Giving Advice *vs.* Making Decisions: Transparency, Information, and Delegation[*]

Sean Gailmard[†] John W. Patty[‡]

January 12, 2018

Abstract

We generalize standard delegation models to consider policymaking when both information and authority are dispersed among multiple actors. In our theory, the principal may delegate *partial* authority to a privately-informed agent while also reserving some authority for the principal's use after observing the agent's decision. Counterintuitively, the equilibrium amount of authority delegated to the agent is *increasing* in the preference divergence between the principal and agent. We also show that the amount of authority delegated depends upon whether the agent can observe the principal's own private information (a condition we refer to as "top-down transparency"): this form of transparency *increases* the authority that must be delegated to the agent to obtain truthful policymaking. Accordingly, such transparency can result in less-informed policymaking. Nonetheless, the principal will sometimes but not always voluntarily choose such transparency.

[*]We gratefully acknowledge multiple helpful conversations with Maggie Penn and feedback from seminar audiences at the University of Chicago, University of Pittsburgh, University of Michigan, New York University, and Washington University in St. Louis. All remaining errors are our own.

[†]Charles & Louise Travers Department of Political Science, University of California, Berkeley. Email: *gailmard@berkeley.edu*.

[‡]Department of Political Science, University of Chicago. Email: *jwpatty@uchicago.edu*.

In bureaucratic organizations, conflict over policy goals creates incentive problems that undermine information aggregation and, ultimately, informed policy making (*e.g.*, Hammond and Miller (1985), Hammond and Thomas (1989), Miller (1992), Ting (2002, 2003, 2008, 2011)). Strict hierarchies, requiring subordinates to funnel information to hierarchical superiors, make this problem particularly intense (*e.g.*, Crawford and Sobel (1982)). Therefore, greater information can usually be extracted by lodging decision authority firmly and exclusively in the agents with the requisite information (*e.g.*, Holmström (1984), Gilligan and Krehbiel (1987), Epstein and O'Halloran (1999), Dessein (2002), Bendor and Meirowitz (2004), Gailmard and Patty (2012*a*,*b*)). A natural precept of organizational design is then to delegate decisions to agents with the relevant information to make them.

However, in complex bureaucratic organizations, this precept may be too simplistic. It may be infeasible or impracticable to concentrate both a fixed set of decisions and all information relevant to them in a single individual or unit. The resulting set of decisions may be too large, or it may be impossible to confine information in this way. Instead, both information and authority over smaller decisions are often dispersed among multiple agents. In exercising its delegated authority, an agent may then reveal information relevant to the decisions of other actors—and the agent may have preferences over those decisions as well. This in turn can affect their incentives to reveal their information through their choices in the first place. In this paper, we develop a model to explore these incentives, and in turn the optimal design of power- and information-sharing arrangements.

More specifically, we consider a model of policy making in which decentralized policy-relevant information is possessed by two or more individuals, at least one of whom is authorized to make policy decisions that affect all of the individuals. In this setting, we first consider the incentives for these individuals to share their information with one another as a function of the number of and preference similarities between the agents. We then consider the possibility that the individual possessing the authority (the "principal") to make policy decisions may delegate some portion (or all) of this authority to one or more other individuals (the "agents") in pursuit of eliciting these agents' private information. We show that such power-sharing can occur in equilibrium for moderate preference disagreement, but that no power is shared with any agent whose preferences are either sufficiently divergent or sufficiently similar to the principal's. Finally, we consider the impact of information-sharing by the principal—a practice we refer to as "top-down transparency"—on both the incentive of the agents to reveal their information to, and the degree of power shared by, the principal.[1]

As this description suggests, our model is closely related to the now vast literature on delegation. In the "standard" delegation model, the principal chooses a set of policies, from which the agent then freely chooses in turn.[2] In other words, existing models typically presume that exactly one actor will eventually make the final policy choice, and then consider which actor the principal prefers this to be.[3] We refer to this aspect of the standard model as "exclusive delegation."

Our model extends the exclusive delegation so that more than one actor may (and in equilibrium, often does) contribute meaningfully to a policy choice, which we refer to as "partial delegation." The literature on optimal delegation (*e.g.*, Alonso and Matouschek (2008) and Gailmard (2009)) allows the principal to constrain the set of actions available to the agent (*i.e.*, choose a "delegation set"). Our framework does not allow the principal to restrict the set of actions from which the agent may choose. Rather, the principal chooses the degree to which the agent's policy choice will affect the payoffs of *both* the principal and agent.[4] One advantage of this conception of partial delegation is that it is simpler to discuss when the principal has delegated "more" or "less" authority to an agent: doing so in the standard models of exclusive delegation require that one compare delegation sets, which need not be unambiguously ordered.

Substantively, the partial delegation framework reveals a novel insight about the incentive effects of delegation. The classic story in the analysis of formal vs. real authority (Aghion and Tirole (1997)), derived from the exclusive delegation model, is that delegation of authority to the agent induces the agent to exert effort beneficial to the principal. In our partial delegation model, delegation of more authority raises the cost to the agent of dissembling to the principal about her signal. In

[1]As we make clear below, one practical instantiation of top-down transparency is active, independent, and on-going policymaking by the principal that is observable by the agent. Thus, the effects of this informational arrangement indicated by our theory extend to considerations of the implications of what one might call "active oversight" by the principal, even if not directed in a punitive or auditing fashion at the agent, *per se*.

[2]This includes, as a special case, models of delegation as an all-or-nothing choice, *i.e.*, the principal ether makes policy on his or her own or grants the agent plenary control over policy.

[3]By "final policy choice," we mean a choice of policy that affects all actors' payoffs.

[4]We discuss how to interpret our notion of partial delegation below (p. 9).

order to mislead the principal about her signal, an agent with more authority must take an action she herself believes to be wrong on a larger set of issues. To our knowledge, this is a new idea in the formal theory of delegation of authority.[5]

Partial Delegation and Signaling. Our analysis is entirely based on the signaling framework developed in Galeotti, Ghiglino and Squintani (2013). We extend the results of Galeotti, Ghiglino and Squintani (2013) by focusing on institutional design. That is, for any given choice of "partial delegation," the principal's and agent's equilibrium beliefs and behaviors are those identified by Galeotti, Ghiglino and Squintani (2013): the new theoretical innovation in this article is in deriving the principal's preferences (as induced by equilibrium behavior) over different degrees of partial delegation.[6]

The partial delegation approach is more applicable than models of exclusive delegation to real-world policy making, in which various different agents each unilaterally render policy decisions at various times. In such processes, the principal and agent(s) each exert/impose independent policy effects through these various decisions. Whereas models of exclusive delegation make sense for processes in which legal authority can be delineated and handed over to another actor, our partial delegation framework is more relevant to decision making when "public policy" is (as in its everyday meaning) not a simple unitary decision, but rather the amalgamation of various implementation and administrative decisions made by various actors in, and possibly outside, government. Thus, an important goal of this article is to extend the theory of delegation to these types of situations.

The partial delegation framework allows us to consider a principal's incentives to delegate a *portion* of policymaking authority to an agent, with the attendant possibility of learning that agent's information for later unilateral use by the principal.[7] Models of exclusive delegation consider the

[5]We thank anonymous reviewers for drawing our attention to the points in this and the previous paragraph.

[6]By focusing on institutional design within the Galeotti, Ghiglino and Squintani (2013) framework, this article is also related to those of Dewan and Squintani (2012) and Penn (2015). The principal formal innovations in this article are the allowance for a continuum of partial delegation decisions, and the consideration of the top-down transparency of the policy-making process.

[7]Our formulation of partial delegation is formally equivalent to "probabilistic delegation." Specifically, one can interpret the principal delegating a portion, $\alpha \in [0,1]$, of his or her authority to the agent as committing to implementing the action recommended by the agent with probability α and otherwise (with probability $1-\alpha$) retaining policymaking authority and making his or her preferred policy choice based on the information provided by the agent's recommendation. We are comfortable with either interpretation of the model.

possibility of granting control over policy so that the agent has an incentive to utilize its information (Aghion and Tirole (1997), Dessein (2002)), but neglect that the principal may be able to use information so revealed in other aspects of a policy choice.[8] In our setting, that possibility is central. As a result, exclusive delegation, while possible in our model, is almost never optimal.[9]

Partial delegation leads naturally to a focus on top-down transparency: the extent to which the principal should reveal its information to agents. This is typically irrelevant in standard delegation models. First, the principal in these models is typically assumed to have no information that the agent does not also have. Second, if the principal did have such information, she should simply reveal it to the exclusive decision maker because this would result in more informed (hence, better for all actors by assumption) decisions. From a vaguely Weberian perspective, top-down transparency seems not just irrelevant but clearly contrary to the principal's interest: if knowledge is power, then to share the one is to give away the other. Under partial delegation, top-down transparency is more subtle and, surprisingly, possibly perverse. It can mitigate the incentive for the agent to reveal his private information, in an attempt to manipulate the choices ultimately made by the principal. Thus, when power is shared between the principal and the agent in the sense of partial delegation, the ultimate welfare impact of imposing partial transparency of information can be counterintuitive. In a nutshell, *top-down transparency (weakly) increases the minimal level of authority that the principal must grant his or her agent in order to induce that agent to reveal/utilize his or her own private information when making policy.*

We now turn to describing the model.

1 The Model

We consider a model of decentralized decision-making with asymmetric and decentralized information. There are two players, the Principal P and Agent A.[10] Each actor $i \in \{P, A\}$ will ultimately

[8] But see Callander (2008), who presents a novel model of expertise in which this "invertibility" of an agent's information from his or her policy choice is not possible.

[9] Specifically, exclusive delegation is optimal within the partial delegation framework only if the agent has exactly the same preferences as the principal.

[10] We briefly remark on an n-agent extension at the conclusion of the analysis.

make a unidimensional policy choice, $y_i \in \mathbf{R}$, and then receive a payoff, u_i, that depends on the policy chosen and a latent *state of nature*, θ, which we presume to be a number between 0 and 1. Each player has some policy-making authority, which is decentralized in the sense that each player's policy decision is made unilaterally.[11] Information is decentralized in the sense that each individual i *privately* observes a partially informative signal about θ. The "policy-making process" is a sequence of observable policy choices by A, followed by P, such that P can (possibly) make inferences about A's signal. This structure is displayed in Figure 1.

The interesting tensions (and strategic incentives) in the model arise from the combination of two facts: first, both P and A desire information about the state of nature θ, and, second, each player's payoff is potentially affected by the choices of the other. Of course, the first fact—that there is uncertainty about a policy-relevant state—is what makes information aggregation welfare enhancing. The second fact is what potentially undermines the achievement of such aggregation. Information aggregation requires P to base its policy decision on the message about θ implicitly conveyed by A's policy choice. However, when P and A have different state-dependent policy goals, this can create an incentive for A not to reveal its information through its policy choice. Alleviating this tension is the essentially the "institutional design" goal in this setting. With an institutional arrangement that satisfies this goal, we can understand the limits of information sharing in decision processes where partial delegation is a possibility.

1.1 Formal Primitives

Prior to making their policy choices, each individual i privately observes a a binary signal (*e.g.*, "low/high" or "bad/good"), denoted by $s_i \in \{0, 1\}$, conditional on θ. We assume[12] that θ is $U[0, 1]$

[11]We set aside issues of collective choice in this article. Including a collectively chosen policy choice represents an obvious extension. In addition to the traditional constraints of time and space, technical concerns also dissuade us from including such an extension in this article. Specifically, incorporating collective choice not only requires that one make modeling choices about the type of collective choice procedures to examine, but also immediately requires that strategic actors condition their vote choice (*i.e.*, their input into the collective choice process) on being pivotal. As is well-known (*e.g.*, Austen-Smith and Banks (1996), Feddersen and Pesendorfer (1996, 1998)), this type of reasoning can quickly become quite complicated. Of interest in such an extension, however, is that there will be situations in which an individual will vote for a policy other than the one he or she would choose on his or her own.

[12]Further details on the information structure are in the appendix.

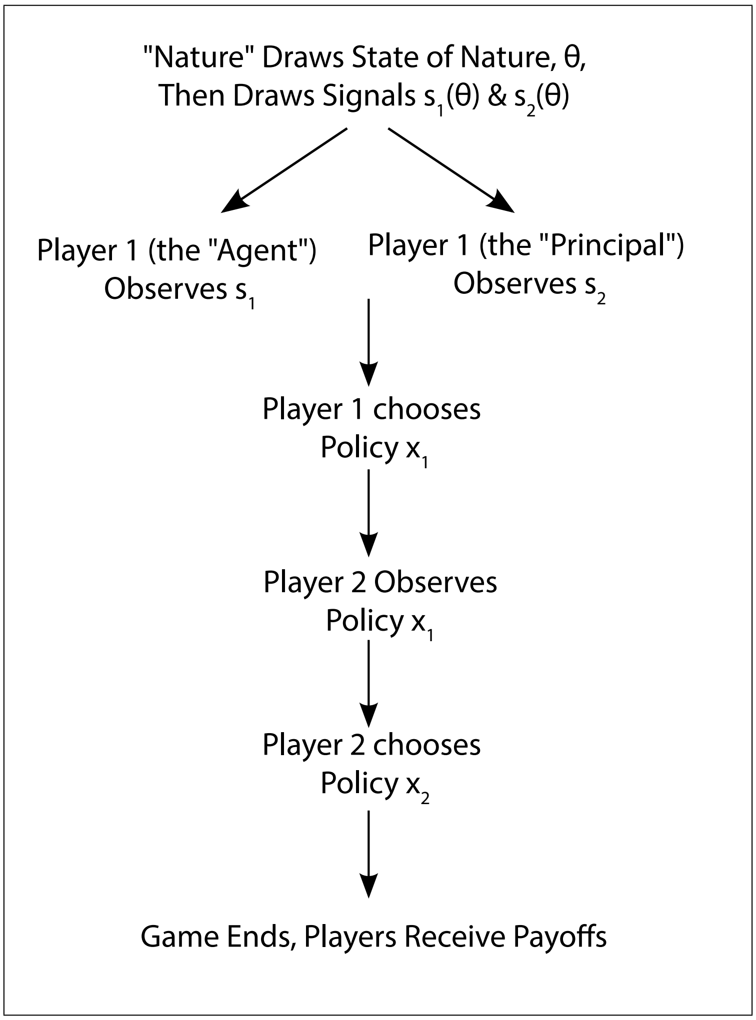

Figure 1: Example of Sequential Policymaking Process

random variable, and given a signal $s_i \in \{0,1\}$, player i's conditional expected value of θ is

$$E(\theta|s_i) = \frac{1+s_i}{3} = \begin{cases} {}^1\!/_3 & \text{if } s_i = 0, \\ {}^2\!/_3 & \text{if } s_i = 1. \end{cases}$$

More generally, upon observing $m > 0$ signals, $\{s_1, \ldots, s_m\}$, with $k = \sum_{i=1}^m s_i$, the conditional expected value and conditional variance of the resulting beliefs are

$$\begin{aligned} E(\theta|k,m) &= \frac{k+1}{m+2}, \text{ and} \\ V(\theta|k,m) &= \frac{(k+1)(m-k+1)}{(m+2)^2(m+3)}. \end{aligned}$$

Intuitively, as the proportion of observed signals that equaled 1 (as opposed to equaling to 0) increases, one's belief about the realized expected value of the state of nature, θ, also increases. In a spatial sense, then, if an individual i observes a signal of 1 ($s_1 = 1$), one thinks that the state of nature (and hence, the optimal policy choice) is "farther to the right" than if i observed a signal equal to 0.

Policy-making. Each player $i \in \{P, A\}$ ultimately chooses a policy $y_i \in \mathbf{R}$, with $y = (y_P, y_A)$. The agent always chooses its policy y_A prior to the choice y_P by the principal. The information available to the agent when choosing y_A depends on whether the policymaking process is characterized by top-down transparency or not. If the process does *not* have top-down transparency—hereafter, the "opaque case"—the agent chooses y_A knowing only his or her private signal, s_A. If, on the other hand, the process does have top-down transparency, then y_A is chosen by the agent with knowledge of both his or her private signal, s_A, as well as that of the signal of the principal, s_P.

Regardless of the top-down transparency of the process, the principal chooses y_P after observing his or her own signal s_P, and the agent's policy choice y_A. Thus, in neither case does the principal ever observe the agent's signal directly. Rather, the principal may be able in equilibrium to infer the agent's signal from the agent's policy choice, y_A.

Payoffs. Each player $i \in \{P, A\}$ has a payoff function of the following form:

$$u_i(y, \theta; \beta) = -\alpha_A(y_A - \theta - \beta_i)^2 - \alpha_P(y_P - \theta - \beta_i)^2, \tag{1}$$

where $\beta_i \in \mathbf{R}$ denotes the policy preference (or bias) of agent i and $\beta \equiv \{\beta_P, \beta_A\}$ denotes the profiles of all biases. We assume throughout that these biases are exogenous and common knowledge to all of the players.[13]

The payoff function in equation 1 is of the classic "quadratic loss" formulation and standard arguments imply that the optimal policy choice for a an individual i who has observed (or inferred) m signals, $\{s_1 \ldots, s_m\}$, with $k = \sum_{j=1}^{m} s_j$, is:

$$y_i^*(k, m) = \frac{k+1}{m+2} + \beta_i. \tag{2}$$

1.2 Shared Authority: The "Partial Delegation" Framework

In the baseline model, the "authority parameters" $\alpha = (\alpha_P, \alpha_A)$, such that $\alpha_P > 0$, $\alpha_A \geq 0$, and $\alpha_P + \alpha_A = 1$, are treated as exogenous.[14] These parameters represent the degree of policy-making *authority* possessed by each player. For each player $i \in \{P, A\}$, the parameter α_i measures the degree to which player i's choice of policy, y_i, will affect the payoff of *both* players. Thus, the decisions of an individual with "greater authority" have greater impact on all players than the choices of an individual with less authority.

This conception of authority differs crucially from the standard delegation literature. Existing models of delegation treat delegation either as an all-or-nothing decision or, in the case of the "discretionary interval" models (Holmström (1984)), as the width of the range of policies that the delegate may choose from. One can think of α_i as representing the proportion of some set of identical decisions that i is allowed to make "on her own."[15] That is, our "partial delegation" representation

[13] All players care about the distance between policy actions of all players and the state of the world. They do not care about the distance between the policy actions themselves. This is in keeping with the standard model of state contingent utilities, though preference for policy consistency is an interesting possibility to explore in future work.

[14] We endogenize the choie of α's below.

[15] Presuming that there are "lots" of decisions to make, we set aside the integer/indivisibility constraints and treat the "unit of α" as infinitely divisible. Relaxing this assumption would not alter any of the substantive conclusions but would

can be thought of as an approximation for "multiple delegations" with two caveats: first, the impacts of the various delegated policy decisions are determined by a common latent state of nature θ, and, second, the total payoff from the various decisions is an additive function of the impacts of the individual decisions. The dependence on a common state of nature roughly approximates the idea that all of the multiple delegated decisions are in the same policy realm, directed at the same purposes. The presumption that the payoff is additive across these decisions is merely for tractability. Further, since each of the multiple decisions is a concave function of the distance between the decision and a fixed θ implies that each player will make each of its decision identically. Thus, for simplicity, we represent the multiple decisions as a single, overall policy decision.

1.3 Sequence of Play: Baseline Game

The policy-making process described above, referred to as the *baseline game*, proceeds as follows:

1. The decision-making authorities of the two players are revealed: $\alpha_P > 0$ and $\alpha_A \geq 0$, with $\alpha_P + \alpha_A = 1$.
2. Nature determines the state of nature θ and players' signals, $s = \{s_P, s_A\}$.
3. - OPAQUE CASE. Each player $i \in \{P, A\}$ privately observes his or her signal, s_i.
   - TOP-DOWN TRANSPARENT CASE. The agent, A, observes his or her signal, s_A, and *both* players observe the principal's signal, s_P.
4. The agent, A, sets policy $y_A \in \mathbf{R}$.
5. Principal P observes y_A and sets policy $y_P \in \mathbf{R}$.
6. Game concludes, players receive payoffs.

In the appendix, we present a formal definition of strategies, beliefs, and equilibrium in these games. Now, however, we turn to the analysis of equilibrium behavior.

greatly increase the notational heft of the model's presentation.

1.4 Decentralized Authority & Communication

We analyze the equilibrium behavior in the baseline model, first considering the "opaque case," in which the agent A observes only s_A prior to making his policy choice, y_A, and then turning to the case of top-down transparency in which the agent A observes the principal's information, s_P, at the same time A observes his own signal s_A, and prior to choosing y_A. We then compare the conditions for information aggregation (*i.e.*, truthful policy choice by the agent) in the two settings.

The Opaque Case: Superior's Information Hidden from Agent. We first suppose that s_P is not observed by A. In this case, if the principal believes that A is being truthful, note that agent A effectively has only two possible policy choices:[16]

$$y_A \in \{ {}^1\!/_3 + \beta_A, {}^2\!/_3 + \beta_A \},$$

where "truthful" policy-making by A involves A setting policy as follows:

$$y_A^*(s_A) = \begin{cases} {}^1\!/_3 + \beta_A & \text{if } s_A = 0, \\ {}^2\!/_3 + \beta_A & \text{if } s_A = 1. \end{cases} \tag{3}$$

Thus, if the principal believes that the agent is truthful when choosing y_A—*i.e.*, that he or she chooses according to (3)—the principal's optimal behavior as a function of the inferred value of the agent's signal, $\tilde{s}_A$, and the principal's own signal, s_P, is:

$$y_P^*(\tilde{s}_A, s_P) = \frac{1 + \tilde{s}_A + s_P}{4} + \beta_P. \tag{4}$$

The behavior described in Equation (4) is the basis of the agent's incentive to possibly set policy insincerely: that is, to *not* behave as described in Equation (3). Specifically, when choosing y_A, the principal's behavior described by (4) illustrates that the agent must consider two potentially

[16] We formally demonstrate in the appendix the principal's beliefs that justify this (Section A.3, Proposition 4), as well as discussing the technical underpinnings of our choice of such beliefs.

countervailing incentives. For example, suppose that $\beta_A > \beta_P$ and A observes $s_A = 0$.[17] In the absence of the principal, P, the agent's optimal choice is clear: $y^*_A(0) = {}^1\!/_3 + \beta_A$. However, if we include the principal and he or she believes that the agent is setting policy truthfully, then choosing ${}^1\!/_3 + \beta_A$ will result in the principal setting y_P lower than he or she would if the agent chose $y_A = {}^2\!/_3 + \beta_A$, which would lead the principal to infer that the agent had observed $s_A = 1$ instead of $s_A = 0$. Because $\beta_A > \beta_P$, the agent always wants to insert a (sufficiently small) positive bias in the principal's decision-making,[18] so that, in some circumstances (depending in this setting, as we will see below, on α and β), the agent will gain more from setting a policy so as to manipulate the principal's choice than the agent will lose from setting policy suboptimally. We now turn to deriving the conditions under which this does *not* hold. In other words, we derive the conditions under which the agent has an incentive to set policy truthfully. These are known as the agent's *incentive compatibility* conditions.

To keep the presentation succinct, we will maintain the supposition that $\beta_A > \beta_P$ and accordingly focus only on the case of $s_A = 0$.[19] The agent's expected payoff from setting y_A optimally such that the principal interprets y_A as implying that A received $s_A = 0$ is:[20]

$$U_A(y^*_A(0); \beta, s_A = 0) \quad = \quad -(1-\alpha_A)(\beta_A - \beta_P)^2 - \frac{\alpha_A}{72} - \frac{1}{24}, \tag{5}$$

which represents the agent's loss from the principal's informed decision-making. The first of the three terms on the righthand side of (5) represent the agent's disutility from the principal's sequentially rational decision-making: this is simply the square of the distance between the agent's and principal's ideal points, weighted by the principal's degree of authority. The second and third terms jointly represent the agent's disutility from his or her residual uncertainty about the state of nature, θ. On the other hand, if the principal believes that the agent is telling the truth, then the agent's

[17] The problem is symmetric for $\beta_A < \beta_P$ and $s_A = 1$.

[18] This ubiquitous incentive is the reason that perfect information transmission is possible in the canonical cheap-talk model of Crawford and Sobel (1982) *only if* $\beta_A = \beta_P$.

[19] It is simple to verify that, if $\beta_A > \beta_P$ (*resp.*, $\beta_P > \beta_A$), the agent always has a strict incentive to behave truthfully following observing $s_A = 1$ (*resp.*, following observing $s_A = 0$).

[20] These expressions are derived in the appendix. Note that A's expected payoff does not incorporate possible errors A makes in inferring the policy action P will take. These do not matter because payoffs depend solely on the distance between the actions and the state.

expected payoff from choosing $y_A = {}^2\!/_3 + \beta_A$ (and leading P to believe that A observed $s_A = 1$ instead of $s_A = 0$) is

$$U_A(y_A^*(1); \beta, s_A = 0) \quad = \quad -(1-\alpha_A)\left((\beta_A - \beta_P)^2 - \frac{\beta_A - \beta_P}{2}\right) - \frac{\alpha_A}{16} - \frac{5}{48}. \tag{6}$$

Again, the first term on the righthand side of (6) represents the agent's disutility from the principal's sequentially rational policy choice. Because $\beta_A > \beta_P$, this disutility is lower than the analogous disutility in (5). However, the disutility represented by the second and third terms—which represent both the agent's residual uncertainty *and the disutility from choosing a policy that is different than the expected value of the state of nature*, θ—is larger in absolute value than that represented by the second and third terms of (5). Thus, the agent's incentive to "tell the truth" after $s_A = 0$ depends on the relative sizes of (1) the difference between the first terms in (5) and (6) and (2) the difference between the sum of the second and third terms of (5) and the sum of those terms in (6).

Formally, comparing (5) and (6) yields the following incentive compatibility condition for the agent:

$$\begin{aligned} U_A(y_A = y_A^*(s_A); \beta, s_A) \quad &\geq \quad U_A^0(y_A = y_A^*(1 - s_A); \beta, s_A), \\ \beta_A - \beta_P \quad &\leq \quad \frac{2}{9}\frac{\alpha_A}{1 - \alpha_A} + \frac{1}{8} \end{aligned} \tag{7}$$

Transparent Case: Agent Observes Superior's Information. Moving to the top-down transparency case, agent A observes the principal's information, s_P, prior to choosing y_A.[21] From an *ex ante* perspective, there are three possible policy choices that agent A will choose with positive probability in equilibrium (because A could observe 0, 1, or 2 successes). These choices are as

[21] A observes P's information perfectly in this case. We do not consider this as resulting from strategic communication by P, though it would be interesting to do so. Our analysis takes a simpler approach by assuming that P is exogenously required to disclose her information truthfully. This is motivated by the substantive context of bureaucratic policy making, in which an organizational principal's information may come from on-the-record communication with experts, whereas staff agents with their own expertise possess nonverifiable information that must be endogenously revealed.

follows:

$$
\begin{aligned}
y_A &= {}^1\!/_4 + \beta_A, \\
y_A &= {}^1\!/_2 + \beta_A, \text{ and} \\
y_A &= {}^3\!/_4 + \beta_A,
\end{aligned}
$$

where "truthful" policy-making by A involves A setting policy as follows:

$$
y_A^*(s_A, s_P) = \begin{cases} {}^1\!/_4 + \beta_A & \text{if } s_A + s_P = 0, \\ {}^1\!/_2 + \beta_A & \text{if } s_A + s_P = 1, \\ {}^3\!/_4 + \beta_A & \text{if } s_A + s_P = 2. \end{cases} \tag{8}
$$

However, once A observes both s_A and s_P, the realization of s_P—which A is aware that P will also observe and condition his or her beliefs on—effectively reduces A's strategic calculation to two choices.[22] Specifically,

$$
\begin{aligned}
y_A &\in \{ {}^1\!/_4 + \beta_A, {}^1\!/_2 + \beta_A \} && \text{if } s_P = 0, \text{ and} \\
y_A &\in \{ {}^1\!/_2 + \beta_A, {}^3\!/_4 + \beta_A \} && \text{if } s_P = 1.
\end{aligned}
$$

Continuing as in the analysis of the opaque case, the expected payoff from setting y_A optimally such that the principal interprets y_A as implying that A received $s_A = 0$ is:[23]

$$
\begin{aligned}
U_A(y_A^*(0, s_P); \beta, s_P, s_A = 0) &= -E_{G(0)}[(1-\alpha_A)(y_P^*(0, s_P) - \theta - \beta_A)^2 + \alpha_A (y_A^*(0, s_P) - \theta - \beta_A)^2], \\
&= -\left[(1-\alpha_A)(\beta_A - \beta_P)^2\right],
\end{aligned} \tag{9}
$$

[22]The details of this restriction are provided in the appendix (Section A.3, Proposition 4).

[23]Note that these expressions omit the conditional variance terms for θ, which are invariant to the choice of y_A and will accordingly cancel in the derivation of the incentive compatibility conditions. Note that this obfuscates the fact that the agent A would prefer to observe P's information prior to setting y_A.

and the expected payoff from setting y_A as if A observed $s_A = 1$ instead of $s_A = 0$ is

$$\begin{aligned} U_A(y^*_A(1,s_P);\beta,s_P,s_A=0) &= -E_{G(0)}[(1-\alpha_A)(y^*_P(1,s_P)-\theta-\beta_A)^2+\alpha_A(y^*_A(1,s_P)-\theta-\beta_A)^2], \\ &= -\left[(1-\alpha_A)({}^1\!/_4+\beta_A-\beta_P)^2+\frac{\alpha_A}{16}\right]. \end{aligned} \tag{10}$$

The Impact of Top-Down Transparency. The impact of top-down transparency arises from the difference between Equations (10) and (6). They differ only in their second righthand side terms: $\alpha_A/9$ in the opaque case and $\alpha_A/16$ in the transparent case. This difference follows from the fact that in the opaque case, the agent must move policy by ${}^1\!/_3$ (from ${}^1\!/_3+\beta_A$ to ${}^2\!/_3+\beta_A$) to successfully change the principal's inference about s_A but, in the transparent case, the agent need move policy only by ${}^1\!/_4$ (either from ${}^1\!/_4+\beta_A$ to ${}^1\!/_2+\beta_A$ if $s_P=0$ or from ${}^1\!/_2+\beta_A$ to ${}^3\!/_4+\beta_A$ if $s_P=1$).

Because the agent's policy payoff is quadratic loss, the direct impact on the agent's expected payoff from deviating $\delta > 0$ away from the agent's subjective expected value of θ is to reduce it by $\alpha_A\delta^2$. Thus, *top-down transparency reduces the direct cost A must pay in order to manipulate P's beliefs*. This difference leads to a more demanding incentive compatibility condition in the transparent case than in the opaque case.

Specifically, comparing (9) and (10), truthful revelation is incentive compatible for A in the top-down transparency setting so long as the following holds:

$$\beta_A-\beta_P \leq \frac{1}{8}\frac{\alpha_A}{1-\alpha_A}+\frac{1}{8}. \tag{11}$$

Comparing inequality (11) with the condition for the opaque case (inequality (7)):

$$\beta_A-\beta_P \leq \frac{2}{9}\frac{\alpha_A}{1-\alpha_A}+\frac{1}{8},$$

it follows that, for a fixed combination of preference biases β_P and β_A and level of discretionary authority α_A, if incentive compatibility holds in the transparent case, it also holds in the opaque case, but the converse does not necessarily hold. This difference implies that attempting to manipulate the principal's choice is "cheaper" for the agent in the transparent case than in the opaque case, a point that plays a central role in the discussion of endogenous transparency (Section 3).

Robustness of the Impact of Top-Down Transparency. Top-down transparency makes incentive compatibility more demanding because of the common knowledge that the agent has more information than in the opaque case, which implies that he or she can credibly manipulate the principal's beliefs with a smaller deviation. This effect can be generalized in several ways. For example, this effect continues to hold if the agent and/or principal observe more than a single signal and is robust to assuming that the prior distribution of θ is $\mathrm{Beta}(a,b)$ for any $a > 0$ and $b > 0$.[24] Furthermore, this effect of top-down transparency holds in a much broader class of informational environments.[25]

In other words, in the transparent case, the principal knows that the agent's own beliefs about the state of nature only "move" by $1/4$ based on the agent's private information. In the opaque case, the agent's beliefs about the state of nature move by a larger amount, $1/3$, based on his or her signal. This means that, to manipulate the principal's beliefs, the agent does not need to choose as costly a deviation from the policy he or she would choose if the principal did not observe the agent's decision.

The effect does rely to some degree on the assumption of "quadratic loss" preferences. We have not explored how much one might relax this assumption, but it seems clear that the agent and principal wanting to "target" the state of nature, θ, in similar ways is fundamental to the nesting of the incentive compatibility conditions for the transparent and opaque cases.

2 Incorporating Delegation: Endogenous Authority

Having considered the equilibrium behavior for every possible delegation situation (*i.e.*, all $\alpha_P > 0$ and $\alpha_A \geq 0$ with $\alpha_P + \alpha_A = 1$), we are now in a position to consider the principal's optimal delegation. That is, we now consider an extended model in which the principal, *prior* to observing his or her own signal, chooses how much authority to delegate to the agent. Formally, P chooses $\alpha_A \in [0,1]$,

[24] We focus on the special case of $a = b = 1$, which is the Uniform[0,1] distribution.

[25] Specifically, it can be shown that this property holds whenever the signals are distributed according to a likelihood function in a regular exponential family and the distribution of θ is the conjugate prior for the likelihood function; proof is available from the authors upon request.

retaining $\alpha_P = 1 - \alpha_A$ for himself or herself. The baseline game analyzed above then proceeds as before.

Precisely because agent A's policy decision is observed by the principal P *prior to P's policy decision*, granting positive unilateral authority to A (*i.e.*, $\alpha_A > 0$) can expand the possibility for truthful information-revelation in equilibrium. Connecting the agent's "message" with real-world policy consequences requires the agent to "put her money where her mouth is," and leads to more trustworthy messages from the agent to the principal.

Not so obvious, however, are the incentives of the principal in allocating such authority. On one hand, the principal will have a first-order incentive to retain decision-making authority in his or her own hands whenever the agent's preference bias differs from the principal's (*i.e.*, $\beta_A \neq \beta_P$). Thus, delegation of a positive degree of decision-making authority to an agent A will occur only to the degree that such delegation is required to support truthful revelation by that agent.

For very small degrees of divergence between an agent and the principal (*i.e.*, small values of $\beta_A - \beta_P$), no delegation will occur, as truthful revelation will prevail even without such delegation, and such delegation necessarily entails some agency loss for the principal.[26]

However, because the minimal discretionary authority required to elicit truthful revelation from a given agent A to the principal, P, is an increasing function of the difference between their preference biases (*i.e.*, $(\beta_A - \beta_P)$), the preference divergence between the principal and agent might be sufficiently extreme to outweigh (from the principal's standpoint) the increased expected payoff from truthful revelation by the agent. In such situations, the principal's optimal allocation of decision-making authority grants zero discretion to the agent. Accordingly, the degree of decision-making discretion delegated to an agent will possess a non-monotonic and counterintuitive relationship to the preference alignment between the principal and the agent.

Sequence of Play: Delegation Game. The *delegation game* proceeds as follows.

1. Player P assigns decision-making authority $\alpha_i \geq 0$ to each player $i \in \{P, A\}$, with $\alpha_P + \alpha_A = 1$.

[26]Obviously, this presumes that divergence is nonzero. Furthermore, the definition of "small enough" will depend on other characteristics of the situation, including most importantly how many *other* informative signals the principal will possess in equilibrium.

2. Nature determines the state of nature θ and players' signals, $s = \{s_P, s_A\}$.

3.
   - Opaque Case. Each player $i \in \{P, A\}$ privately observes his or her signal, $s_i \in \{0, 1\}$.
   - Top-Down Transparent Case. Each player $i \in \{P, A\}$ privately observes his or her signal, s_i, *and* the agent observes the principal's signal, s_P.

4. The agent, A, sets policy $y_A \in \mathbf{R}$.

5. Principal P observes y_A and sets policy $y_P \in \mathbf{R}$.

6. Game concludes, both players receive their payoffs.

As in the previous section, we analyze first the opaque case and then turn to the top-down transparency setting.

Endogenous Authority: The Opaque Case. Inequality (7) is satisfied if and only if

$$\alpha_A \geq \frac{8(\beta_A - \beta_P) - 1}{7/9 + 8(\beta_A - \beta_P)},$$

so that, for any $\beta = \{\beta_P, \beta_A\}$, the minimal feasible level of authority that the principal can grant to the agent and secure truthful policymaking in the opaque case is

$$\alpha_A^{*O}(\beta) = \max\left[\frac{8(\beta_A - \beta_P) - 1}{7/9 + 8(\beta_A - \beta_P)}, 0\right]. \tag{12}$$

Thus, $\alpha_A^{*O}(\beta)$ represents the minimal delegation of authority required to secure truthfulness by agent A when principal P's information is opaque to agent A. Of course, $\alpha_A^{*O}(\beta)$ is the "cheapest" means for principal P to secure truthfulness from agent A, but for sufficiently different preferences β_P and β_A, it might represent too costly a power sharing arrangement for principal P in terms of the induced bias of final policy, relative to principal P simply proceeding on the basis of his or her own signal, s_P. Proposition 6, presented in the appendix, formally describes optimal delegation in the opaque case. For the purposes of presentation, we first turn to the top-down transparency case and then explicitly compare optimal delegation in the two cases.

Endogenous Authority: The Top-Down Transparency Case. Inequality (11) is satisfied so long as

$$\alpha_A \geq 1 - \frac{1}{8(\beta_A - \beta_P)}.$$

From this it follows that, for any $\beta = \{\beta_P, \beta_A\}$, the minimal feasible level of authority that the principal can grant to the agent and secure truthful policymaking is

$$\alpha_A^{*T}(\beta) = \max\left[0, 1 - \frac{1}{8(\beta_A - \beta_P)}\right]. \tag{13}$$

Thus, $\alpha_A^{*T}(\beta)$ is the "least expensive" means by which the principal can incentivize agent A to apply/reveal his or her private information when the principal's information is transparent to the agent. As the preferences of the agents diverge, the minimum discretion required to elicit truthful revelation by agent A increases. Sufficiently divergent preferences implies that principal P would prefer to make policy based only on his or her own information and delegate zero authority to agent A. As described earlier, principal P will find it in his or her interest to share decision-making authority with agent A only for an intermediate range of preference divergence. Proposition 7, presented in the appendix, formally describes the optimal delegation strategy for the principal in this setting. Now, however, we present and compare the optimal delegations in the two cases side-by-side. This multifaceted comparison represents one of the two central conclusions of the article: top-down transparency increases the minimal level of authority that the principal must delegate in order to achieve truthful (*i.e.*, "informed") policymaking from the agent and, concomitantly, can lead the principal to retain exclusive policymaking authority and forego the agent's expertise/information entirely.

The Effect of Exogenous Transparency on Delegation. Regardless of the observability of the principal's information to the agent, there always exists a nonempty set of cases in which the principal will delegate positive authority to the agent in pursuit of truthful reporting of the agent's information.

Proposition 1 *For any pair of preference biases $\beta = \{\beta_P, \beta_A\}$, if the principal delegates positive* *authority in equilibrium to the agent in the top-down transparent case, then the principal delegates strictly more authority to the agent than he or she would in equilibrium in the opaque case:*

$$\alpha_A^{*T}(\beta) > 0 \Rightarrow \alpha_A^{*T}(\beta) > \alpha_A^{*O}(\beta).$$

Proof: Follows immediately from comparison of Propositions 6 and 7 (in the appendix). ∎

Proposition 1 implies that the situations (*i.e.*, sets of preference biases, $\beta = \{\beta_P, \beta_A\}$) in which the principal will find it optimal to delegate any nontrivial policymaking authority to the agent will differ between the opaque and top-down transparent cases. In fact, these sets of cases (in terms of the preference biases, $\beta = \{\beta_P, \beta_A\}$) are nested, as summarized in the next proposition.

Proposition 2 *For any pair of preference biases $\beta = \{\beta_P, \beta_A\}$, if the principal delegates positive authority in equilibrium to the agent in the top-down transparent case, then the principal also delegates positive authority in equilibrium to the agent in the opaque case:*

$$\alpha_A^{*T}(\beta) > 0 \Rightarrow \alpha_A^{*O}(\beta) > 0,$$

but the converse does not hold: there are pairs of preference biases $\beta = \{\beta_P, \beta_A\}$ such that

$$\alpha_A^{*O}(\beta) > 0 \text{ and } \alpha_A^{*T}(\beta) = 0.$$

Proof: Follows immediately from comparison of Propositions 6 and 7. ∎

Proposition 2 implies an odd, though partial, path through which information *begets* power: an agent with access to the information held by his or her principal faces a lower cost of manipulating that principal's subsequent policy decisions: the agent can obfuscate his or her own information through less extreme deviations from what the agent would do if he or she held complete authority. Furthermore, allowing the principal to conceal his or her own information from the agent will never eliminate, and indeed might engender, an incentive to delegate nontrivial, though partial, decision-making authority to the agent. Thus, as we return to below, *while top-down transparency will* *increase the precision and efficiency of the agent's policy decision, imposition of such transparency can somewhat ironically negate the principal's incentive to grant the agent any authority at all.*

The results regarding endogenous delegation in the two transparency regimes (Propositions 1 and 2) are portrayed in Figure 2. If delegation of positive discretionary authority is required to

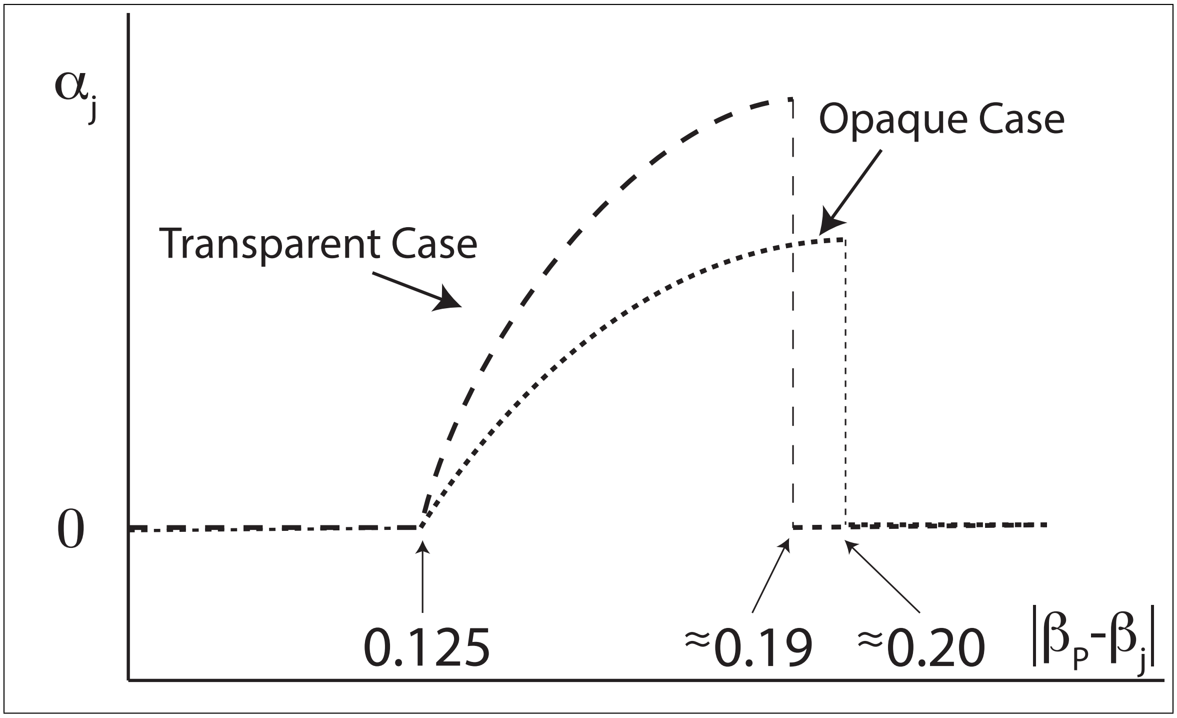

Figure 2: The Effect of Transparency and Preference Divergence on Delegation of Authority

elicit truthfulness from the message sender, *the minimal required level of discretionary authority is always greater when the principal's information is transparent and observable by the message sender.* This conclusion can be understood in a variety of ways, including the fact that *ex ante* incentive compatibility is always less restrictive than interim incentive compatibility. In other words, when the message sender is unaware of the principal's additional information about θ, he or she is less certain about whether he or she will gain from deviating from truthfulness.

The principal's incentive to delegate authority is greater when the principal's information is privately held. This conclusion has two important implications for any policy-interested party outside the model (*i.e.*, a third party, such as a voter or legislator) who might seek to impose top-down transparency on a hierarchical policy-making institution arrangement with endogenously delegated discretionary authority similar to that modeled here. Namely,

1. imposing such transparency will reduce the likelihood that the principal will choose to grant any such authority, but

2. if delegation of authority occurs in the transparent environment, the degree of authority that is delegated will be strictly larger than it would have been in the opaque case.

These implications present the foundations of a classic trade-off for third parties charged with delineating/circumscribing the institutional details of a hierarchical policymaking organization where the possibility of sub-delegation is inherent. Space constraints clearly prohibit the explication of a model that would fully incorporate an outsider/third party faced with such choices, so we leave this extension for future work.[27] We now extend the theory to encompass the possibility of endogenously determined top-down transparency.

3 Endogenous Top-Down Transparency

The previous subsection considered fixed game forms in which the agent A either could or could not observe s_P. We now consider the principal's preferences over this observability. To do this, we extend the game form so that the principal P can choose whether or not the agent A observes P's signal s_P, and then the game proceeds as in the delegation game analyzed above:

Sequence of Play: Delegation Game with Endogenous Transparency. The *delegation game with endogenous transparency* proceeds as follows.

1. Player P chooses whether to adopt top-down transparency ($t = T$) or opacity ($t = O$)

2. Player P assigns decision-making authority $\alpha_i \geq 0$ to each player $i \in \{P, A\}$, with $\alpha_P + \alpha_A = 1$.

3. Nature determines the state of nature θ and players' signals, $s = \{s_P, s_A\}$.

4. • OPAQUE CASE ($t = O$). Each player $i \in \{P, A\}$ privately observes his or her signal, s_i.

[27] As an aside, note that the tension presented by these two implications suggest that a third-party, B, whose preferences were more aligned with those of the agent than with those of the principal (*i.e.*, ($|\beta_B - \beta_A| < |\beta_B - \beta_P|$) might nevertheless actively work *against* attempts to require that the principal reveal his or her information to the subordinate).

- TOP-DOWN TRANSPARENT CASE ($t = T$). Each player $i \in \{P, A\}$ privately observes his or her signal, s_i, *and* the agent observes the principal's signal, s_P.

5. The agent, A, sets policy $y_A \in \mathbf{R}$.

6. Principal P observes y_A and sets policy $y_P \in \mathbf{R}$.

7. Game concludes, players receive payoffs.

The previous section's analysis illustrated that top-down transparency increases the amount of authority that the principal needs to share with agent A and simple reflection clearly indicates that such transparency is strictly preferred by the principal only if the principal delegates some authority to agent A. In fact, it turns out that the welfare impact of transparency (from the principal's point of view) is ambiguous: on one hand, the imposition of transparency weakly raises the "delegation price" of eliciting truthfulness. On the other hand, transparency reduces the "variance cost" of the agent's policy-making by giving him or her access to the same information as the principal. When $(\beta_A - \beta_P) \leq 1/8$, transparency is clearly in the principal's interest: transparency will effectively allow the the principal and agent to communicate truthfully and simultaneously. When $(\beta_A - \beta_P)$ is sufficiently large, however, opacity is clearly in the principal's interest. The most interesting case of this is in the region between $\frac{3+\sqrt{41}}{48}$ and ρ^* in Figure 2, where the principal voluntarily chooses to shut down (*i.e.*, not delegate so as to elicit) truthful communication when his or her information is transparent, but voluntarily grants significant discretionary authority to elicit truthful communication when his or her information is opaque. Thus, in this region, it is clear that the principal *and agent* strictly prefer opacity to transparency in these situations. The next proposition formally characterizes the principal's preference for transparency.[28]

Proposition 3 *When the principal, P, can delegate discretionary authority to agent A, the principal prefers that agent A observes the principal's information, s_P, prior to choosing y_A* (i.e.*, trans-*

[28] The proof of Proposition 3 is omitted, as it follows from simple calculations included in the statement of the proposition.

parency is preferred to opacity) whenever

$$(\beta_A - \beta_P) \leq {}^1\!/_7.$$

Thus, in line with the intuition described above, Proposition 3 confirms that the principal prefers transparency from an *ex ante* perspective for sufficiently small differences in biases, including some, but not all, situations in which the principal would prefer to delegate authority to the agent to elicit truthfulness from the agent.

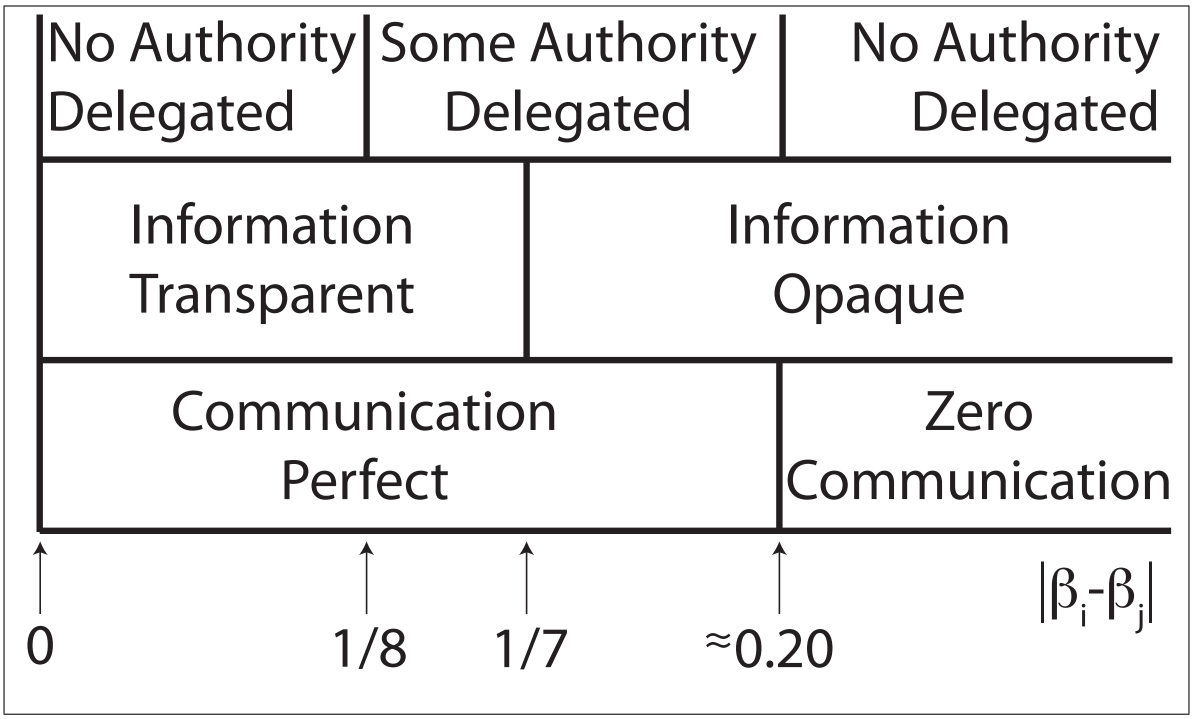

Figure 3: Summary of Results

Figure 3 summarizes our results. It illustrates the variety of potential organizational structures that characterize optimal designs by the principal. Communication can be perfect between the agent and principal, though this communication might require delegation of decision-making authority to the agent. In equilibrium, authority is delegated only if communication will be perfect as a result. Similarly, transparency will be chosen by the principal only if the ensuing communication is perfect. However, transparency might not be chosen even though the principal chooses to delegate authority to the agent and the ensuing communication is perfect. Finally, the absence of delegation of authority does not completely determine the information transmission from agent to principal. Rather, when authority is not delegated, the key issue for information transmission is the principal's choice of transparency. If the principal does not share authority with the agent, then the principal will choose to make his or her information transparent to the agent if and only if the ensuing communication will be perfect. That is, *transparency in the absence of power-sharing occurs only if the principal and agent can credibly communicate with each other through cheap talk.*

While our analysis considers a principal and a single agent, a natural extension involves multiple agents choosing in sequence. This has important implications for two key issues in the theory of delegation. The first concerns an emergent "chain of command," in the sense of increasing authority as one moves up the chain. When later agents act after observing the choices of earlier agents, they necessarily possess some of the principal's information when making their own decisions (as in the case of top down transparency). Accordingly, these intervening agents will be more tempted to manipulate the principal's decisions through their own decisions. This immediately implies that, when there are two agents with identical biases reporting in sequence to the principal, the first agent will have less authority than the second. The second issue is an organization's "span of control," or the maximum number of agents to which the principal wishes to delegate authority. It is generally true in this model that elicitating information from an additional agent becomes increasingly costly—requires greater delegated authority—as the number of agents increases.[29] Thus, there are situations in which the principal will grant positive discretionary authority to some, but not all, agents.

4 Conclusion

In this article, we have presented a model of decision making with dispersed information and decentralized authority. A key element of the model are that decisions are observable and made in sequence; thus, early decisions can reveal information to later actors. An early actor's degree of policy authority directly affects its incentives to reveal information through its policy choices, by

[29]The principal source of this tension is analogous to the "congestion effect" discussed in Galeotti, Ghiglino and Squintani (2013): manipulation/insincere policymaking becomes more tempting for any given agent as the principal becomes more informed in the sense of possessing more signals and/or observing more truthful messages/policy decisions.

giving it some "skin in the game" of policy making. In light of this, we consider the incentives of the principal (the last decision maker in the sequence) to delegate partial authority to the agent (the first decision maker). Optimal delegation balances its effects on the agent's incentives to reveal information through its policy, and the principal's desire to retain some residual authority to make use of the information so revealed.

Our results speak to a variety of considerations in institutional design involving delegated authority and transparency. The theory illustrates how both "top-down" transparency and wider spans of control can hinder "bottom-up" information aggregation. These problems can be mitigated through delegation of discretionary authority to informed agents. When such delegation occurs, the principal's optimal power-sharing arrangement is a non-monotonic function of the preference divergence between the delegate agent and the principal. This can lead to observed power structures in which the principal's closest (preference-based) "allies" have significantly less (or even zero) power, while other agents with moderately different preferences from the principal are granted significant individual decision-making power.

References

Aghion, Philippe and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy* 105(1):1–29.

Alonso, Ricardo and Niko Matouschek. 2008. "Optimal Delegation." *Review of Economic Studies* 75:259–293.

Austen-Smith, David and Jeffrey S. Banks. 1996. "Information Aggregation, Rationality and the Condorcet Jury Theorem." *American Political Science Review* 90:3445.

Bendor, Jonathan and Adam Meirowitz. 2004. "Spatial Models of Delegation." *American Political Science Review* 98(2):293–310.

Callander, Steven. 2008. "A Theory of Policy Expertise." *Quarterly Journal of Political Science* 3(2):123–140.

Crawford, Vincent P. and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50(6):1431–1451.

Dessein, Wouter. 2002. "Authority and Communication in Organizations." *Review of Economic Studies* 69(4):811–838.

Dewan, Torun and Francesco Squintani. 2012. "The Role of Party Factions: An Information Aggregation Approach." Working Paper, University of Warwick.

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

Feddersen, Timothy J. and Wolfgang Pesendorfer. 1998. "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts." *American Political Science Review* 92:23–36.

Feddersen, Timothy and Wolfgang Pesendorfer. 1996. "The Swing Voters Curse." *American Economic Review* 86(3):408–424.

Gailmard, Sean. 2009. "Oversight and Agency Problems in Legislative-Bureaucratic Interaction." *Journal of Theoretical Politics* 121(2):161–186.

Gailmard, Sean and John W. Patty. 2012*a*. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15:353–377.

Gailmard, Sean and John W. Patty. 2012*b*. *Learning While Governing: Information, Accountability, and Executive Branch Institutions*. Chicago, IL: University of Chicago Press.

Galeotti, Andrea, Christian Ghiglino and Francesco Squintani. 2013. "Strategic Information Transmission in Networks." *Journal of Economic Theory* 148(5):1751–1769.

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* 3:287–335.

Hammond, Thomas H. and Gary J. Miller. 1985. "A Social Choice Perspective on Expertise and Authority in Bureaucracy." *American Journal of Political Science* 29(1):1–28.

Hammond, Thomas H. and Paul A. Thomas. 1989. "The Impossibility of a Neutral Hierarchy." *Journal of Law, Economics, and Organization* 5(1):155–184.

Holmström, Bengt. 1984. On the Theory of Delegation. In *Bayesian Models in Economic Theory*, ed. Marcel Boyer and Richard Khilstrom. New York: North Holland.

Miller, Gary. 1992. *Managerial Dilemmas: The Political Economy of Hierarchy*. New York, NY: Cambridge University Press.

Patty, John W. and Elizabeth Maggie Penn. 2014. "Sequential Decision-Making & Information Aggregation in Small Networks." *Political Science Research & Methods* 2(2):249–271.

Penn, Elizabeth Maggie. 2015. "Engagement, Disengagement, or Exit: A Theory of Equilibrium Associations." *American Journal of Political Science* Forthcoming.

Ting, Michael M. 2002. "A Theory of Jurisdictional Assignments in Bureaucracies." *American Journal of Political Science* 46(2):364–378.

Ting, Michael M. 2003. "A Strategic Theory of Bureaucratic Redundancy." *American Journal of Political Science* 47(2):274–292.

Ting, Michael M. 2008. "Whistleblowing." *American Political Science Review* 102:249–267.

Ting, Michael M. 2011. "Organizational Capacity." *Journal of Law, Economics, and Organization* 27(2):245–271.