocks. Random variable a with distribution function $A(\cdot)$ is a **mean-preserving spread** of b having distribution $B(\cdot)$ if (i) $\int a dA - (a) = \int b dB(b)$ (we denote this mean z) and (ii) for any $\delta > 0$, $\int_{\{\|a-z\|>\delta\}} dA(a) \geq \int_{\{\|b-z\|>\delta\}} dB(b)$ with strict inequality if the left-hand side is not zero. Define the real-valued random variables $a^* := \|a - z\|$ and $b^* := \|b - z\|$.

Lemma 1. *If random variable a with distribution function $A(\cdot)$ is a mean-preserving spread of b having distribution $B(\cdot)$, then the distribution functions $A^*(\cdot)$ of a^* and $B^*(\cdot)$ of b^* satisfy the ordering $A^*(\cdot) \leq B^*(\cdot)$ with strict inequality on the interior of the union of the supports of a^* and b^* .*

Proof. Assume that a is a mean-preserving spread of b . By definition we have $\int_{\{\|a-z\|>\delta\}} dA(a) \geq \int_{\{\|b-z\|>\delta\}} dB(b)$ for any $\delta > 0$ with strict inequality if the left-hand side is nonzero. But by definition $A^*(\delta) = 1 - \int_{\{\|a-z\|>\delta\}} dA(a)$ and $B^*(\delta) = 1 - \int_{\{\|b-z\|>\delta\}} dB(b)$ for $\delta \geq 0$. Thus $A^*(\cdot) \leq B^*(\cdot)$ with strict inequality on the interior of the union of the supports of a^* and b^* . ■

Proposition 2. *Assume the assumptions of A', that a is a mean preserving spread of b and that under either lottery $-z$ is an optimal policy for the boss, then under random shock a the delegation set is a strict superset of the delegation set for shock b .*

Proof. By assumption $p_a^* = p_b^* = -\int a dA(a) = -\int b dB(b) = -z$. Thus, the expected utility to the boss from not delegating in each case is $\int h(\|a - z\|) dA(a) = \int h(a^*) dA^*(a^*)$ and $\int h(\|b - z\|) dB(b) = \int h(b^*) dB^*(b^*)$. But, by the previous lemma we have $A^*(\cdot) \leq B^*(\cdot)$ with strict inequality on the interior of the union of the supports of a^* and b^* . Since $h(\cdot)$ is strictly decreasing it is well known that a (first order) stochastically dominated random variable is preferred to the random variable that dominates it, so that we have $\int h(a^*) dA^*(a^*) < \int h(b^*) dB^*(b^*)$. But, this means that the boss's expected utility from not delegating is lower under a than under b . Since the utility to delegation is the same under a and b the result follows. ■

Extensions to A': (3) Multiple Principals. We denote the set of bosses by \mathbf{L} . Assume the following (1) In the initial period the bosses simultaneously vote whether or not to delegate. A weakly Paretian voting rule is used to determine whether or not delegation occurs. (2) If the decision is not to delegate, then a finite period Baron-Ferejohn bargaining game ensues in which a policy $p \in P$ is chosen without knowledge of ε . (3) If delegation is chosen, then a finite period Baron-Ferejohn bargaining game ensues over to which agent to delegate.

Lemma 2. *The subgame that begins with a decision to delegate has a subgame-perfect Nash equilibrium in weakly undominated strategies.*

Proof. We establish the claim by backward induction. At any history in which delegation has occurred to agent m and ε is revealed, the best response for agent m is to select policy $y_m - \varepsilon$. Given this, the decision to delegate to any m results in a degenerate lottery of attaining outcome y_m with probability 1. Since the bosses have complete, reflexive, and transitive preferences over this set of outcomes, the finite bargaining game over these outcomes has a subgame-perfect equilibrium in weakly undominated strategies. ■

Given a multiple-boss game, a strategy profile satisfies **condition α** if at any information set in which any agent i has a binary choice between action a and action b in which the expected utility resulting from every profile of strategies that involve the choice of a by agent i is less than the expected utility resulting from the play of best responses by agents at histories following i 's choice of b , agent i plays action b .

Proposition 3. *For fixed boss and agent ideal points, there exists a distribution $F(\cdot)$ of the shock ε for which any strategy profile satisfying condition α results in unanimous votes to delegate by the bosses.*

Proof. The proof is constructive. Given fixed boss and agent ideal points and arbitrary distribution $F(\cdot)$ we consider L translated games constructed as follows. For boss $l \in \mathbf{L}$ the l -translated single-boss game has shock $\varepsilon^l = \varepsilon - y_l$, agent ideal points $y_m^l = y_m - y_l$ for each $m \in M$, and boss ideal point $\mathbf{0}$. Clearly, each l -translated game satisfies the assumptions of the basic game of model A'. In each l -translated game either all agent ideal points are in the delegation set or they are not. If not, then by the risk-comparative static we can take mean-preserving spreads of ε^l until the delegation set covers all agent ideal points in the l -translated game. Thus, a mean-preserving spread of ε^l exists for which all agent ideal points are in the delegation set. Denote this shock ε^{l+} . Now, in the multiple-boss game with random shock $\varepsilon^{l+} + y_l$ the boss with ideal point y_l prefers to delegate to any

agent than to select policy herself (facing uncertainty about shock $\varepsilon^{l+} + y_l$). Repeating these operations for each $l' \in \mathbf{L}/l$ we can construct mean-preserving spreads of $\varepsilon^{l+} + y_l$ such that every boss in \mathbf{L} prefers delegating to any agent in \mathbf{M} to selecting policy herself (under uncertainty). Since the best expected utility that a boss can attain in the nondelegation subgame is from selecting policy herself, this implies that if a strategy profile satisfies condition α , then delegation must occur. ■

REFERENCES

- Baron, David P., and Adam Meirowitz. 2002. "Relations between Fully-Revealing Equilibria of Multiple-Sender Signaling and Screening Models." Typescript. Princeton University.
- Battaglini, Marco. 2002. "Multiple Referrals and Multidimensional Cheap Talk." *Econometrica* 70 (4): 1379–1401.
- Bawn K. 1995. "Political Control Versus Expertise: Congressional Choices about Administrative Procedures." *American Political Science Review* 89 (March): 62–73.
- Bendor, Jonathan, Amihai Glazer, and Thomas H. Hammond. 2001. "Theories of Delegation." *Annual Review of Political Science* 4: 235–69.
- Bendor, Jonathan, and Adam Meirowitz. 2003a. "Spatial Models of Delegation." *Graduate School of Business Working Paper Series*, Stanford University.
- Bendor, Jonathan, and Adam Meirowitz. 2003b. "Supplementary Appendix to Spatial Models of Delegation." Typescript Princeton University.
- Epstein, David. 1997. "An Informational Rationale for Committee Gatekeeping Power." *Public Choice* 91 (3–4): 271–99.
- Epstein, David. 1998. "Partisan and Bipartisan Signaling in Congress." *Journal of Law, Economics, and Organization* 14 (October): 183–204.
- Epstein, David. 1999. "Legislating from Both Sides of the Aisle: Information and the Value of Bipartisan Consensus." *Public Choice* 101 (1–2): 1–22.
- Epstein, David, and Sharyn O'Halloran. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers*. Cambridge: Cambridge University Press.
- Ferejohn, John. 1986. "Incumbent Performance and Electoral Control." *Public Choice* 50 (1–3): 5–25.
- Gilligan, Thomas, and Keith Krehbiel. 1987. "Collective Decision-Making and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures." *Journal of Law, Economics, and Organization* 3 (Fall): 287–335.
- Gilligan, Thomas, and Keith Krehbiel. 1989. "Asymmetric Information and Legislative Rules with a Heterogeneous Committee." *American Journal of Political Science* 33 (May): 459–90.
- Gilligan, Thomas, and Keith Krehbiel. 1990. "Organization of Informative Committees by a Rational Legislature." *American Journal of Political Science* 34 (May): 531–64.
- Hamilton, Alexander, James Madison, and John Jay [1788] 1961. *The Federalist: The famous papers on the principles of American government*, ed. Benjamin F. Wright, New York: Metrobooks.
- Huber, John, and Charles Shipan. 2002. *Deliberate Discretion? The Institutional Foundations of Bureaucratic Autonomy*. Cambridge: Cambridge University Press.
- Krishna, Vijay, and John Morgan. 2001. "Asymmetric Information and Legislative Rules: Some Amendments." *American Political Science Review* 95 (June): 435–52.
- Mas-Colell, Andreu, Michael Whinston, and Jerry Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.
- McCubbins, Mathew, and Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols versus Fire Alarms." *American Journal of Political Science* 2 (February): 165–79.
- McNollGast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law, Economics and Organization* 3 (Fall): 243–77.
- Simon, Herbert, Donald Smithburg, and Victor Thompson. 1950. *Public Administration*. New York: Knopf.
- Wilson, James Q. 1989. *Bureaucracy: What Government Agencies Do and Why They Do It*. New York: Basic Books.