Competence and Advice*

Short Title: Competence and Advice

Anna Denisenko† Catherine Hafer‡ Dimitri Landa§

May 6, 2024

Keywords: principal-agent, advice, expertise, information, bureaucracy, persuasion

*We thank Ethan Bueno de Mesquita, Wiola Dziuda, Sandy Gordon, Federica Izzo, Carlo Prato, Matthew Stephenson, participants at the PECO 2022, Kellogg Political Economy Conference, and the NYU PE Workshop, as well as the *AJPS* editors and anonymous referees for valuable feedback on earlier versions of this paper.

†Harris School of Public Policy, University of Chicago, e-mail: ad4205@nyu.edu

‡Associate Professor, Wilf Family Department of Politics, NYU, e-mail: catherine.hafer@nyu.edu

§Professor, Wilf Family Department of Politics, NYU, e-mail: dimitri.landa@nyu.edu

Abstract

We develop a theory of policy advice that focuses on the relationship between the competence of the advisor (e.g., an expert bureaucracy) and the quality of advice that the leader may expect. We describe important tensions between these features present in a wide class of substantively important circumstances. These tensions point to the presence of a trade-off between receiving advice more often and receiving more informative advice. The optimal realization of this trade-off for the leader sometimes induces her to prefer advisors of limited competence – a preference that, we show, is robust under different informational assumptions. We consider how institutional tools available to leaders affect preferences for advisor competence and the quality of advice they may expect to receive in equilibrium.

Word count: 9693

Government leaders are inevitably defined by the outcomes of policies they choose, yet those policies are rarely a product of the leaders' *sui generis* judgments. Instead, leaders must often rely on advice from government officials better positioned to develop expertise and competence with respect to the questions of interest. It is tempting to assert, in light of this, that the leaders' political and policy successes owe a lot to the competence of their advisors. But this seemingly obvious claim obscures fundamental strategic complexities at the core of the leader-advisor relationship.

We develop a theory of policy advice that focuses on the connection between central elements of that relationship – advisor competence (agency expertise) and the quality and quantity of advice that the leader may expect – and calls into question some strongly held intuitions about the quality of advice and about the politics of policy-making more broadly.

Our theory highlights an underappreciated tension between advisor (agency) competence and the informativeness of advice in equilibrium: while more competent advisors sometimes give more informative advice, in a wide class of substantively important circumstances we characterize, we should expect the opposite to be the case. More broadly, our theory sheds light on when the trade-off between between better information and more information revelation is likely to be acute and on what institutional features can mitigate it.

The policy-making setting to which our theory most readily applies is the interaction between political leaders and bureaucrats. In such settings, the leader might be an elected official or a politically appointed head of an established government agency, and the informed advisors are the agency's career civil servants. The following examples from such settings, intentionally somewhat stylized, evoke recent events and help motivate our analysis.

1. The political appointees of a U.S. President are seeking specific information about corrupt government practices in foreign country X. The U.S. Ambassador to X, who is a recognized expert on business and political practices in X, is in a particularly good position to have such information. The ambassador sees the long-term goal of his work to be supporting democratic institutions and political development in X, which

he believes to be central to the U.S. interests in the region and to be best achieved by strengthening the rule of law in X. He also suspects that the political leadership he advises places greater weight on the President's own immediate electoral fortunes than on the success of democratic practices in X, and that the leadership's sudden intense interest in corruption in X is borne of a hope of finding a pressure point against a political rival. This advisor, then, chooses to risk being reassigned rather than reveal to the leader what he knows, anticipating that if he is replaced with a more compliant appointee, the latter will also likely to be less knowledgeable and so less harmful to X.

2. A regulatory agency has the mission of protecting consumers from predatory business practices and a staff consisting mainly of civil servants attracted to that mission. The agency administrator is a political appointee, whose preferences reflect the political interests of the sitting president, whose agenda is contrary to the agency mission and who wishes to minimize government regulation of business and to promote their profitability. The administrator (the leader) seeks to identify the regulatory rule that will be most advantageous to a particular subset of businesses, while the agency staff (the advisor) believes that that should not be a determinant of the policy choice, given the agency's mission. Members of the staff who are convinced that the agency would be regarded as relatively incompetent by the adverse administration are willing to be forthcoming with the information relevant to the administration's aims and avoid the retaliation against the agency, while those members who are expecting that the staff would be perceived as competent on the policy question are more willing to withhold information and risk the retaliation.

3. The head of country Y's government strongly supports restarting the program of targeted assassinations of known leaders of terrorist organizations, which he values, in part, because they give him a chance to look tough and effective in protecting the country. Suppose that Y's head of an intelligence agency, on the other hand, has come

to believe that targeted killings of terrorist leaders provide relatively little benefit in disrupting terrorist activity and distract the agency from its mission of developing intelligence about potential terrorist threats. While the leader presses the head of the intelligence agency (the advisor) to identify the location of a desired target, the advisor may prefer to forgo gathering the relevant information, even at the risk of angering the leader.

In these and many other instances, civil servants have preferences that suggest an interest in maintaining the status quo. Lawrence J. Peter (of the Peter Principle) puts this in a characteristically flip fashion: "Bureaucracy defends the status quo long past the time when the quo has lost its status." But while Peter's comment implies a certain degree of criticism, Huq and Ginsburg (2018) view this conservatism as a guarantor of political stability, and describe as essentially symmetric with respect to possible challenges to status quo: "Just as bureaucracy may make progressive reform difficult to achieve, it also slows down rapid shifts away from liberal democratic norms" (p. 129). The bureaucracies' status-quo bias need not imply their primitive conservatism: as another astute observer of bureaucracies observes, bureaucrats "think of their temporary political superiors as birds of passage (and of themselves as the guardians of stability and continuity)... The incentive structure engenders bureaucratic resistance to abrupt, hasty changes in direction (and, indeed, was designed in part to encourage such steadiness)" (Kaufman (1981), p. 4).

In our model, an advisor who, like the bureaucrats in the above description, has a bias toward the status-quo and has access to superior information, chooses whether to reveal to the leader her privately observed signal about the state of the world. The leader prefers to set the policy by the value of the state of the world. The effective gap in the policy preferences between the advisor and the leader is, thus, contingent on the state: when the state is closer to the status quo, the gap is smaller, but as the state gets farther away from the status quo, the difference in preferences grows. Assuming the advisor shares her signal, the leader updates more strongly if the advisor is commonly known to be more competent and, conditional on such an update, shifts policy farther. Given this expectation and the gap between the most preferred policies of the leader and of her advisor, a more competent advisor stands to lose more from revealing her signal than does a less competent one. In general circumstances, the less competent advisors will, thus, have stronger incentives to reveal their information to the leader than will more competent ones. The leader, then, faces a trade-off between the quality of advice he receives and the likelihood of receiving it.

We describe conditions that influence the advisors' incentives to share and obtain information, as well as conditions that determine when leaders may be better off having advisors with lower competence, even though it means getting advice that is less reliable. We also show that the mandating disclosure of advisors' information (a common consideration in the context of relationships between political leaders and bureaucratic agencies) has an equivocal effect on the leader's utility. When the advisor's information is fixed, it allows the leader to extract more information and is more beneficial when the advising agency is more competent. But the opposite is true when the advisor's information is endogenous: mandating disclosure encourages more competent advisors not to acquire information, and the leader should limit requiring disclosure to when the advisor is relatively incompetent – that is, when the value of information that the enforced disclosure could reveal is, ultimately, not too great.

The rest of the paper proceeds as follows. Following a brief review of the prior literature, we introduce our model. We then begin its analysis by studying conditions that generate, in equilibrium, the trade-off between advisor competence (i.e., the quality of the information available to the advisor) and the strategically supportable quantity of advice offered. After that, we extend the model to include endogenous information acquisition, modeled as an experiment with non-contractible precision, and let the Advisor choose both whether to conduct the experiment and the experiment's precision. We show, first, that the only precision consistent with robust equilibrium play is the maximum possible precision, and that the Advisor's competence affects her incentives to obtain information in a way that echoes her incentives to reveal information in the fixed-information environment, but that such incentives are moderated by the publicity of the experiment. In the last substantive section, we introduce the possibility of mandating disclosure of the Advisor's information and study its effects on the Advisor's incentives to acquire information and the Leader's equilibrium utility.

Connection to the Literature

Information revelation through advice has been the focus of a substantial literature in political economy, including in the political economy of bureaucracy. For a review focusing on the themes of delegation and communication within hierarchies, see Gailmard and Patty (2012). One branch of this literature models communication as "cheap talk" in which an advisor's set of feasible messages is independent of her information. Early influential studies in this branch include Gilligan and Krehbiel (1989), Austen-Smith (1990), and Austen-Smith (1993). (For an important review of this literature, see Sobel (2013).) A key result in this literature is that the divergence in the actors' preferences limits revelation, and successful communication requires that the advisor's and the leader's preferences be sufficiently aligned. (Important exceptions are Battaglini (2002) and Aybas and Callander (2023), who identify policy-relevant settings in which communication fully favors, respectively, the receiver, and the sender.)

In contrast with this literature, this paper models the communication of hard evidence and focuses on the effects of the advisor's competence on her incentives to reveal existing, and to obtain additional, information. We model the advisor's communication as verifiable messaging for both pragmatic and substantive reasons. To understand the pragmatic reason, first note that if the advisors can distort the signal in messaging the leader, then, in the setting we study, advisors of all types will have an incentive to do so maximally – that is, to send messages that are maximally uninformative. Given our interest in understanding the implications of the variation in advisor competence for the quality of their advice to the leader, unverifiable messages are, thus, not the right informational setting for the analysis. In contrast, in the context of verifiable messages, the variation in advisor type induces important differences in the nature of communication, suggesting that that is an appropriate setting for our analysis. The second set of reasons is substantive. First, the evidence available to the bureaucratic agencies is typically in the form of internal reports, distorting which would create substantial administrative and legal jeopardy. And second, the leader's policy action based on the received advice must often engage directly with the evidentiary basis of the message (e.g., verifying the presence and location of an assassination target before and after the hit, or showing the evidence of corruption in an attempt to coerce a foreign leader to take a particular action). Given such constraints, a more natural way to think of the possibility of distortion is with respect to the (im-)precision of the information that the agency chooses to acquire, and we study such a possibility in an extension of our baseline model below.

Another branch of the literature on strategic communication focuses precisely on this kind of hard information. A touchstone finding in this literature is that all private information is revealed in equilibrium (Milgrom, 1981, 2008), and a major focus of subsequent work has been on understanding when the unraveling logic of the full revelation result does not hold. An important early paper is Shin (1994), which shows that the leader will not be able to infer perfectly the advisor's private information in the event of "no news" when the advisor's knowledge is imperfect; see also Wolinsky (2003). Dziuda (2011) shows that, in a setting where the fixed expert's preferences are different and unknown to the decision-maker, there is never full disclosure, but the expert offers pros and cons for the advocated alternative in order to pool with the honest/non-strategic type. In the present model, there is, similarly, no unraveling. A key reason is that, given the sender's optimal choice of strategy, the receiver's optimal choice of action is the same for every null signal that he might receive in equilibrium, for every possible sender competence.[1]

A seminal paper on public information acquisition by an interested agent is Kamenica and

[1] One could rank null signals by quality in the sense of Milgrom (1981).

Gentzkow (2011). Unlike the setting in that paper and in the subsequent extensive literature, the public experiment setting we study is one in which the precision of the experiment is not (directly) observable. Di Pei (2015) and Gentzkow and Kamenica (2017) study cheap-talk and verifiable-information models, respectively, in which the agent acquires information covertly and then decides whether to reveal it (in the latter paper, with the relevant received evidence). A key result in these papers is that requiring disclosure does not alter equilibrium play. As indicated above, this result does not hold in our model. Because explicating the reasons for this contrast builds on the detailed analysis below, we defer the discussion of this contrast until later in the paper, in the section on disclosure requirement.

A key downside of transparency explored in the literature is that transparency encourages the agents to pander to their principal's prior beliefs (in political settings, see, e.g., Canes-Wrone, Herron and Shotts (2001); Fox (2007); Patty and Turner (2021); and Turner (2022)). The mechanism that underlies the negative consequences of the disclosure mandate in our model is different – it comes from the effect of discouraging information acquisition by the agent due to information leakage and is related to the "stovepiping" studied in Gailmard and Patty (2013).[2]

Our focus on the strategic implications of the relationship between advisors' competence and the quality of advice they could give to the leaders distinguishes our model from previous studies that uncover a possible downside for the principal of the agent's higher competence. Egorov and Sonin (2011, 2023) explore the competence-loyalty trade-off in the relationship between leaders and advisors. In their model, the higher is the advisor's competence, the more confident the advisor is that his betrayal of the leader will lead to the enemy's victory, and, thus, the costlier it is for the leader to enforce the loyalty of the more qualified advisors.

[2]Somewhat more distantly, our model is also related to the literature on the effects and value of biased information. In the present study, we abstract away from the possibility of ideological bias on the part of either principals or agents (they are explored in detail in a companion paper). In a seminal paper in that literature, Calvert (1985) shows that the policy-maker may prefer a biased adviser when the policy-maker is concerned about extreme utility realizations from policy choices. Patty (2009), who examines the incentives stemming from the supply side of the advice, finds that a policy-maker's bias leads to more biased information collection by the advisor and bias-reinforcing advice.

While the advice of a less competent advisor is, thus, less reliable, such advisors are more loyal. (For related models, see Terai and Glazer (2018); Lagerlöf (2012); Zakharov (2016).) Buisseret and Prato (2016) show that, in a multi-district model of political accountability, expectation of strategic responses by voters in other districts can make increasing competence of a representative unattractive to her own voters.

Closer to our focus on informational effects associated with agent competence, Prendergast and Topel (1996) show that, in a three-level hierarchy, the principal may prefer low-powered incentives (which can lead to the selection of lower-competence subordinates) as a way of limiting distortions resulting from mid-level supervisors' favoritism, but the principal's utility is, in equilibrium, always increasing in the latter's competence. Bhattacharya and Mukherjee (2013) find that improving an advisor's quality in a panel of experts can lower the decision-maker's utility, but, in their model, revealed signals are quality-invariant, and the the negative effect of greater advisor competence is driven by the strategic effects on revelation by other advisors, and so possible only in the panel of multiple advisors. Sobel (1993) shows that when the agent is uninformed, she might exert more effort to achieve an outcome than would a better informed agent, leading to the possibility of the former being more attractive to the principals. Gailmard and Patty (2007) model bureaucratic competence as an exogenous cost of acquiring expertise and show that when the cost is too high, the legislature does not reward the advisor for expertise acquisition. The novel trade-off that we uncover between the quality and the quantity of information in the advisor's strategic communication choices to the leader is not present in these environments, and, further, we show this trade-off to be present in the *binary* relationship between the advisor and the leader, without turning on strategic action by third parties.

Finally, in the present model, the divergence of the policy preferences between the advisor and the leader is state-contingent both in its magnitude and direction. This distinguishes this model from the Crawford and Sobel (1982) modeling framework, in which the absolute value of the "bias" of the sender (with respect to the receiver) is constant in the value of the state. Other models that assume state-contingency of “bias” include Ashworth and Sasso (2019) and Che and Kartik (2009).

The Baseline Environment

We analyze a strategic interaction between a Leader (he) and an Advisor (she) of known competence θ. The Advisor's competence determines the informativeness of the signal about the state of the world that she privately observes. By way of interpretation, consider competence θ as a characteristic not of a single individual, but, instead, of a bureaucratic agency, whose information is constrained by formal and informal institutional rules.

The Advisor may share her information about the state with the Leader, who does not directly observe the state. The Advisor's messages are verifiable; for the reasons discussed above, this is the natural environment in which to study the implications of advisor competence for the quality of advice to the leader.The timeline of the game is as follows:

1. Nature determines the state of the world $w \in \mathbf{R}$, where w is a draw from a standard normal distribution $N(0, 1)$.

2. The Advisor of commonly known competence θ observes signal s about the state of the world w, $s = w + \varepsilon$. The variable ε represents random noise drawn from a normal distribution with mean 0 and precision θ (s.t. $\theta \in \mathbf{R}^+$), $\varepsilon \sim N(0, 1/\theta)$.

3. The Advisor chooses whether or not to disclose her signal to the Leader in message $m \in \{s, \varnothing\}$.

4. The Leader observes message m and decides which policy $a \in \mathbf{R}$ to implement.

We denote the Leader's and the Advisor's preferences by $u_L(\cdot)$ and $u_A(\cdot)$ correspondingly. We assume that the Leader wants to match the state of the world (w) and suffers a disutility from her policy choice's departures from w:

$$u_L(a, w) = -(a - w)^2. \tag{1}$$

Note that we assume that the Leader does not derive utility from the Advisor's competence directly; the Advisor's competence affects the Leader's utility only indirectly, as it alters the informativeness of advice (see below).

The Advisor exhibits more conservative preferences than the Leader and wishes the policy to match $c \cdot 0 + (1 - c) \cdot w$, where $c \in [0, 1]$ measures the Advisor's conservatism, with higher c corresponding to a more conservative preference. In other words, the Advisor desires policy that is closer to the distribution mean than that favored by the Leader. The Advisor suffers a disutility from the departures of the Leader's policy a from her own, more conservative, preference. Apart from depending on the Leader's policy choice, the Advisor's utility also depends on rewards she receives when the signal she sends is informative. We assume, in particular, that Advisor receives a finite and positive office-based benefit Ψ if and only if she sends an informative signal (i.e., if $m \neq \varnothing$). We interpret this benefit as a bonus, a value of future influence or latitude, or a promotion within the Leader's administration. For convenience, we will refer to Ψ below as the "reward."[3] Formally, the Advisor's utility is given by

$$u_A(m, w, a) = \begin{cases} -(a - (1 - c) \cdot w)^2 & if\ m = \varnothing, \\ -(a - (1 - c) \cdot w)^2 + \Psi & \text{else.} \end{cases} \tag{2}$$

The state-contingent absolute value of the policy bias separating the Advisor and the Leader is consistent with the account, in the Introduction, of the differences between the in-

[3] In order to focus on the information revelation aspect, the baseline model abstracts from the optimal reward-granting mechanism design. In a reduced way, we assume that the Advisor's reward Ψ is fully contingent on revelation. It is straightforward to extend the baseline model by introducing explicitly the Leader's decision on whether to award Ψ to the Advisor in the event of a lack of revelation. We show in Online Appendix (page 6) that, consistent with the reduced-form assumption, the plausible Leader payoff-dominant equilibrium in this extended game is characterized by the Leader always rewarding the Advisor if and only if she reveals her signal.

centives of bureaucrats and elected officials in many existing democracies. Broadly speaking, while the time horizons of elected officials tend to be short – they tend to be strongly incentivized to pander to voters' opinions or otherwise bolster their immediate electoral fortunes, and discount the longer-term implications of their actions – bureaucrats' relatively long time horizons make them apt to discount the value of responding to short-term trends, as they often, and not unreasonably, expect to bear the costs of policy implementation and policy change. The assumption, in effect, is motivated by the idea that such costs are increasing in the distance between the policy chosen by the leader and average policy (set at status quo).[4]

The solution concept is weak Perfect Bayesian Equilibrium, which imposes the following requirements. Following the Advisor's message m, the Leader forms posterior belief $\mu(w|m,\theta)$ derived from the Advisor's strategy using Bayes's rule, where $\mu(\cdot)$ denotes the probability that the state of the world is w conditional on the message m the Leader observes and the competence of the Advisor θ. The Leader chooses $a^*(m,\theta)$ such that it maximizes his expected utility, given his posterior beliefs

$$a^*(m,\theta) = \arg\max_a \int u_L(a,w)d\mu(w|m,\theta). \tag{3}$$

Further, the Advisor's message $m^*(\theta)$ maximizes the Advisor's expected utility, given her (correct) expectation of the Leader's policy choice

$$m^*(\theta) = \arg\max_{m\in\{s,\varnothing\}} u_A(a^*(m,\theta)). \tag{4}$$

Analysis

To characterize the equilibrium behavior in our model, we proceed by backward induction.

[4] The Advisor's utility can be microfounded as identical to the Leader's save for the added term of the quadratic costs of policy implementation. Under that representation, it leads to the equilibrium revelation interval that is isomorphic to the equilibrium revelation interval in the present specification. See Online Appendix (page 7) for details.

When the Leader observes the (directly) informative message $m = s$, he adopts a policy equal to the mean of the posterior distribution conditional on the message he observes. When the Advisor does not convey her signal, i.e., $m = \varnothing$, the Leader chooses policy equal to the mean of his posterior belief conditional on the Advisor's having concealed her signal. Note that both $m = s$ and $m = \emptyset$ may be informative in equilibrium; but the former is so intrinsically, whereas the informational content of the latter is determined by its use in equilibrium. Henceforth, we will use the term "informative message" solely to indicate the intrinsically informative message $m = s$. Given the Leader's anticipated policy response, the Advisor sends an informative message when her utility from revelation exceeds her utility from concealment. The following proposition characterizes the Leader's and the Advisor's equilibrium strategies:

Proposition 1. *In the unique equilibrium, the Advisor sends an informative signal to the Leader if and only if*

1. *she is not too conservative,* $c \leq 1/2$*; or*

2. *she is sufficiently conservative,* $c \geq 1/2$ *and the signal is close enough to the mean state,* $-\hat{s}(\theta, \cdot) < s < \hat{s}(\theta, \cdot)$*, where* $\hat{s}(\theta, \cdot) \equiv \sqrt{\Psi} \cdot (1 + \frac{1}{\theta}) \cdot \frac{1}{\sqrt{2 \cdot c - 1}}$.

Otherwise, the Advisor conceals her information.

The Leader implements policy

$$a^* = \begin{cases} \frac{m}{1+1/\theta}, & \text{if } m = s; \\ 0, & \text{if } m = \varnothing. \end{cases} \tag{5}$$

Proof. See Appendix A. □

The final policy the Leader adopts following information revelation depends on the information shared by the advisor and the advisor's competence, θ. The higher the Advisor's competence θ, the more the Leader relies on the Advisor's message when selecting policy.

Thus, holding fixed s, the greater the Advisor's competence, the larger the difference between the policy the Leader selects after receiving the message and Advisor's ideal point. This, in turn, more strongly discourages the Advisor of higher θ from disclosing s in the first place. This observation highlights the main trade-off that determines the Leader's utility: more competent Advisors are less likely to reveal information, forcing the Leader to balance the quality of advice and its availability.[5]

Furthermore, it's important to note that the conditions for revelation in equilibrium are such that when the advisor chooses $m = \varnothing$, the Leader's posterior belief is symmetric about the prior mean, albeit more dispersed than the prior. This observation justifies the Leader's policy choice when there's a lack of revelation ($a^*(m = \varnothing) = 0$).[6]

Figure 1 shows the thresholds $\pm\hat{s}(\theta, \cdot)$ as a function of the Advisor's competence θ assuming $c = 1$. The shaded area depicts the signals that an Advisor of competence θ reveals to the Leader. As we can see in Figure 1, the higher the Advisor's competence, the smaller the range of informative messages that the Advisor sends to the Leader.

[5] The incentives underlying this result can be instructively contrasted with incentives in an alternative model with additive Advisor bias in the sense of Crawford and Sobel (1982). In such a model, the Advisor would reveal her information to the Leader if and only if the signal realization falls significantly far from the status quo, and the revelation would increase with the Advisor's competence, as opposed to decreasing, as our model predicts. Formal details can be found in the Online Appendix (page 9).

[6] The symmetry of the revelation interval around the status quo is a consequence of our assumption that the prior mean is at the status quo. One might imagine a setting in which that is not so. Assuming that the Advisor is biased toward the prior (with the background interpretation that shifts in the state are "short-term," and the Advisor prefers a "long-term" average), the revelation interval would shift away from the status quo (and the status quo would become irrelevant), but would remain symmetric around the priorand the no-unraveling result would remain.

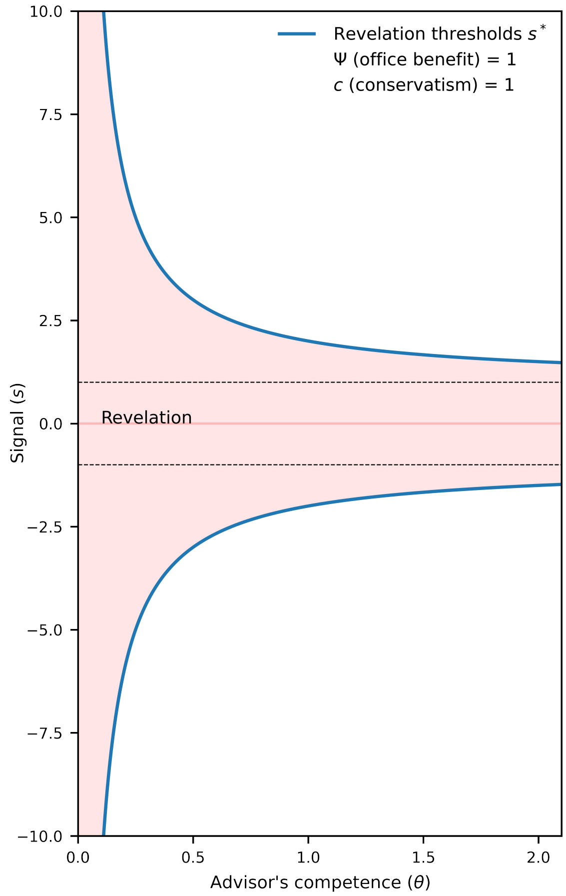

Figure 1: Revelation thresholds as a function of the Advisor's competence. Notes: The solid curves show the signal thresholds $\pm\hat{s}(\theta)$ as a function of the Advisor's competence (θ) when $c = 1$. When the Advisor receives a signal $s \in [-\hat{s}(\theta), \hat{s}(\theta)]$, she reveals her signal to the Leader. Dashed lines are asymptotes of the thresholds $\pm\hat{s}(\theta)$ when $\Psi = 1$.

The next proposition characterizes how the Advisor's incentives to share information with the Leader, given the Leader's anticipated policy response, vary with the parameters of the model.

Proposition 2. *In the equilibrium, the Advisor's incentive to send an informative signal to the Leader*

1. *weakly decreases with her competence;*

2. *weakly increases in the Advisor's valuation of the reward, Ψ; and*

3. weakly decreases in the Advisor's conservatism, c.

Proof. Follows from Proposition 1 by taking derivatives with respect to θ, Ψ, and c. $\square$

The reward Ψ directly enters the Advisor's utility, encouraging her to reveal more information to the Leader. The Advisor's conservatism c captures the Advisor's preference for the prior mean vs. the current state. The higher is the Advisor's conservatism, the farther is her ideal point from the policy the Leader chooses after $m = s$, on average, thereby discouraging revelation. Figure 2 illustrates these relationships at a fixed value of c (panel (a)) and fixed value of Ψ (panel (b)).

Figure 2: Advisor's equilibrium strategy based on competence (θ), reward (Ψ), and conservatism (c)

(a) Advisor's revelation thresholds as a function of competence at $c = 1$

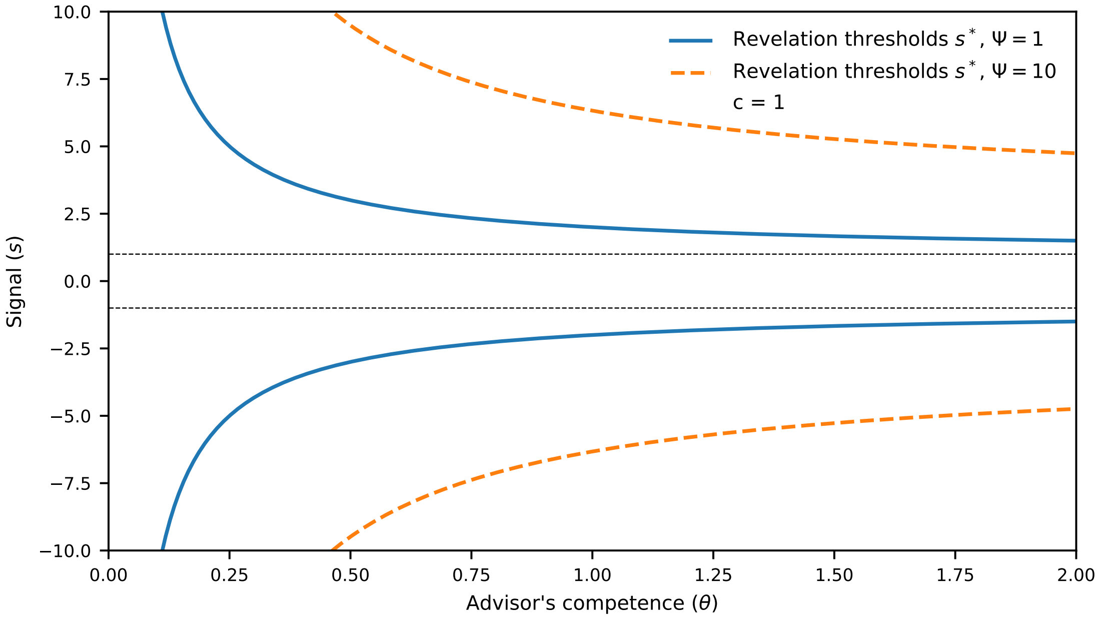

(b) Advisor's revelation thresholds as a function of competence at $\Psi = 1$

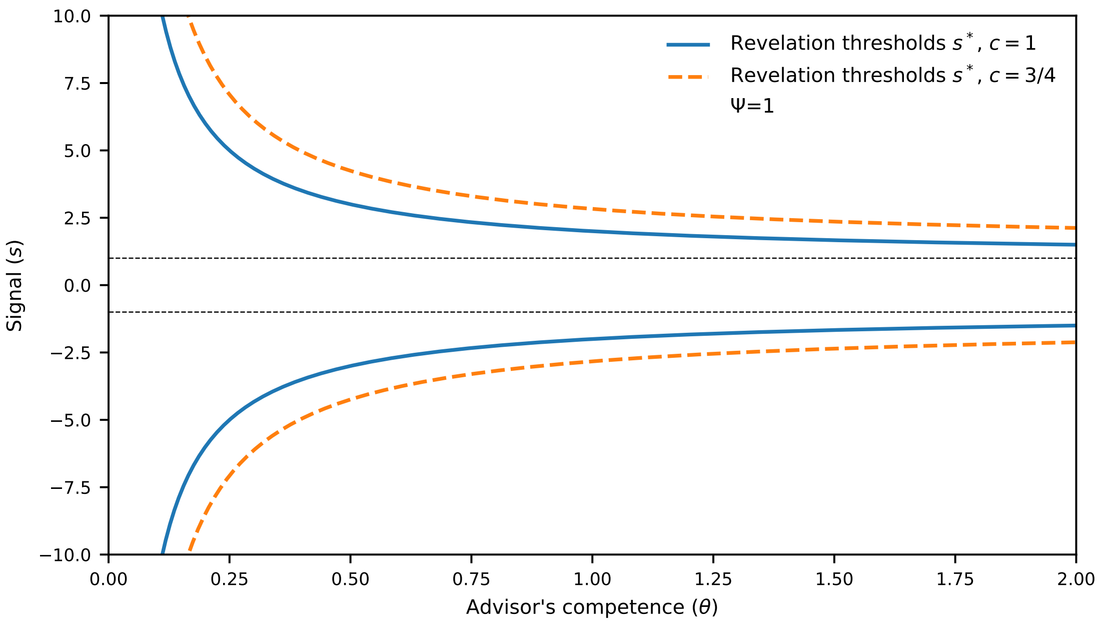

Note: The dashed curves in panel (a) show the Advisor's revelation thresholds for $\Psi = 10$. The solid curves in panel (a) show the Advisor's revelation thresholds for $\Psi = 1$. The dashed curves in panel (b) show the Advisor's revelation thresholds for $c = 3/4$. The solid curves in panel (b) show the Advisor's revelation thresholds for $c = 1$.

Denote the equilibrium probability that an Advisor of competence θ reveals her informative signal to the leader as $r(\theta) \equiv Pr[s \in \big(-\hat{s}(\theta), \hat{s}(\theta)\big)]$. The Leader's equilibrium expected utility with an Advisor of competence θ is thus

$$E[u_L^*(\theta)] = U_L(\theta) = r(\theta) \cdot \frac{-1}{1+\theta} + (1 - r(\theta)) \cdot E[-w^2 | s \notin \big(-\hat{s}(\theta), \hat{s}(\theta)\big)]. \tag{6}$$

Recall that, conditional on not observing an informative message from the Advisor, the Leader infers that the signal the Advisor observed is not in $\big(-\hat{s}(\theta), \hat{s}(\theta)\big)$. For the closed form of the Leader's expected utility in equilibrium, see Appendix B.

The following proposition shows a key property of the Leader's expected utility as a function of Advisor competence.[7]

Proposition 3. *1. For $c \leq 1/2$, the Leader's equilibrium welfare is always weakly increasing in the Advisor's competence.*

2. For $c > 1/2$ and strictly positive Ψ, there exists a unique threshold $\Psi^*(c)$ such that

(a) for all Ψ such that $\Psi < \Psi^*(c)$, the Leader's equilibrium welfare is maximized with an Advisor of limited competence;

(b) for all $\Psi \geq \Psi^*(c)$, the Leader's equilibrium welfare can always be improved with an Advisor of sufficiently higher competence.

3. For $c > 1/2$ and $\Psi = 0$, the Leader's equilibrium welfare does not depend on the Advisor's competence.

Proof. For proof see Appendix C. □

When the Advisor's policy preference closely aligns the Leader's ($c \leq 1/2$), she reveals information regardless of the value of the reward benefit Ψ. Consequently, the Leader does

[7] In the Online Appendix (page 5) we extend the findings presented below and demonstrate that the Leader's preference for a less competent Advisor, driven by the trade-off between quality and quantity of information, is not an artifact of the specific functional forms of the utilities and signal distributions. Further, our central results can be obtained with a discrete state space as well. However, the continuous model we study facilitates smooth comparative statics analysis, in contrast to what a discrete model may offer.

not face a trade-off between the quality and quantity of information and derives the greatest utility from an Advisor of maximal possible competence.

When, instead, the Advisor's ideal point is closer to the status quo than to the Leader's ideal point (when $c \geq 1/2$), absent additional revelation incentives – when revealing information does not come with a reward, $\Psi = 0$, – the Advisor is always better off withholding her information from the Leader, thereby promoting the implementation of the status quo, regardless the absolute proximity of the observed signal to the status quo.

The most interesting case arises when a conservative Advisor ($c \geq 1/2$) is subjected to additional revelation incentives ($\Psi > 0$). While the Leader always favors having better (more precise) information, conservative advisors with higher quality information do not always deliver higher quality advice. Instead, as Proposition 2 shows, they are the most likely to conceal their information to avoid significant policy changes, and that effect would only be exacerbated if the reward for revelation Ψ is low. Proposition 3 shows, then, that in a setting with low reward benefit (low Ψ), the Leader is better off with an advisor of limited competence who reveals her lower quality information, rather than with a more competent advisor who rarely provides her higher quality information.[8]

Figure 3 shows the Leader's expected utility as a function of her Advisor's competence (θ) for different values of the reward (Ψ) and Advisor conservatism (c). Other things being equal, when the advisors value the reward highly (panel (a)), the Leader's expected utility increases in the Advisor's competence. However, as the reward valuation decreases, the advisors of high competence begin to conceal more information from the Leader. Panel (b) illustrates the case in which the value of the reward Ψ is present but low, and thus the Leader gets more useful advice (and higher expected utility) from a relatively low competence Advisor.

[8] Note that the result is robust to the Leader's having some uncertainty about the Advisor's type, θ. To see this, first note that the Leader's expected utility (6) will then take expectations over θ as well as ω. Now suppose that there are two possible distributions of θ, one stochastically dominating the other. There are conditions under which the Leader would have higher expected utility with an Advisor drawn from the dominated distribution. For a more general discussion of the robustness of results to Leader uncertainty regarding θ, see Online Appendix (page 8).

Figure 3: Expected Leader's utility based on Advisor's competence, reward (Ψ), and conservatism (c).

(a) Leader's expected utility for different values of the Advisor's conservatism (c) when $\Psi = 5$

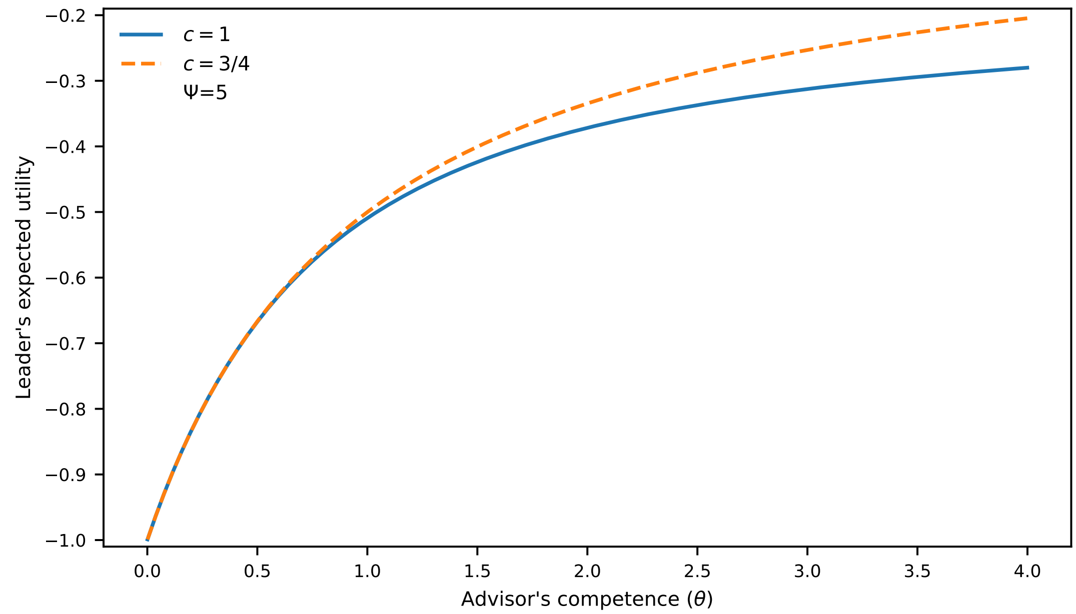

(b) Leader's expected utility for different values of the Advisor's conservatism (c) when $\Psi = 0.1$

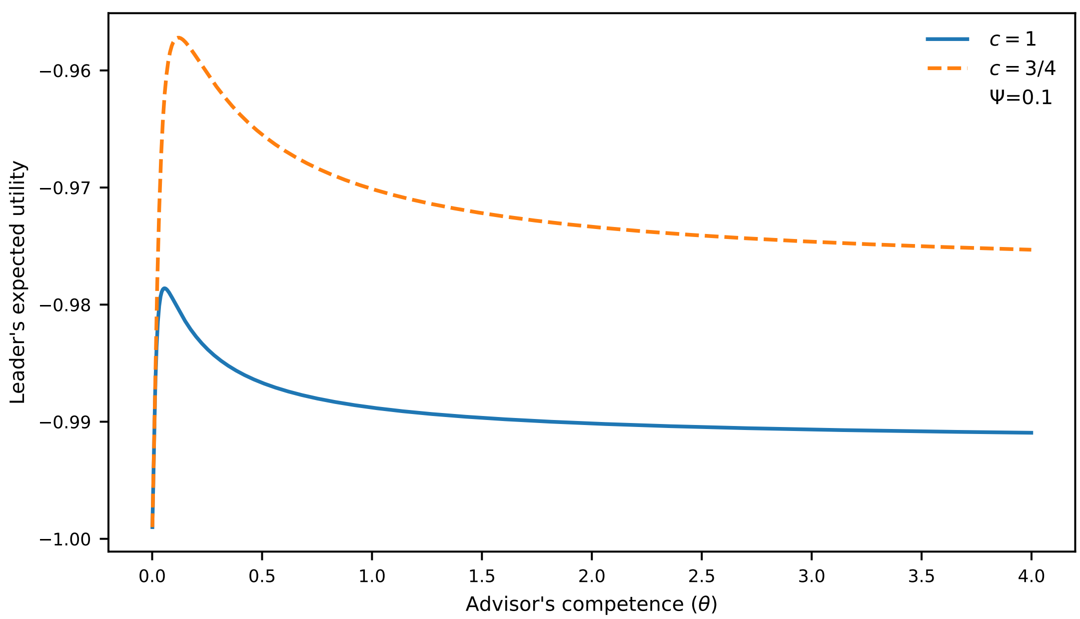

Note that the Advisor's valuation of the reward (Ψ) and her conservatism (c) essentially describe how co-aligned the Advisor is with the Leader. While the former directly measures the Advisor's incentives to please the leadership by supplying information, the latter affects the Advisor's valuation of the policies the Leader implements, which affects her revelation strategy and, thus, indirectly, affects the Leader. Our next result is consistent with the results from other models of communication: greater alignment of the Advisor's preferences with those of the Leader encourages revelation and benefits the Leader.

Proposition 4. *Both (a) the Leader's equilibrium utility holding fixed the Advisor's type, and (b) for $\Psi < \Psi^*(c)$, the type of Advisor θ that maximizes the Leader's equilibrium utility*

1. *weakly increase in the Advisor's value of the reward, Ψ;*
2. *weakly decrease in the Advisor's conservatism, c.*

Proof. See Appendix D for proof. □

We conclude our analysis by considering, in the context of our model, the possibility of bureaucrats being intentionally vague or imprecise in their communication with political leaders. We can capture and explore such a possibility by allowing the Advisor to send a (truthful) message of the form "s is positive but no more than $\bar{m}$" or "s is negative but no less than $\underline{m}$," in addition to $m = s$ and $m = \varnothing$. The use of these additional messages can be sustained in equilibrium, however, only if the Leader regards them with "skepticism," i.e., concludes from, e.g., "s is positive but no more than $\bar{m}$" that $Pr(s = \bar{m}) = 1$.[9] Given these beliefs, the Advisor chooses to send that message only when $s = \bar{m}$. Thus, in equilibrium, the meaning of the message cannot be vague. Equivalently: on the subset of possible states on which the Advisor prefers $m = s$ to $m = \varnothing$, her attempts to withhold information by sending vague messages would unravel – "less news" will be "bad news" (in the sense of Milgrom (1981)).

[9]See Online Appendix (page 8) for the proof.

Endogenous Information Acquisition

In this section, we consider an extension of our model that incorporates the Advisor's endogenous information acquisition. Maintaining the interpretation of the Advisor as a bureaucratic agency, we can think of the decision to run an information-generating policy experiment as the commissioning of a report, and the choice of the precision of the experiment as reflecting decisions about how much of the agency's resources to devote to the project.

Consider the following motivating example. The U.S. President (the Leader) wants to identify ways to save money by streamlining business operations in the Department of Defense, but does not know what specifically can be done or how much money could be saved by doing so. The department administration (the Advisor) responds by commissioning a high-powered study based on full access to all relevant data throughout the department, while taking measures to ensure that it will control the revelation of the study analysis. The results of the study come as a shock to the department leadership, identifying extraordinary levels of wasteful spending and recommending a course of action projected to save a sum far in excess of what was expected, and the department administration moves to further tighten the distribution of the study and its recommendations.[10] The games with endogenous information acquisition that we analyze in this and the next section capture some of the key analytical elements of this example, including, in equilibrium, the choice to run the maximally informative experiment and the revelation decisions after observing its results.

The game we analyze below proceeds as follows. The Advisor, of competence θ, decides whether to run a *policy experiment* which, if initiated, always provides the Advisor with an informative signal about the state of the world. If the Advisor does not initiate the experiment, she observes no informative signal about the state, the Leader chooses a policy to

[10] In the historical events on which this example is based (Whitlock and Woodward (2016)), one of DoD's key concerns appears to have been that the public revelation of the study recommendations would force the hand of the then-President and/or other principals who would influence the approval of the department's budget, including important Congressional coalitions and future Presidents. We abstract away from these details in our narrative by identifying a single principal as the Leader. As the Congressional hearings in the wake of media reports about the department's seeking to suppress the results of the study indicate, it is clear that the department's leadership was willing to risk Congress' ire by withholding the study.

implement, and the game ends. If the Advisor chooses to initiate the experiment, she chooses its precision $\tau \in [0, \theta]$ and then observes its outcome: signal $s = w + \varepsilon$, where $\varepsilon \sim N(0, 1/\tau)$. The Advisor, thus, chooses whether to run an experiment that could be relatively more or less informative, but cannot run a better experiment than her own competence permits.[11]

While the Advisor's decision to initiate the experiment is public information, the Leader does not observe the experiment's precision τ. The Leader, however, has an equilibrium conjecture, which we denote as $\hat{\tau}$, about the experiment's precision, given θ, that is correct in equilibrium. Importantly, although the Leader "knows" the distribution of τ played in equilibrium, given θ, the value of τ is not contractible. This assumption is particularly appropriate in the current setting, as agency bureaucracies have considerable discretion and relative insulation within their information-gathering and processing functions (see, e.g., the discussion of the above DoD study example in Whitlock and Woodward 2016; see also Nou, 2012; Roberts, 2009). After the Advisor observes the experiment's outcome, the Advisor chooses to reveal it, or not, to the Leader. Once the revelation stage is complete, the Leader updates the policy based on the information he has received.

In order to maintain parsimony and simplify computation of the Advisor's payoff, we set $c = 1$ in what follows – that is, we assume that the Advisor is maximally conservative,[12] so that her most preferred policy is the status quo.

From Proposition 1, we have that the Advisor's optimal revelation strategy, taking as given the running of the experiment, is to share her signal s with the Leader if and only if

$$-\sqrt{\Psi} \cdot (1 + \frac{1}{\hat{\tau}}) < s < \sqrt{\Psi} \cdot (1 + \frac{1}{\hat{\tau}}) \equiv \tilde{s}(\hat{\tau}, \Psi). \tag{7}$$

[11] We assume that there is no cost associated with information acquisition. This assumption is in line with our general focus on the indirect constraints on informative communication between the principal and the agent, in contrast to direct constraints such as the direct costs that the Advisor incurs in case of information acquisition. We relax the assumption of costless information in the Online Appendix (page 3) and show that information costs have heterogeneous effects on Advisors of different competencies, disproportionately affecting more competent Advisors, which further reinforces the quality-quantity of information trade-off we present.

[12] Note that the results that follow will be robust to marginal departures from $c = 1$.

Note that, during the revelation stage, the Advisor's competence θ does not explicitly affect the Advisor's revelation strategy, which entirely depends on Leader's interpretation of the signal.

If the Advisor runs the experiment, she chooses its precision, anticipating that she will reveal the signal if and only if (7) is satisfied. Two substantially different classes of equilibria might arise in the aftermath of this decision. In one, the Advisor selects the highest possible precision available to her, and the Leader's beliefs are consistent with this strategy ($\hat{\tau} = \theta$). In the other, the Advisor of every competence level chooses zero precision, and the Leader always believes that the precision of the signal he observes must be zero ($\hat{\tau} = 0$). While both of them constitute weak Perfect Bayesian Equilibria, only the former one is *consistent* in the sense of Kreps and Wilson (1982) – i.e., only the former is a *sequential equilibrium.*[13]

Our next proposition summarizes the Advisor's equilibrium strategy and associated comparative statics.[14]

Proposition 5. *When the information acquisition is endogenous, there is a unique sequential equilibrium in which*

1. *the Advisor always runs the experiment with maximum precision; and*

2. *the Advisor reveals the result of the experiment when it falls within the interval* $[-\tilde{s}(\theta, \Psi), \tilde{s}(\theta, \Psi)]$ *and conceals if otherwise.*

In effect, in equilibrium, holding the Leader's conjecture $\hat{\tau}$ fixed, it immediately follows from Blackwell's theorem that the Advisor has a weak preference for the most informative

[13]Formally, we refine the weak Perfect Bayesian Equilibrium and analyze the (unique) sequential equilibrium, which imposes the additional requirement of consistency

$$\Big(\sigma^*, \mu^*\Big) = \lim_{n \to +\infty} \Big(\sigma^n, \mu^n\Big), \tag{8}$$

where σ denotes the strategy profile of the Leader and the Advisor, $\Big(\sigma^n, \mu^n\Big)$ is some sequence in the set of all assessments $\Big(\sigma, \mu\Big)$, where, for every k, σ^k is totally mixed and belief μ^k is derived from σ^k using Bayes' rule.

[14]For the proofs of this and subsequent results, see Online Appendix (page 1).

experiment. To highlight this logic within the framework of the model, note that increasing the precision of the experiment increases the likelihood that its realization will fall within the disclosure interval, resulting in the Advisor earning the disclosure benefit. Thus, the Advisor will always have an incentive to choose the maximally precise experiment.

Consider next the Advisor's decision to initiate the experiment, given she and the Leader behave sequentially rationally in all subsequent stages. When the Advisor chooses not to run the experiment, she receives zero utility. However, if she decides to initiate the experiment, she can then decide whether to conceal the experiment's outcome from the Leader upon observation. In this case, the Advisor considers the Leader's reaction to every possible signal she might reveal and may opt to conceal outcomes that lead to overly drastic policy changes. Therefore, when faced with the decision of whether to run an experiment, the Advisor always chooses to do so (assuming, as we have, that running the experiment is costless).

The analysis we present in this section reinforces the competence-revelation trade-off introduced in the baseline model. Conditional on having acquired information, an Advisor of higher competence continues to have a greater incentive to conceal this information. Given that the experiment realization remains private until the Advisor decides to share its realization, the Advisors of all types always acquire information and later conceal unfavorable, for them, outcomes.

Mandated Disclosure

In our model, the trade-off between the quality of advice and the likelihood of receiving it raises the possibility that the Leader is sometimes better off with a less competent advisor. Of course, this trade-off is a consequence of particular behavioral choices by the advisors, and the Leader has a preference over those: she prefers receiving advice to not receiving it and, furthermore, values receiving higher rather than lower quality advice. This raises the question of what institutional tools the Leader could use to induce the behavioral choices she prefers to help mitigate the trade-off. In this section, we explore an institutional feature that requires the Advisor to reveal her information – the disclosure mandate – and study its equilibrium effects. We show that requiring disclosure can, indeed, mitigate the trade-off between the quality and the quantity of advice – but that that effect is fundamentally contingent on the nature of the Advisor's information. When the information is exogenously fixed, a disclosure mandate makes higher type advisors more beneficial to the Leader. But when the advisor's information is endogenous, a disclosure mandate has the opposite effect – it makes higher competence advisors (weakly) less beneficial and lowers the Leader's expected utility.

To model the disclosure mandate, we modify the baseline environment to assume that the Leader immediately learns all information the Advisor observes, and the Advisor receives the reward Ψ if and only if she is informed.

Remark 1. *Holding fixed the Advisor's information, the disclosure mandate increases the Leader's marginal utility from the Advisor's competence.*

Holding fixed the Advisor's information, a disclosure mandate fully eliminates the competence-information trade-off we saw in the baseline model. It forces the revelation of the information most valuable to the Leader, and it has the greatest impact precisely for the most competent advisors, who would, otherwise, have concealed the most information.

Panel (a) of Figure 4 contrasts the Leader's utility with no disclosure mandate and the Leader's utility with disclosure mandate, holding fixed the Advisor's information. When Advisor's information is exogenously given, the disclosure requirement always serves to improve Leader information and utility.

Figure 4: Expected Leader's utility based on Advisor's competence (θ) and reward (Ψ)

(a) Leader's expected utility with and without disclosure mandate, given exogenous information acquisition

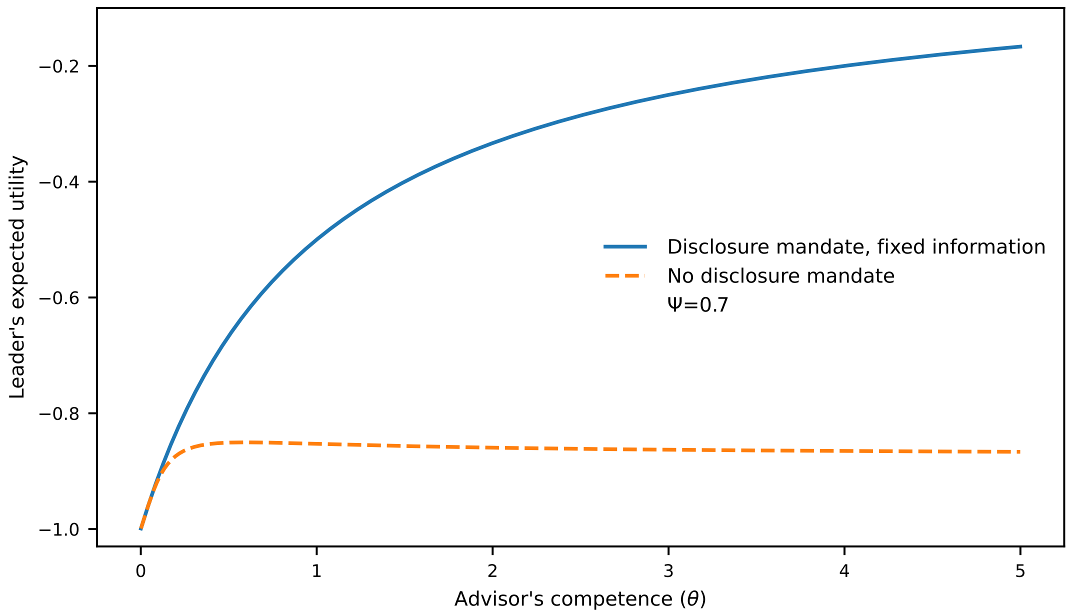

(b) Leader's expected utility with and without disclosure mandate, given endogenous information acquisition

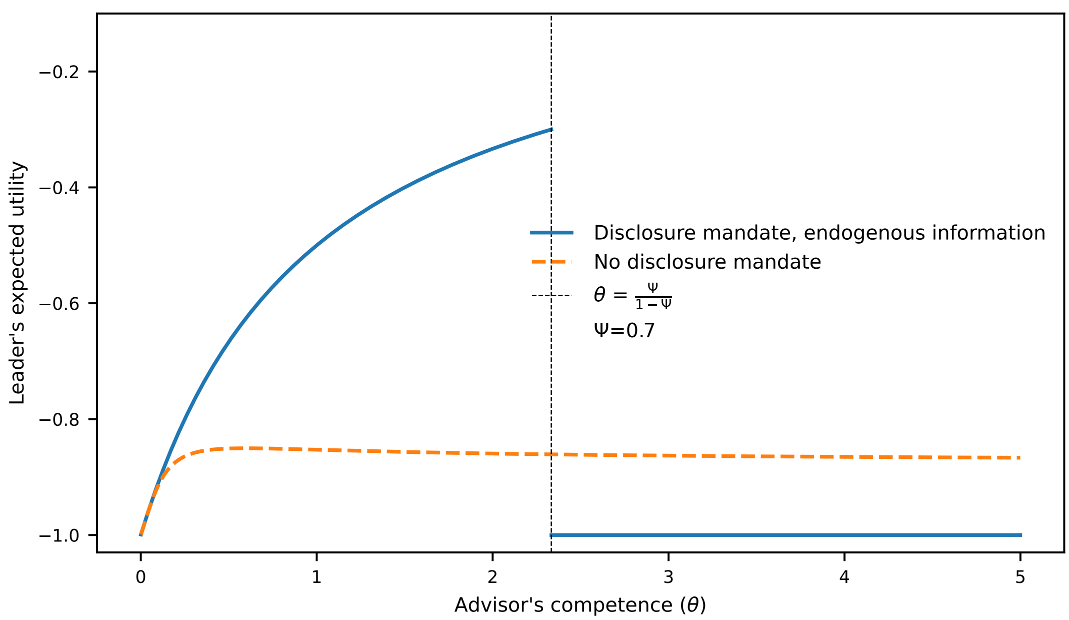

However, when we take into account the impact of the disclosure requirement on the Advisor's incentives to acquire information, the conclusion changes. As in the previous section, to introduce endogenous information acquisition, we assume that the Advisor can decide whether to acquire information (to initiate an experiment) and if she does, she chooses its quality, τ. Under the disclosure mandate, then, if she chooses to initiate the experiment, the realization of the experiment is immediately observed by the Leader. When the Advisor initiates the experiment, she always receives the reward Ψ, and she gets no reward otherwise.

Our next proposition summarizes the Advisor's equilibrium strategy and associated comparative statics.

Proposition 6. *When information acquisition is endogenous and there is a disclosure mandate,*

1. *the Advisor runs the most precise experiment possible ($\tau^* = \theta$) if her competence is below the threshold*

$$\hat{\theta}(\Psi) \equiv \begin{cases} \frac{\Psi}{1-\Psi} & \text{if } \Psi \leq 1 \\ \infty & \text{if } \Psi > 1 \end{cases} \tag{9}$$

and does not run the experiment otherwise; and

2. *the threshold $\hat{\theta}(\Psi)$ (weakly) increases in the Advisor's valuation of the reward Ψ.*

Proposition 6 highlights the re-emergence of the equilibrium trade-off between competence and advice. Importantly, the disclosure mandate changes the set of equilibrium outcomes, discouraging more competent Advisors from acquiring information relative to settings with no disclosure mandate. A key intuition is that if the high-competence Advisor is not obligated to reveal, she may still choose to reveal when the signal turns out to be close enough to the status quo. (Although, in our model, running the experiment does not have a direct cost for the Advisor, there is a range of costs that the Advisor would be willing to incur to run the experiment because the Advisor values the reward Ψ. This is true even if the disclosure of the information she acquires is relatively unlikely.) The disclosure mandate takes that discretion away from the Advisor, whose decision to run the experiment now turns on the ex ante distribution of signals. For sufficiently competent advisors, running the experiment under the disclosure mandate is a lottery with a lower expected value than the certain forfeiture of Ψ.

Note that the finding of the equilibrium impact of the disclosure mandate contrasts with the key result in Di Pei (2015). Di Pei extends the canonical cheap-talk settings to allow for endogenous information acquisition where the sender can select the information structure prior to the communication stage. Under the assumption that the acquisition cost is increasing in information precision, the sender never obtains information she might then choose to conceal; instead, she lowers the precision, reducing acquisition costs. Thus, the game always resolves in full communication, and the disclosure mandate in Di Pei (2015) does not affect the sender's equilibrium play. Gentzkow and Kamenica (2017) establish a similar result in settings with verifiable messages. A key element that explains the difference between our predictions and those of Gentzkow and Kamenica and Di Pei is the extent of flexibility in the sender's choice of information structure; the expectation of full revelation in those papers depends on such flexibility. In contrast, in our model, because the Advisor lacks the ability to commit credibly to less-than-maximal information precision, she prefers to conceal certain obtainable information, and, therefore, enforced disclosure always discourages information acquisition.

As the analysis above demonstrates, the disclosure mandate does not necessarily serve to improve Leader utility when information acquisition is endogenous. Instead, it can backfire and result in the Leader receiving less information specifically when that information would be most valuable, i.e., when the Advisor is highly competent. Thus, while one might, ignoring the strategic effects, think that the setting with highly competent advisors is the setting in which the disclosure mandate would be most valuable to the leaders, this intuition is exactly wrong when information is endogenous: leaders should be less eager to require disclosure when the advising agency is highly competent, lest it destroy the advisor's willingness to acquire information.

Our next proposition generalizes the above intuition and characterizes conditions under which the Leader is better off with the disclosure mandate than without it.

Proposition 7. *When information acquisition is endogenous and disclosure is required,*

1. *Leader utility strictly increases in Advisor competence if and only if the Advisor competence is below the threshold $\hat{\theta}(\Psi)$; and*

2. *the Leader favors having a disclosure mandate when the Advisor competence is below threshold $\hat{\theta}(\Psi)$, and prefers not having it otherwise.*

Panel (b) of Figure 4 shows Leader utility with and without the disclosure mandate, assuming endogenous information acquisition. When Advisor competence is sufficiently low, the disclosure requirement improves Leader utility. But if Advisor competence exceeds $\hat{\theta}(\Psi)$ threshold, the Advisor does not acquire information, and Leader utility is lower than it would be with no disclosure.

Discussion and Conclusion

We establish the existence of a trade-off between the quality of the information the advisor has to offer and her willingness to reveal it. The setting for this result – an advisor/agent with more conservative preferences than has the policy-maker, verifiable information, and limited rewards/punishments for the agent – is ubiquitous in modern, expertise-reliant governance. It is robust to endogenous information acquisition and is even exacerbated by endogenous information acquisition with sufficient publicity. These results suggest the need for a more nuanced approach in place of the common assumption that a policy-maker will fair better with, and thus prefer, a more competent advisor, and calls into question the presumption that the utilization of less competent advisors when more competent ones are available is a sign of corruption or dysfunction.

To be sure, especially in the context of advisors as bureaucratic agencies, advice-giving is but one of the functions that agents perform. Although, as we detailed above, one way to think about the Advisor's preferences for more conservative policies in our model is as a consequence of her bearing the costs of implementing policy changes, we abstract away from explicitly modeling policy implementation. We bracket the latter function not because we believe it to be secondary, but to focus on a strategic mechanism that has received little prior attention. A subsequent analysis would do well to consider interactions between advising and implementing explicitly, which may uncover important relationships between the strategically contingent quality of advice and the quality of policy implementation based on such advice, but such an inquiry goes beyond the scope of this paper.

As we sought to emphasize in the motivating examples and the discussion of the assumptions underlying our model, the trade-off between the quality and the quantity of information that the leader may expect from the privately informed advisor helps shed light on observations across different settings in the political economy of advice. As such, it should be of interest to political scientists beyond the broad point that more competent agents may not be optimal for the principals. How and why the trade-off we detail manifests should matter for formal analysis of political institutions, including whether or when disclosure mandates are desirable, which is one of the questions we pursue in an extension of our baseline model.

Our analysis of the disclosure mandate suggests that it has limited value from the standpoint of improving the informational quality of leaders' choices, given the incentives faced by the more competent advisors. A different possibility worth considering is that of multiple independent information sources – e.g., multiple, potentially competing, agencies. Such a possibility presents a new and complex set of strategic challenges. To see a key intuition, consider an advisor's expectation of when revealing her signal would reinforce or, alternatively, undermine the Leader's (intermediate) posterior based on the messages from other advisors. If another advisor's message aligns with one's own stance relative to the status quo, then, all else being equal, the value for an advisor in revealing her signal becomes negative. On the other hand, if another advisor's message is on the opposite side of the status quo from one's own, then the value of revealing could become positive if it prompts the Leader to conclude that the realization of the state is closer to the status quo. As this makes clear, advisors, then, have substantial incentives to collude in order to avoid sending mutually reinforcing messages. While a detailed examination of the effects of multiple advisors is beyond the scope of our analysis here, the above discussion also indicates why adding advisors is unlikely to be a magic bullet for resolving the trade-off between the quality and the quantity of advice: when advisors can coordinate, the case in which the set of advisors are high-competence and their individual signals are far from the status quo (and so, likely to be on the same side of the status quo) is one in which the advisors' incentives against revelation are particularly strong.

Another intuitive approach to mitigating the problem facing the Leader revolves around commitment in policy choice. The idea of such a commitment in the context of our model can be instructively compared to the commitment-based approach analyzed in the seminal work by Gilligan and Krehbiel (1989). Gilligan and Krehbiel show, in the context of an interaction between a legislative committee and the less informed median voter on the legislative floor, that closed rule, which limits the amendments that the floor can make to the committee's proposal before the up-or-down vote, can increase the committee's incentives to reveal its private information. The stark case of a perfectly conservative Advisor makes it clear, however, why a commitment of this kind cannot help the Leader in the setting we analyze above: Because the Advisor's preference is to maintain the status quo, and so is against revealing information to the Leader, the option of deferring to the sender's proposal is moot. But there is a different kind of commitment-based institution that could work for the Leader in our setting. If the Leader could credibly commit to making only small changes to the status quo – in other words, if the Leader could be constrained to adopt a conservative posture, in effect, bringing his induced actions closer to the Advisor's most preferred actions – then the downside of revelation for the Advisor would decrease. This approach, however, entails substantial welfare losses for the Leader: notably, unlike in Gilligan and Krehbiel's setting, the Leader is unable to take advantage of the Advisor's expertise to avoid big utility losses when the state of the world is a large departure from the prior (in our model, very high or very low). Thus, although the commitment approach may appear intuitively appealing here, it does not, at least on the basis of these very preliminary considerations, seem particularly promising from the standpoint of leader welfare. Of course, a more definitive conclusion requires a considerably more detailed analysis, which we leave to subsequent work.

Appendix

A Equilibrium Characterization

The Leader's expected utility from implementing a policy a after observing an informative message m will be

$$E[u_L(a,w)] = \int_{-\infty}^{\infty} -(x-a)^2 \cdot f_{w|w+\varepsilon}(x|m)dx$$
$$= \int_{-\infty}^{\infty} -(x-a)^2 \cdot \frac{f_{w,w+\varepsilon}(x,m)}{f_{w+\varepsilon}(m)}dx$$

where the numerator of the fraction above represents the joint probability distribution of the state of the world (w) and random noise (ε), and the denominator shows the probability distribution of the message the advisor observes and shares with the Leader ($s = w + \varepsilon$). Next,

$$E[u_L(a,w)] = \int_{-\infty}^{\infty} -(x-a)^2 \cdot \frac{f_{w,\varepsilon}(x,m-x)}{f_{w+\varepsilon}(m)}dx$$
$$= \int_{-\infty}^{\infty} -(x-a)^2 \cdot \frac{f_w(x)\cdot f_\varepsilon(m-x)}{f_{w+\varepsilon}(m)}dx =$$
$$= \int_{-\infty}^{\infty} -(x-a)^2 \cdot \frac{\frac{1}{\sqrt{2\cdot\pi}} \cdot e^{-\frac{1}{2}\cdot x^2} \cdot \frac{\sqrt{\theta}}{\sqrt{2\cdot\pi}} \cdot e^{-\frac{1}{2}\cdot(m-x)^2\cdot\theta}}{\frac{1}{\sqrt{2\cdot\pi}} \cdot \frac{1}{\sqrt{1+1/\theta}} \cdot e^{-\frac{1}{2}\cdot\frac{m^2}{1+1/\theta}}}dx$$
$$= \int_{-\infty}^{\infty} -(x-a)^2 \cdot \frac{\sqrt{1+\theta}\cdot e^{-\frac{(m-x)^2\cdot\theta+x^2}{2}}}{\sqrt{2\cdot\pi}\cdot e^{-\frac{1}{2}\cdot\frac{m^2}{1+1/\theta}}}dx,$$

The expected utility the Leader gets is

$$E[u_L(a,w)] = -\frac{1+\theta+(a\cdot(1+\theta)-m\cdot\theta)^2}{(1+\theta)^2}.$$

Because

$$\frac{\partial E[u_L(a,w)]}{\partial a} = -\frac{2\cdot(a\cdot(1+\theta)-s\cdot\theta)}{1+\theta},$$

the Leader who observes $m \neq \varnothing$ selects a policy $a^*(m \neq \varnothing) = \frac{m}{1+1/\theta}$. Let us now denote the policy the Leader selects in the absense of revelation as $d(\theta)$.

The Advisor sends an informative message when her expected utility from revealing this information exceeds her expected utility from concealing it.

$$E[u_A(m = s, w, a^*(m))] = -(\frac{s}{1+1/\theta} - (1-c) \cdot \frac{s}{1+1/\theta})^2 + \Psi$$
$$> E[u_A(m = \varnothing, w, a^*(m))] = -(d(\theta) - (1-c) \cdot \frac{s}{1+1/\theta})^2.$$

The Advisor reveals her information to the Leader when

$$-\sqrt{\Psi + d(\theta)^2} \cdot (1 + \frac{1}{\theta}) \cdot \frac{1}{\sqrt{2 \cdot c - 1}} < s < \sqrt{\Psi + d(\theta)^2} \cdot (1 + \frac{1}{\theta}) \cdot \frac{1}{\sqrt{2 \cdot c - 1}}$$

note that because the Advisor's strategy is symmetric around 0, in the absence of informative message, the Leader's optimal strategy is to implement $d^*(\theta) = 0$. Finally, when $c < 1/2$, the Advisor reveals all her information to the Leader, because $-c^2 \cdot (\frac{s}{1+1/\theta})^2 + \Psi$ always exceeds $-(1-c)^2 \cdot (\frac{s}{1+1/\theta})^2$ for $c < 1/2$.

B Leader Welfare

The Leader's equilibrium expected utility

$$E[u_L^*(\theta)] = \underbrace{Pr[s \in (-\hat{s}, \hat{s})]}_{\text{Advisor sends informative message}} \cdot \underbrace{\frac{-1}{1+\theta}}_{\text{Leader's expected utility after informative message}} + \underbrace{Pr[s \notin (-\hat{s}, \hat{s})]}_{\text{Advisor does not send informative message}} \cdot E[-w^2 | s \notin (-\hat{s}, \hat{s})],$$

where

$$A \equiv Pr[s \in [-\hat{s}, \hat{s}]] \cdot \frac{-1}{1+\theta}$$
$$= \big(\Phi(\hat{s}/\sqrt{1+1/\theta}) - \Phi(-\hat{s}/\sqrt{1+1/\theta})\big) \cdot \frac{-1}{1+\theta}.$$

and

$$B \equiv Pr[s \notin (-\hat{s}, \hat{s})] \cdot E[-w^2 | s \notin [-\hat{s}, \hat{s}]]$$

$$= Pr[s < -\hat{s}] \cdot E[-w^2 | s < -\hat{s}] + Pr[s > \hat{s}] \cdot E[-w^2 | s > \hat{s}]$$

$$= Pr[s < -\hat{s}] \cdot (\int_{-\infty}^{\infty} \int_{-\infty}^{-\hat{s}-y} -x^2 \frac{f_{w,\varepsilon}(x,y)}{Pr[s < -\hat{s}]} dxdy)$$

$$+ Pr[s > \hat{s}] \cdot (\int_{-\infty}^{\infty} \int_{\hat{s}-y}^{\infty} -x^2 \frac{f_{w,\varepsilon}(x,y)}{Pr[s > \hat{s}]}) dxdy)$$

$$= (\int_{-\infty}^{\infty} \int_{-\infty}^{-\hat{s}-y} -x^2 \frac{1}{2\pi} \frac{1}{\sqrt{1/\theta}} e^{-\frac{1}{2}(x^2 + \frac{y^2}{1/\theta})} dxdy$$

$$+ \int_{-\infty}^{\infty} \int_{\hat{s}-y}^{\infty} -x^2 \frac{1}{2\pi} \frac{1}{\sqrt{1/\theta}} e^{-\frac{1}{2}(x^2 + \frac{y^2}{1/\theta})} dxdy)$$

$$= \frac{1}{2\pi} \frac{1}{\sqrt{1/\theta}} \int_{-\infty}^{\infty} -\sqrt{2\pi} \cdot e^{\frac{\theta \cdot y^2}{2}} \Big((\hat{s} - y) \cdot \phi\big(\hat{s} - y\big)$$

$$+ (\hat{s} + y) \cdot \phi\big(\hat{s} + y\big) + 2 - \Phi(\hat{s} - y) - \Phi(\hat{s} + y) \Big) dy$$

$$= \int_{-\infty}^{\infty} - \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{1/\theta}} \cdot e^{-\frac{\theta \cdot y^2}{2}} \Big((\hat{s} - y) \cdot \phi\big(\hat{s} - y\big)$$

$$+ (\hat{s} + y) \cdot \phi\big(\hat{s} + y\big) + 2 - \Phi(\hat{s} - y) - \Phi(\hat{s} + y) \Big) dy.$$

Let us denote

$$g(a) \equiv \int_{-\infty}^{\infty} - \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{1/\theta}} \cdot e^{-\frac{\theta \cdot y^2}{2}} \Big((\hat{s} - y) \cdot \phi\big(\hat{s} - y\big) + (\hat{s} + y) \cdot \phi\big(\hat{s} + y\big)$$

$$+ 2 - \Phi(\sqrt{a} \cdot (\hat{s} - y)) - \Phi(\sqrt{a} \cdot (\hat{s} + y)) \Big) dy.$$

Because $g(1) = B$ we need to compute $g(1)$. We begin with $g(0)$

$$g(0) = \int_{-\infty}^{\infty} - \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{1/\theta}} e^{-\frac{\theta \cdot y^2}{2}} \Big((\hat{s} - y) \cdot \phi\big(\hat{s} - y\big) + (\hat{s} + y) \cdot \phi\big(\hat{s} + y\big) + 1 \Big) dy$$

$$= -1 - \frac{e^{\frac{-\hat{s}^2}{2(1+1/\theta)}} \sqrt{2/\pi} \hat{s}}{\sqrt{(1 + 1/\theta)^3}}. \tag{10}$$

Now we compute $\frac{\partial g(a)}{\partial a}$. By Leibniz integral rule[15]

$$\begin{aligned}
\frac{\partial g(a)}{\partial a} &= \frac{\partial}{\partial a}\int_{-\infty}^{\infty} -\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}}\cdot e^{-\frac{\theta\cdot y^2}{2}}\big((\hat{s}-y)\cdot\phi(-(\hat{s}-y))+(\hat{s}+y)\cdot\phi(-(\hat{s}+y)) \\
&\qquad +2-\Phi(\sqrt{a}\cdot(\hat{s}-y))-\Phi(\sqrt{a}\cdot(\hat{s}+y))\big)dy \\
&= \int_{-\infty}^{\infty}\frac{\partial}{\partial a}\Big(-\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}}\cdot e^{-\frac{\theta\cdot y^2}{2}}\big((\hat{s}-y)\cdot\phi(-(\hat{s}-y))+(\hat{s}+y)\cdot\phi(-(\hat{s}+y)) \\
&\qquad +2-\Phi(\sqrt{a}\cdot(\hat{s}-y))-\Phi(\sqrt{a}\cdot(\hat{s}+y))\big)\Big)dy \\
&= \int_{-\infty}^{\infty} -\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}}e^{-\frac{\theta\cdot y^2}{2}}\Big(-\phi(\sqrt{a}\cdot(\hat{s}-y))\frac{1}{2\sqrt{a}}(\hat{s}-y)-\phi(\sqrt{a}\cdot(\hat{s}+y))\frac{1}{2\sqrt{a}}(\hat{s}+y)\Big)dy \\
&= \int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}}e^{-\frac{\theta\cdot y^2}{2}}\frac{1}{2\sqrt{a}}\Big(\frac{e^{-\frac{a(\hat{s}-y)^2}{2}}}{\sqrt{2\pi}}(\hat{s}-y)+\frac{e^{-\frac{a(\hat{s}+y)^2}{2}}}{\sqrt{2\pi}}(\hat{s}+y)\Big)dy \\
&= \int_{-\infty}^{\infty} e^{-\frac{\theta\cdot y^2}{2}}\sqrt{\theta}\cdot\frac{e^{-\frac{a(\hat{s}-y)^2}{2}}(\hat{s}-y)+e^{-\frac{a(\hat{s}+y)^2}{2}}(\hat{s}+y)}{4\pi\sqrt{a}}dy \\
&= \hat{s}\cdot\frac{e^{-\frac{a\hat{s}^2\theta}{2(1+\frac{a}{\theta})}}}{\sqrt{2\pi}\sqrt{a}\cdot(1+\frac{a}{\theta})^{3/2}}.
\end{aligned}$$

We now integrate $\frac{dg(a)}{da}$ wrt a

$$\begin{aligned}
g(a) &= \int \hat{s}\cdot\frac{e^{-\frac{a\hat{s}^2\theta}{2(a+\theta)}}}{\sqrt{2\pi}\sqrt{a}\cdot(1+\frac{a}{\theta})^{3/2}}da \\
&= 2\Phi(\frac{\hat{s}}{\sqrt{\frac{1}{a}+\frac{1}{\theta}}})-1+C,
\end{aligned} \tag{11}$$

where C is unknown constant.

Note that

$$g(a=0) = -1-\frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}} \tag{12}$$

and

$$g(a=0) = \lim_{a\to 0} 2\cdot\Phi(\frac{\hat{s}}{\sqrt{\frac{1}{a}+\frac{1}{\theta}}})-1+C = C, \tag{13}$$

where equation 12 follows from equation 10 and equation 13 follows from equation 11. There-

[15]Leibniz integral rule is applicable because the integral of partial derivative converges.

fore

$$C = -1 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}}.$$

Finally,

$$B = \frac{Pr[s \notin [-\hat{s}, \hat{s}]]}{Pr[s < -\hat{s}]} \cdot g(a=1)$$

$$= \lim_{a \to 1} 2 \cdot \Phi(\frac{\hat{s}}{\sqrt{\frac{1}{a} + \frac{1}{\theta}}}) - 2 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}}$$

$$= (2 \cdot \Phi(\frac{\hat{s}}{\sqrt{1 + \frac{1}{\theta}}}) - 2 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}})$$

Thus

$$E[u_L^*(\theta)] = \big(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}}) - 1\big) \cdot \frac{-1}{1+\theta}$$

$$+(2\Phi(\frac{\hat{s}}{\sqrt{1 + \frac{1}{\theta}}}) - 2 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}}).$$

C Welfare-Maximizing Competence

1. $c \leq 1/2$:

$$[-\hat{s}(\theta, c, \cdot), \hat{s}(\theta, c, \cdot)] = (-\infty, \infty).$$

$$\Rightarrow \frac{\partial E[u_L^*(\theta)]}{\partial \theta} > 0 \;\; \& \;\; \arg\max_{\theta}(E[u_L^*(\theta)]) = \{\infty\}.$$

2. $c > 1/2$:

It will help simplify exposition to denote $K \equiv \frac{\Psi}{2 \cdot c - 1}$. Note that $c > 1/2$ implies $K \geq 0$.

Parts a, b. Let $\Psi > 0$ ($\Rightarrow K > 0$), $D(\theta, K) \equiv \frac{\partial E[u_L^*(\theta)]}{\partial \theta}$.

Claim 1. $\forall K : K \in (0, 1)$, *the Leader's equilibrium welfare is greatest with an Advisor of finite competence*

Proof.

$$D(\theta, K) = \frac{1}{2(1+\theta)^2}\Big(-\frac{\sqrt{2}e^{-\frac{K\cdot(1+1/\theta)}{2}}\sqrt{\frac{K\cdot\theta\cdot(1+\theta)}{\pi}}(2\cdot\theta + K\cdot(1+\theta))}{\theta^2} + \underbrace{2(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}})-1)}_{>0}\Big). \tag{14}$$

To establish that, for $K \in (0,1)$, an interior maximum of $E[u_L^*(\cdot)]$ exists and is the global maximum, we first examine the limits of $D(\theta, K)$ as θ approaches 0 and as θ approaches ∞. When the former is positive and the latter negative, an interior maximum must exist and, furthermore, there cannot be a local maximum at $\theta = \infty$. First,

$$\lim_{\theta\to 0} D(\theta, K) = 1.$$

Next,

$$\begin{gathered}\lim_{\theta\to\infty} D(\theta, K) = \pm 0, \text{ and } \operatorname{sgn}(\lim_{\theta\to\infty} D(\theta, K)) = \operatorname{sgn}(F(K)),\\ \text{where } F(K) \equiv -2 - e^{-\frac{K}{2}}\sqrt{\frac{2}{\pi}}\sqrt{K}(2+K) + 4\Phi(\sqrt{K}).\end{gathered} \tag{15}$$

To establish that $\exists\theta :\ D(\theta, K) = 0$, we must identify conditions on K s.t. $F(K) < 0$. From (15), $F(K = 0) = 0$ and $\frac{dF(K)}{dK} < 0 \iff K \in (0,1)$;

$$\Rightarrow \forall K : K \in (0,1),\ F(K) < 0$$

$$\Rightarrow \forall K : K \in (0,1),\ \lim_{\theta\to\infty} D(\theta, K) < 0.$$

By continuity of $E[u_L^*(\theta)]$

$$\forall K \in (0,1)\ \exists\,\hat{\theta} : \arg\max_{\theta} E[u_L^*(\theta)] = \{\hat{\theta} : \hat{\theta} < \infty\}. \tag{16}$$

□

Claim 2. *When $K \to \infty$, the Leader's equilibrium welfare is greatest with an Advisor of* *infinite competence.*

Proof. From (14)

$$\lim_{K\to\infty} D(\theta, K) = \frac{2}{2\cdot(1+\theta)^2} > 0 \Rightarrow \arg\max_{\theta} E[u_L^*(\theta, K\to\infty)] = \{\infty\}. \tag{17}$$

□

Next, to establish that

$$\exists! K^* : \arg\max_{\theta} E[u_L^*(\theta)] = \begin{cases} \hat{\theta} : \hat{\theta} < \infty \; if \; K < K^*, \\ \infty \; if \; K > K^*. \end{cases} \tag{18}$$

We require the following lemma:

Lemma 1. *1.* $\exists\hat{\theta} \Rightarrow \frac{d\hat{\theta}(K)}{dK} > 0$*;*

2. $E[u_L^*(\theta)]$ *has at most one interior maximum* $(\hat{\theta})$*.*

Proof. By the implicit function theorem

$$\frac{d\hat{\theta}(K)}{dK} = -\frac{\partial_K D(\theta, K)}{\partial_\theta D(\theta, K)}.$$

Because $\hat{\theta}(K)$ is a maximum, $\partial_\theta D(\theta, K) < 0$.

$$\Rightarrow \operatorname{sgn}(-\frac{\partial_K D(\theta, K)}{\partial_\theta D(\theta, K)}) = \operatorname{sgn}\big(\partial_K D(\theta, K)\big).$$

$$\partial_K D(\theta, K) = \underbrace{\Big(\frac{e^{-\frac{K\cdot\theta\cdot(1+\theta)}{2\cdot\theta}}\cdot K}{2\cdot\theta^2\cdot\sqrt{2\pi}\cdot\sqrt{K\cdot\theta\cdot(1+\theta)}}\Big)}_{>0}\cdot(K\cdot(1+\theta)-\theta),$$

$$\Rightarrow \operatorname{sgn}(\frac{\partial\hat{\theta}}{\partial K}) = \operatorname{sgn}(K\cdot(1+\theta)-\theta),$$

$$\Rightarrow \begin{cases} \frac{\partial\hat{\theta}}{\partial K} > 0 \iff K\cdot(1+\theta)-\theta > 0, \\ \frac{\partial\hat{\theta}}{\partial K} < 0 \iff K\cdot(1+\theta)-\theta < 0, \end{cases}$$

$$\Rightarrow \arg\min_K D(\theta, K) = \{\frac{\theta}{1+\theta}\}.$$

From (14), $D(\theta=0,K)=\frac{1}{2}>0$ & $D(\theta=\frac{K}{1-K},K)=-\frac{(1-K)^2(3\sqrt{2}+\sqrt{e\cdot\pi}\cdot 2\cdot(1-2\cdot\Phi(1))}{2\cdot\sqrt{e\cdot\pi}}<0 \Rightarrow$

$$\forall K > \frac{\theta}{1+\theta},\ \exists\ \hat{\theta} : \hat{\theta} < \frac{K}{1-K}\ \ \&\ \ \frac{\partial\hat{\theta}}{\partial K} > 0. \tag{19}$$

Next, we prove that if $\exists\ \hat{\theta}$ then $\exists!\ \hat{\theta}$. Because $\forall K > \frac{\theta}{1+\theta}$, $\exists\ \hat{\theta} : \hat{\theta} < \frac{K}{1-K}$ & $\frac{\partial\hat{\theta}}{\partial K} > 0$, once we prove uniqueness of $\hat{\theta}$, then $\forall\hat{\theta}$, $\frac{\partial\hat{\theta}}{\partial K} > 0$. Note that

From (14),

$$\begin{aligned} &\operatorname{sgn}(D(\theta,K)) \\ &= \operatorname{sgn}(-\frac{\sqrt{2}e^{-\frac{K\cdot(1+1/\theta)}{2}}\sqrt{\frac{K\cdot\theta\cdot(1+\theta)}{\pi}}(2\cdot\theta+K\cdot(1+\theta))}{\theta^2} \\ &+ 2(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}})-1)). \end{aligned}$$

Let

$$Interior(\theta) \equiv -\frac{\sqrt{2}e^{-\frac{K\cdot(1+1/\theta)}{2}}\sqrt{\frac{K\cdot\theta\cdot(1+\theta)}{\pi}}(2\cdot\theta+K\cdot(1+\theta))}{\theta^2} + 2(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}})-1)$$

$$\frac{\partial Interior(\theta)}{\partial\theta} = (\frac{e^{-\frac{K\cdot(1+1/\theta)}{2}}\sqrt{\frac{K\cdot\theta\cdot(1+\theta)}{2\pi}}}{\theta^4})\cdot K\cdot(\theta-K\cdot(1+\theta)). \tag{20}$$

$$\operatorname{sgn}(20) = \operatorname{sgn}(\theta-K(1+\theta)).$$

From (20),

$$\begin{cases} \frac{\partial Interior(\theta)}{\partial\theta} < 0\ \forall\ \theta : \theta < \frac{K}{1-K} \\ \frac{\partial Interior(\theta)}{\partial\theta} > 0\ \forall\ \theta : \theta \geq \frac{K}{1-K}. \end{cases}$$

It follows that $Interior(\theta)$ and, thus, $D(\theta, K)$ can switch sign from positive to negative only once. Therefore, $E[u_L^*(\theta)]$ has at most one interior maximum $\Rightarrow \partial_K \hat{\theta} > 0$. $\square$

From (19) and the Lemma (1), $\exists \hat{\theta} \iff \theta < \frac{K}{1-K} \iff K > \frac{\theta}{1+\theta}$. Thus, conditional on the interior maximum's existence,

$$\partial_K D(\theta, K) = \frac{e^{-\frac{K \cdot \theta \cdot (1+\theta)}{2 \cdot \theta}} \cdot K}{2 \cdot \theta^2 \cdot \sqrt{2\pi} \cdot \sqrt{K \cdot \theta \cdot (1+\theta)}} \cdot (K \cdot (1+\theta) - \theta) > 0. \tag{21}$$

From Claim (1), (17), and monotonicity of $D(\theta, K)$ wrt K from inequality (21), we obtain that (18) must hold. Because $K \equiv \frac{\Psi}{2c-1}$, the statement of the proposition follows.

Part c. For $\Psi = 0$, $[-\hat{s}(\cdot), \hat{s}(\cdot)] = \{\varnothing\} \; \forall \, \theta$.

D Comparative Statics

Part b of Proposition 4 for $c > 1/2$ & $\Psi > 0$, follows from Lemma 1. When $c \leq 1/2$,

$$\arg\max_{\theta} E[u_L^*(\theta, \cdot)] = \{\infty\}.$$

When $\Psi = 0$, the Leader's utility does not depend on θ.

References

Ashworth, Scott and Greg Sasso. 2019. "Delegation to an overconfident expert." *The Journal of Politics* 81(2):692–696.

Austen-Smith, David. 1990. "Information transmission in debate." *American Journal of political science* pp. 124–152.

Austen-Smith, David. 1993. "Interested experts and policy advice: Multiple referrals under open rule." *Games and Economic Behavior* 5(1):3–43.

Aybas, Yunus, C. and Steven Callander. 2023. "Efficient Cheap Talk in Complex Environments." *Working paper* .

Battaglini, Marco. 2002. "Multiple referrals and multidimensional cheap talk." *Econometrica* 70(4):1379–1401.

Bhattacharya, Sourav and Arijit Mukherjee. 2013. "Strategic information revelation when experts compete to influence." *The RAND Journal of Economics* 44(3):522–544.

Buisseret, Peter and Carlo Prato. 2016. "Electoral control and the human capital of politicians." *Games and Economic Behavior* 98:34–55.

Calvert, Randall L. 1985. "The value of biased information: A rational choice model of political advice." *The Journal of Politics* 47(2):530–555.

Canes-Wrone, Brandice, Michael C Herron and Kenneth W Shotts. 2001. "Leadership and pandering: A theory of executive policymaking." *American Journal of Political Science* pp. 532–550.

Che, Yeon-Koo and Navin Kartik. 2009. "Opinions as incentives." *Journal of Political Economy* 117(5):815–860.

Crawford, Vincent P and Joel Sobel. 1982. "Strategic information transmission." *Econometrica: Journal of the Econometric Society* pp. 1431–1451.

Di Pei, Harry. 2015. "Communication with endogenous information acquisition." *Journal of Economic Theory* 160:132–149.

Dziuda, Wioletta. 2011. "Strategic argumentation." *Journal of Economic Theory* 146(4):1362–1397.

Egorov, Georgy and Konstantin Sonin. 2011. "Dictators and their viziers: Endogenizing the loyalty–competence trade-off." *Journal of the European Economic Association* 9(5):903–930.

Egorov, Georgy and Konstantin Sonin. 2023. "Why Did Putin Attack Ukraine? A Theory of Degenerative Autocracy." *NBER Working Paper 31187* April.

Fox, Justin. 2007. "Government transparency and policymaking." *Public choice* 131(1-2):23–44.

Gailmard, Sean and John W Patty. 2007. "Slackers and zealots: Civil service, policy discretion, and bureaucratic expertise." *American Journal of Political Science* 51(4):873–889.

Gailmard, Sean and John W Patty. 2012. "Formal models of bureaucracy." *Annual Review of Political Science* 15(1):353–377.

Gailmard, Sean and John W Patty. 2013. "Stovepiping." *Journal of Theoretical Politics* 25(3):388–411.

Gentzkow, Matthew and Emir Kamenica. 2017. "Disclosure of endogenous information." *Economic Theory Bulletin* 5(1):47–56.

Gilligan, Thomas W and Keith Krehbiel. 1989. "Asymmetric information and legislative rules with a heterogeneous committee." *American journal of political science* pp. 459–490.

Huq, Aziz and Tom Ginsburg. 2018. "How to lose a constitutional democracy." *UCLA L. Rev.* 65:78.

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101(6):2590–2615.

Kaufman, Herbert. 1981. "Fear of bureaucracy: A raging pandemic." *Public Administration Review* pp. 1–9.

Kreps, David M. and Robert Wilson. 1982. "Sequential Equilibria." *Econometrica* 50(4):863–894.

Lagerlöf, Nils-Petter. 2012. "A Dynamic Theory of Competence, Loyalty and Stability in Dictatorships." *The B.E. Journal of Macroeconomics* 12(1).

Milgrom, Paul. 1981. "Good News and Bad News: Representation Theorems and Applications." *The Bell Journal of Economics* 12(2):380–391.

Milgrom, Paul. 2008. "What the seller won't tell you: Persuasion and disclosure in markets." *Journal of Economic Perspectives* 22(2):115–131.

Nou, Jennifer. 2012. "Agency self-insulation under presidential review." *Harv. L. Rev.* 126:1755.

Patty, John W. 2009. "The politics of biased information." *The Journal of Politics* 71(2):385–397.

Patty, John W and Ian R Turner. 2021. "Ex Post Review and Expert Policy Making: When Does Oversight Reduce Accountability?" *The Journal of Politics* 83(1):23–39.

Prendergast, Canice and Robert H Topel. 1996. "Favoritism in organizations." *Journal of Political Economy* 104(5):958–978.

Roberts, Patrick S. 2009. "How security agencies control change: Executive power and the quest for autonomy in the FBI and CIA." *Public Organization Review* 9(2):169–198.

Shin, Hyun Song. 1994. "The burden of proof in a game of persuasion." *Journal of Economic Theory* 64(1):253–264.

Sobel, Joel. 1993. "Information control in the principal-agent problem." *International Economic Review* pp. 259–269.

Sobel, Joel. 2013. "Giving and receiving advice." *Advances in economics and econometrics* 1:305–341.

Terai, Kimiko and Amihai Glazer. 2018. "Rivalry among agents seeking large budgets." *Journal of Theoretical Politics* 30(4):388–409.

Turner, Ian R. 2022. "Reviewing Procedure versus Judging Substance: How Increasing Bureaucratic Oversight Can Reduce Bureaucratic Accountability." *Journal of Political Institutions and Political Economy* 2(4):569–596.

Whitlock, Craig and Bob Woodward. 2016. "Pentagon buries evidence of $125 billion in bureaucratic waste." *The Washington Post, Accessed December 6th* .

Wolinsky, Asher. 2003. "Information transmission when the sender's preferences are uncertain." *Games and Economic Behavior* 42(2):319–326.

Zakharov, Alexei V. 2016. "The loyalty-competence trade-off in dictatorships and outside options for subordinates." *The Journal of Politics* 78(2):457–466.