The Perils of Friendly Oversight*

Dino Gerardi[†]

Edoardo Grillo[‡]

Ignacio Monzón[§]

May 23, 2022

Abstract

Decision makers often rely on experts' evaluations to decide on complex proposals. Proponents want the approval of their proposals and can work to improve their quality. The scrutiny of experts ought to push proponents to work harder, leading to high-quality proposals. Experts, however, have their own agendas: they may favor or oppose the proposals under their scrutiny. We study how the expert's agenda affects the likelihood that proposals are approved and their quality. We show that an expert in favor of a proposal can be detrimental towards its approval. This happens when it is easy to incentivize the proponent to work and when the status quo alternative is not too attractive.

Keywords: information transmission, moral hazard, oversight, persuasion

JEL Classification: D72, D73, D82, D83

^{*}We thank the editor and two anonymous referees for their thoughtful reports and suggestions. We also thank Jacopo Bizzotto, Federico Boffa, Leandro Gorno, Juan S. Morales, Santiago Oliveros, Amedeo Piolatto and Carlo Prato for valuable comments.

[†]Collegio Carlo Alberto and University of Turin (ESOMAS Department); dino.gerardi@carloalberto.org

[‡]University of Padua (Department of Economics and Management "Marco Fanno"); edoardo.grillo@unipd.it

[§]Collegio Carlo Alberto and University of Turin (ESOMAS Department); ignacio@carloalberto.org

1 Introduction

Oftentimes, decision makers must decide on complex proposals but lack the necessary expertise to evaluate them. They can then resort to experts with the technical competence to correctly assess proposals. For example, a local government who must decide whether to authorize the development of a potentially polluting firm relies on specialized agencies to assess the firm's environmental plan. Similarly, when the CEO of a company proposes an acquisition, shareholders decide only after receiving the due diligence from a consulting company. Finally, when the executive branch proposes a bill, the legislature holds hearings where experts evaluate the bill, and only afterwards the legislature votes on it.

In many economic and political environments, there is a conflict of interest between those submitting proposals and decision makers. While those submitting proposals want to push them through, decision makers only want to approve proposals of high quality. In these environments, the scrutiny by expert authorities ought to press proponents to design better proposals and, ultimately, ought to improve the quality of implemented proposals. For example, the oversight of expert agencies should make firms strive to reduce their environmental impact. Similarly, due diligence should push CEOs to work hard and propose acquisitions that create value for shareholders. Finally, hearings during the legislative process should press politicians to craft high-quality bills.

Authorities, however, may have their own agendas. A captured agency may look the other way to help get a development plan approved even if it poses environmental hazards. A standing committee controlled by the minority can block reforms, even when the society would benefit from them.

We study how the agenda of oversight authorities affects the likelihood of approval of proposals. On the one side, when the authority wants the approval of a proposal, its scrutiny is lenient. This directly increases the likelihood of approval. On the other side, when the authority wants the rejection of a proposal, its scrutiny is tough. While a tough scrutiny directly reduces the likelihood of approval, it also forces the proponent to work hard, which improves the quality of the proposal. This, in turn, has the indirect effect of increasing the likelihood of approval. We study these two countervailing forces and characterize the impact of the authority's agenda. We find that friendly authorities can have the unintended effect of preventing the adoption of proposals, and that this happens when high-quality proposals are needed the most.

To fix ideas, we focus on the following simple environment. A politician (e.g., the executive) proposes a reform and a constituency (e.g., the legislature) then decides whether to approve it or to reject it, maintaining a *status quo*. The reform can be either of high quality or of low quality. The constituency prefers the reform to the status quo only if the reform is of high quality. The politician always wants the reform to be approved. At the onset of the game, the politician exerts costly and unobservable effort to draft the reform. This effort determines the expected quality of the reform.

The constituency does not observe the quality of the reform, but must instead rely on a report from an authority (e.g., the legislative committee). The authority has specific expertise and can thus distinguish between reforms of high and of low quality. The authority establishes procedures to issue the report. Authorities, for example, decide the composition and ideological bias of those providing testimony during hearings. Similarly, authorities can submit more or less detailed reports about the consequences of a reform. Typically, constituencies can observe these procedures: interested parties have access to the minutes of the hearings, observe the specific enquiries of the authority and know who is selected to provide testimony. We model these features by assuming that the authority commits to a *reporting strategy*: a mapping from the quality of the reform to a set of reports. A useful interpretation for the reporting strategy is that the authority chooses a procedure which imperfectly assesses the quality of the proposal. We also assume that the constituency observes both the reporting strategy and the actual report issued by the authority.

We consider three types of oversight authorities. First, a *truthful authority* always reports the true quality of the reform. Second, a *friendly authority* is strategic and wants the reform to be approved. Third, an *unfriendly authority* is also strategic but wants the reform to be rejected. In some cases, the constituency may rely on authorities renowned to be non-partisan, like the Congressional Budget Office (truthful authority).¹ In other cases, the authorities that evaluate the reforms may represent the interests of the party of the politician (friendly authority) or of the opposition (unfriendly authority): they may have been appointed by the current administration

¹Even nominally non-partisan authorities are often accused of being politically biased. See, for instance, the accusation from the Trump administration against the NIAID director Dr. Fauci during the COVID-19 pandemic (https://cutt.ly/qvR05xz).

or by a previous one, now at the opposition. The type of the authority is given and known to all players.

The party of the politician is an inactive player representing the interest group pushing for the reform. Similarly to the politician, the party wants the reform to be approved, but differently from her, the party does not internalize the cost of effort. Therefore, the party's payoff is equal to the probability that the reform is approved. To simplify the exposition, we let the *approval threshold* denote the payoff that the constituency obtains from the status quo. Thus, the constituency prefers to approve the reform when its expected quality is greater than the approval threshold.

The type of authority affects the party's payoff through two channels. First, it directly affects the informativeness of the reports available to the constituency. Second, it modifies the politician's incentives to exert effort. Both the informativeness of the report and the level of effort exerted by the politician affect the decision of the constituency and thus the party's payoff.

For a fixed level of effort, the likelihood of approval is higher under the friendly than under the unfriendly authority. The friendly authority commits to a lenient reporting strategy. Its procedure generates false positives: with positive probability the reform is deemed of high quality even when its quality is low. Instead, the unfriendly authority commits to a tough reporting strategy. Its procedure generates false negatives: with positive probability the reform is deemed of low quality even when its quality is high. Then, holding the level of effort fixed, the party would be better off under the friendly authority.

The type of authority also affects the level of effort that the politician exerts in equilibrium. When facing the friendly authority, the politician takes advantage of its favorable reporting strategy by exerting lower effort. When instead facing the unfriendly authority, the politician must work hard to compensate its unfavorable reporting strategy. To sum up, the politician exerts a higher level of effort under the unfriendly authority than under the friendly one.

The party's preference over strategic authorities depends on the relative weight of these countervailing forces. On the one hand, the party benefits from the direct effect of a friendly oversight. On the other hand, the party benefits from the indirect effect of an unfriendly oversight, as it forces the politician to exert a higher level of effort.

Our paper provides a tractable model to quantify these countervailing forces. We identify the environments in which the party prefers the friendly authority and those in which it prefers the unfriendly one. In particular we show that the party prefers

the unfriendly to the friendly authority whenever the following two conditions hold. First, the approval threshold is relatively low. Second, the marginal cost of effort is relatively flat. A low approval threshold prevents the unfriendly authority from issuing false negatives too often, which incentivizes the politician to exert effort. Both the less severe reporting strategy and the higher level of effort benefit the party. A low approval threshold also allows the friendly authority to issue a higher share of false positives, which disincentivizes the politician from exerting effort. While the more favorable reporting strategy benefits the party, the lower level of effort hurts him. When the marginal cost of effort is flat, the politician's level of effort is highly responsive to changes in the reporting strategy: a small decrease in the approval threshold leads to a large decrease in the level of effort under the friendly authority and to a large increase under the unfriendly authority. Together, the two conditions above thus guarantee that the favorable effect of the friendly authority's reporting strategy on the party's payoff does not prevail. The party's payoff under the unfriendly authority is then higher than under the friendly one.

The oversight of a friendly authority can thus be detrimental for the approval of reforms when the adoption of reforms of high quality is needed the most. Indeed, the adoption of high-quality reforms is most valuable when the status quo's payoff is low; that is, when the approval threshold is low. In these situations, an unfriendly authority can increase the probability of approval. The economic forces that we highlight thus provide a novel rationale for requiring a tough oversight for reforms, possibly from an authority controlled by the opposition.

Finally, we also study how a truthful authority affects the politician's effort and the party's payoff. The main message extends to this case: compared to the truthful authority, the friendly authority is also detrimental to the party.

1.1 Related literature

In our paper, the approval probability of a reform depends on the politician's costly and unobservable effort. Our work is thus related to the literature on moral hazard pioneered by Hölmstrom (1979). Within this extensive literature, our paper is closest to models in which incentive provision occurs through retention/approval binary choices, rather than through compensation contracts. Retention rules are particularly relevant in political settings (see the surveys of Ashworth 2012 and Duggan and Martinelli 2017). We depart from this literature by introducing an informed third party (the authority) who issues a report about the quality of the reform. This allows us to study how the authority's objectives affect incentive provision and outcomes.²

There is by now a large literature that studies Bayesian persuasion (Kamenica and Gentzkow, 2011; see Kamenica, 2019 for a review on this topic). In our paper, the authority knows the quality of the reform (but not the level of effort exerted by the politician) and chooses a reporting strategy that is observable to the constituency. Hence, we study information design in the presence of moral hazard.³ A growing line of research within the Bayesian persuasion framework focuses on the interaction between persuasion and moral hazard: Boleslavsky and Cotton (2015), Boleslavsky and Kim (2021), Zapechelnyuk (2020), Rodina (2020), and Rodina and Farragut (2020).⁴ Boleslavsky and Cotton (2015) show that inflated grading policies, despite decreasing the informativeness of grades, may increase the quality of graduating students if schools compete against each other.⁵ Rodina (2020), Rodina and Farragut (2020) and Zapechelnyuk (2020) study the optimal information design to incentivize effort.

Boleslavsky and Kim (2021) study persuasion under moral hazard, as we do in this paper. A key distinction between our work and Boleslavsky and Kim's lies in the timing of the game. In our paper, the authority and the politician move simultaneously. Instead, in Boleslavsky and Kim (2021), a sender (the authority in our setup) moves first by choosing a reporting strategy which is publicly observed both by the agent (politician) and by the receiver (constituency). Boleslavsky and Kim characterize the sender's optimal reporting mechanism. In our paper we consider three types of authorities: those in favor of a reform, those against it, and truthful ones. Both ideologically biased and independent authorities are common in political environ-

²Georgiadis and Szentes (2020) study optimal monitor design when the decision maker can pay a cost to directly acquire information.

³Alonso and Câmara (2016), Bardhi and Guo (2018), Chan et al. (2019) and Heese and Lauermann (2021) study persuasion in political settings. In these papers, a politician persuades an electorate made by multiple receivers into approving a policy. In our model, there is only one receiver, the quality of the policy under consideration is determined endogenously, and the sender is a third party with its own preferences. Our paper also differs from Dragu et al. (2014), who characterize optimal checks and balances from a mechanism design perspective in a setting without uncertainty.

⁴Feng and Lu (2016) and Zhang and Zhou (2016) study the interaction between effort and persuasion in contests. Bizzotto et al. (2021) study information design when the receiver is subject to moral hazard.

⁵Because of the complementarity between students' effort and schools' investment, the result holds true even when students react to a more lenient grading policy.

ments. Our timing allows us to characterize the different mechanisms through which the ideology of the authority affects the approval of reforms.⁶

Our paper is also related to the literature on certification agencies (Lizzeri, 1999 Albano and Lizzeri, 2001, Miklós-Thal and Schumacher, 2013, and Bizzotto and Harstad, 2020). This literature focuses on market environments where sellers choose the quality of their products and agencies certify the quality in exchange of a fee. Certifying agencies improve the quality of products, but such quality never reaches the full information benchmark. In contrast, the authority in our model is motivated by ideology and not by monetary incentives.

There is a broad literature in political economy that investigates the role of informed experts and how they interact with elected officials (see Gailmard and Patty, 2012, for a review). Maskin and Tirole (2004), Alesina and Tabellini (2007, 2008), and laryczower et al. (2013) study the trade-offs associated with the allocation of political tasks to bureaucrats—who are equipped with superior technical expertise and who are not elected—rather than to politicians. Banks (1989), Huber et al. (2001), and Bueno de Mesquita and Stephenson (2007) investigate how politicians can control and incentivize authorities to behave in the politicians' interests. Differently from both these literatures, we assume that the constituency relies on the report issued by an authority and we study how the reporting strategy affects the likelihood of approval of a reform. In Prato and Turner (2022), a politician can cover its malfeasances by distorting the information available to a bureaucrat. The bureaucrat can detect the cover-ups thanks to an overseer. In their model, the politician (rather than the overseer) designs the information structure and the overseer has no ideological bias. Similarly to our model, in Gulotty and Luo (2021) a constituency relies on an authority to monitor a politician. In contrast to our model, the authority's reporting strategy is fixed.

In Taylor and Yildirim (2011), like in our paper, the behavior of a reviewer affects the effort exerted by an agent. Taylor and Yildirim compare blind evaluations—i.e., evaluations in which the identity of the agent is kept secret—against informed evaluations. We instead focus on how the preferences of the reviewer affect outcomes.

Finally, our paper shows that biased experts may lead to inefficient approval of reforms. Our paper is thus linked to the literature on political gridlock (see Krehbiel,

⁶If, as in Boleslavsky and Kim (2021), the politician moves after observing the authority's reporting strategy, then the friendly authority yields the highest approval probability. Intuitively, a friendly authority could always mimic the behavior of any other authority.

1998, Ortner, 2017, Austen-Smith et al., 2019, and the references therein). Our work highlights the role of information transmission as a source of political stalemate.

2 The model

A politician ("she") puts forward a policy reform. The reform can be either of high quality ($\omega = 1$) or of low quality ($\omega = 0$). The reform is of high quality with probability $e \in [0, 1]$ and of low quality with probability 1 - e, where e denotes the politician's costly effort. The politician's effort level is not observable.

The constituency of the politician must decide whether to approve the reform (a = 1) or to reject it (a = 0). The constituency does not observe the quality of the reform. Instead, the constituency must rely on a report from an oversight authority ("it"), which observes the quality ω of the reform. The authority may favor or oppose the reform and can bias its report to persuade the constituency into choosing the authority's preferred action. The preferences of the authority are known to all the players.

The authority commits to a reporting strategy *before* learning the quality ω . The reporting strategy is a mapping from the quality of the reform to reports. As the constituency has two actions (and the authority has commitment power), we assume without loss of generality that there are two available reports, $m \in \{0, 1\}$. Thus, the authority reporting strategy is summarized by a pair $\mu = (\mu_0, \mu_1)$, where μ_{ω} denotes the probability that the report is m = 1 when the quality is ω . In other words, the authority selects a procedure (test) with a share μ_0 of false positives and a share $1 - \mu_1$ of false negatives. We assume that $\mu_1 \ge \mu_0$, so m = 1 constitutes (weak) evidence that the reform is of high quality.⁷

The timing of the game is simple. First, the politician chooses the level of effort. Second, without observing the level of effort, the authority chooses the reporting strategy. Third, the quality of the reform is determined and the authority issues its report. Fourth, the constituency observes both the reporting strategy of the authority

⁷Report m = 1 can thus be interpreted as a recommendation to approve the proposal and report m = 0 as the recommendation to reject it. For every reporting strategy (μ_0, μ_1) with $\mu_1 > \mu_0$ there is a strategically equivalent reporting strategy $(1 - \mu_0, 1 - \mu_1)$ in which the report m = 0 provides instead evidence that the reform is of high quality. We focus on reporting strategies with $\mu_1 \ge \mu_0$ to overcome this trivial source of multiplicity.

and its actual report. Finally, the constituency decides whether to approve or to reject the reform and payoffs are realized.

The politician obtains a payoff of 1 if the constituency approves the reform and a payoff of 0 otherwise. She also pays a cost of effort captured by the function $e \mapsto c(e) \in \mathbb{R}$. Thus, she obtains utility u(a, e) = a - c(e).

ASSUMPTION The cost function c is continuous, strictly increasing and strictly convex. Furthermore, c(0) = 0, c'(0) < 1, c'(1) > 1, and $c'''(e) \ge 0$ for all $e \in [0, 1]$.

The assumption c'(0) < 1 implies that it is possible to incentivize the politician to exert some effort. Next, the assumption c'(1) > 1 implies that it is never optimal for the politician to exert maximal effort and guarantee a reform of high quality.⁸ The assumption on the third derivative of the cost function simplifies the analysis. However, our equilibrium characterization and the ranking over authorities extend beyond it (see Section 3.3 and footnote 17 for details). The remaining assumptions on the cost function are standard.

The constituency obtains a payoff of 1 if it approves a high-quality reform and a payoff of 0 if it approves a low quality reform. If the constituency instead rejects the reform, it obtains a payoff of $q \in (0, 1)$ from a status quo policy.

The authority is of one of three possible types: a friendly authority, which gets a payoff of 1 if the reform is approved and 0 otherwise, an unfriendly authority, which gets a payoff of 1 if the reform is rejected and 0 otherwise, and a non-strategic truthful authority that always reveals the true quality of the reform.⁹

An inactive player in the game, a party ("he"), wants the reform to be approved, but does not pay any effort cost. Hence, the party obtains a payoff of 1 if the reform is approved and a payoff of 0 otherwise.

2.1 Equilibrium concept

We use Perfect Bayesian Equilibrium (PBE) as our solution concept. After observing the report from the authority, the constituency forms a belief about the quality of the

⁸The main findings of the paper extend to the case with $c'(1) \leq 1$.

⁹We do not model the truthful authority as a player as we find it convenient to take its perfectly informative reporting strategy as given. Such truthful reporting strategy can be rationalized by assuming that the authority shares the same preferences as the constituency.

reform.¹⁰ We let $\pi(m)$ denote the probability that the reform is of high quality when the authority issues a report $m \in \{0, 1\}$.

The constituency approves the reform if $\pi(m) > q$, rejects it if $\pi(m) < q$ and is indifferent if $\pi(m) = q$. We refer to q as the *approval threshold*. In equilibrium an indifferent constituency approves the reform if the authority is friendly and rejects it if the authority is unfriendly.¹¹

Given the simplicity of the constituency's behavior, we focus hereafter on the behavior of the politician and of the authority. When the authority chooses its reporting strategy, it does not know the politician's level of effort. Thus, we model the interaction between the politician and the authority as a static game and summarize their behavior with a pure strategy profile (e, μ) . As the cost function is strictly convex and the authority commits to a reporting strategy, all equilibria are in pure strategies.

3 Equilibrium analysis

3.1 Preliminary results

We start the analysis by highlighting that there are always equilibria where the politician exerts zero effort and thus the reform is of low quality with certainty. This is true regardless of whether the authority is friendly or unfriendly. In these equilibria, there is no room for the authority to persuade the constituency.

REMARK 1. EQUILIBRIUM WITH ZERO EFFORT. Both under the friendly and the unfriendly authority, there exist equilibria where the politician exerts zero effort and the reform is never approved.

To see why Remark 1 holds true, suppose that the politician exerts zero effort. For every reporting strategy and for every on-path message, the constituency rejects the reform. As a result, the authority is indifferent among all reporting strategies (μ_0 , μ_1). If the authority chooses a reporting strategy with either $\mu_0 > 0$, or with $\mu_0 = \mu_1 = 0$, then exerting zero effort is indeed optimal for the politician.¹²

¹⁰In a PBE, off-path beliefs are computed according to Bayes rule when the equilibrium level of effort is interior: $e \in (0, 1)$. Footnote 12 describes the restrictions that the solution concept imposes on beliefs when the effort level takes an extreme value.

¹¹This behavior of the constituency guarantees that the best response of the authority is non-empty.

¹²If e = 0, the observed reporting strategy has $\mu_0 = 0$ and $\mu_1 > 0$, and the report is m = 1, then the constituency assigns probability one to the reform being of high quality. For all other cases with e = 0, off-path beliefs can be computed according to Bayes rule.

Henceforth, we turn our attention to the interesting case where persuasion plays a role. We study equilibria with interior effort levels: $e \in (0, 1)$, which do not fully determine the quality of the reform.¹³ We refer to these equilibria as *interior equilibria*.

We next characterize the useful benchmark of the equilibrium under the truthful authority. In this case, the authority reports the quality of the reform truthfully and the constituency approves the reform if and only if the authority reports m = 1. Thus, the approval probability of the reform coincides with the politician's effort level.

REMARK 2. BENCHMARK: TRUTHFUL AUTHORITY. Assume that a non-strategic authority truthfully reveals the quality of the reform: $\mu = (0, 1)$. Then, the politician solves $\max_{e \in [0,1]} e - c(e)$ by choosing $e^* = (c')^{-1}(1)$.

The politician's level of effort e^* under a truthful authority (Remark 2) coincides with the maximum level of effort achievable by varying the authority's reporting strategy (as in Boleslavsky and Kim, 2021, see their Proposition 6).

3.2 Friendly authority

The friendly authority wants the approval of the reform. Thus, it adopts a procedure that generates false positives: the constituency receives report m = 1 not only when the reform is of high quality, but also, with some probability, when the reform is of low quality. The politician takes advantage of this friendly oversight and exerts a level of effort lower than the approval threshold q.¹⁴ Proposition 1 characterizes the interior equilibrium under the friendly authority. We report its proof, as well as all other proofs, in Appendix A.

PROPOSITION 1. FRIENDLY AUTHORITY. Under the friendly authority, there exists a unique interior equilibrium (e^F, μ^F) with

$$c'(e^F) = 1 - \mu_0^F \quad and \tag{1}$$

$$\mu^{F} = (\mu_{0}^{F}, \mu_{1}^{F}) = \left(\frac{1-q}{q}\frac{e^{F}}{1-e^{F}}, 1\right).$$
(2)

The constituency approves the reform only when it observes report m = 1. The politician then maximizes $e + (1 - e)\mu_0 - c(e)$. Equation (1) reflects the politician's

¹³The assumption c'(1) > 1 guarantees that in equilibrium the politician never chooses e = 1.

¹⁴If the equilibrium level of effort was greater or equal than q, the constituency would approve the reform regardless of the report. The politician would then choose zero effort. See the proof of Proposition 1 for details.

first order condition. Relative to the truthful authority, the friendly authority decreases the politician's marginal benefit of effort by μ_0 and thus the politician chooses a level of effort lower than e^* . Equation (2) describes optimal persuasion. The authority chooses a reporting strategy that makes the constituency indifferent between approving or rejecting the reform after receiving report m = 1, that is, it sets $\pi(1) = q$. In this way, the authority sends report m = 1 as often as possible, while still making it credible enough to induce the approval of the reform.

Figure 1a depicts the politician's best response (red dashed curve) and the authority's best response (blue thick curve). An interior equilibrium exists and it is unique. The intuition behind equilibrium uniqueness is simple: a *higher* level of effort allows the authority to send report m = 1 more often when the reform is of low quality. This in turn reduces the marginal benefit of effort, which would then lead to a *lower* level of effort. As Figure 1a illustrates, while the authority's best response is increasing, the politician's best response is decreasing.



Figure 1: The friendly authority

Notes (a): The red dashed curve depicts the politician's best response. It shows the politician's choice of effort level *e* for each reporting strategy (μ_0 , 1). The blue thick curve represents the authority's best response. It shows its choice of μ_0 for each level of effort *e* of the politician. The cost of effort is $c(e) = \frac{2}{3}e^2$ and q = 0.5.

Notes (b): The red dashed curve depicts the party's payoff when $c(e) = e^2$, the blue thick curve when $c(e) = (2e + e^2)/3$ and the black dotted line when $c(e) = -\log(1 - e)/2$.

The equilibrium level of effort increases with the approval threshold q. A higher threshold q shifts the authority's best response down. The friendly authority must limit its misreporting and so μ_0^F decreases. The drop in μ_0^F increases the politician's marginal benefit of effort and thus leads to an increase in e^F . Finally, as a higher approval threshold q reduces μ_0^F , the politician's equilibrium payoff also decreases.

Interestingly, while an increase in the approval threshold always hurts the politician, it may benefit the party. To see how, consider the effect of a higher approval threshold q on the party's payoff. A higher approval threshold decreases the distortion of the friendly authority, which has two countervailing effects on the approval probability. First, it decreases the likelihood of approval of low-quality reforms. Second, it forces the politician to work harder: she exerts a higher level of effort. The degree of convexity of the cost function determines how much the politician increases her level of effort. When the degree of convexity is not too high (see the definition of mildly convex cost functions below), the overall effect of a higher q on the approval probability is positive: the higher level of effort more than compensates the less lenient oversight. Thus, the party is better off with a higher approval threshold q.

To understand the effect of the convexity of the cost function on the party's payoff, recall that the party obtains a payoff equal to: $e^F + (1 - e^F)\mu_0^F$ in equilibrium. Then, the effect on the party's payoff of a marginal increase in q is equal to

$$\left[(1 - \mu_0^F) \frac{de}{d\mu_0} + (1 - e^F) \right] \frac{d\mu_0^F}{dq} = \left[c'(e^F) \frac{de}{d\mu_0} + (1 - e^F) \right] \frac{d\mu_0^F}{dq},$$
(3)

where the equality follows from equation (1) and $de/d\mu_0$ is the slope of the politician's best response. As μ_0^F decreases with q, the party's payoff increases if and only if the square bracket on the right-hand side of equation (3) is negative. Again by equation (1), $de/d\mu_0$ depends on the convexity of the cost function: $de/d\mu_0 = -1/c''(e^F)$. Lemma 1 presents the comparative statics under the friendly authority. In Section 4, we take advantage of this lemma to compare outcomes under different types of authorities.

LEMMA 1. COMPARATIVE STATICS WITH FRIENDLY AUTHORITY. In the interior equilibrium under the friendly authority, the level of effort e^F strictly increases with q and the politician's payoff strictly decreases with q. Moreover, at any given q, the party's payoff strictly increases with q if $c'(e^F) - (1 - e^F)c''(e^F) > 0$, strictly decreases with q if $c'(e^F) - (1 - e^F)c''(e^F) > 0$ and is constant otherwise.

Lemma 1 shows that the local degree of convexity of the cost function determines how the approval threshold q affects the party's payoff. The cost function $c(e) = -\alpha \log(1-e)$ with $0 < \alpha < 1$ acts as a knife-edge case: the condition $c'(e^F) - (1 - e^F)c''(e^F) = 0$ always holds true. When $c(e) = -\alpha \log(1-e)$, the effect of a higher e^F exactly compensates that of a lower μ_0^F . As a result, the party's payoff is constant for all $q \in (0, 1)$.

We then define two classes of cost functions for which the effect of q on the party's payoff is monotonic. First, a cost function is *highly convex* if it is convex relative to $-\log(1-e)$. Second, a cost function is *mildly convex* if $-\log(1-e)$ is convex relative to it.¹⁵

DEFINITION. HIGHLY CONVEX AND MILDLY CONVEX COST FUNCTIONS. We say that a cost function c is highly convex if there exists a strictly convex and increasing function h such that $c(e) = h(-\log(1-e))$ for all $e \in [0, (c')^{-1}(1)]$. Similarly, we say that a cost function c is mildly convex if there exists a strictly convex and increasing function h such that $-\log(1-e) = h(c(e))$ for all $e \in [0, (c')^{-1}(1)]$.

A cost function c is mildly (highly) convex if c''(e)/c'(e) < (>) 1/(1 - e) for all $e \in [0, (c')^{-1}(1)]$ (see footnote 15). When the cost function is mildly convex, the local condition $c'(e^F) - (1 - e^F)c''(e^F) > 0$ always holds true. Lemma 1 then implies that the party's payoff strictly increases with q. The blue thick curve in Figure 1b depicts the party's payoff when $c(e) = (2e + e^2)/3$, which is mildly convex. When instead the cost function is highly convex, the local condition $c'(e^F) - (1 - e^F)c''(e^F) < 0$ always hold true. Then, by Lemma 1, the party's payoff strictly decreases with q. The red dashed curve in Figure 1b depicts the party's payoff when $c(e) = e^2$, which is highly convex. Finally, the black dotted line in Figure 1b corresponds to the knife-edge case of $c(e) = -\log(1 - e)/2$. In this case, the party's payoff is constant in q and it is equal to e^* .

¹⁵A function f is convex relative to a function g if f = h(g) for some increasing and strictly convex function h. For twice differentiable (strictly increasing and convex) functions the following definition is equivalent: f is convex relative to g in [a, b] whenever f''(x)/f'(x) > g''(x)/g'(x) for all $x \in [a, b]$, that is if the (absolute) value of the Arrow-Pratt coefficient of f is larger than the one of g at all points (see Mas-Colell et al., 1995, Proposition 6.C.2, for the proof of the equivalence between these two definitions).

3.3 Unfriendly authority

The unfriendly authority wants the rejection of the reform. Thus, it adopts a procedure that generates false negatives: the constituency receives report m = 0 not only when the reform is of low quality, but also, with some probability, when the reform is of high quality. In response, the politician works hard: her level of effort in an interior equilibrium exceeds the approval threshold q.¹⁶ Proposition 2 characterizes the set of interior equilibria under the unfriendly authority.

PROPOSITION 2. UNFRIENDLY AUTHORITY. Under the unfriendly authority, there exists $q^{\dagger} \in (0, e^*]$ such that interior equilibria exist if and only if $q \leq q^{\dagger}$. When $q < q^{\dagger}$, there are two interior equilibria (e^{ℓ}, μ^{ℓ}) and (e^h, μ^h) with $e^{\ell} < e^h$ and $\mu_1^{\ell} < \mu_1^h$. Moreover, for $k \in \{\ell, h\}$,

$$c'(e^k) = \mu_1^k \quad and \tag{4}$$

$$\mu^{k} = \left(\mu_{0}^{k}, \mu_{1}^{k}\right) = \left(0, \frac{e^{k} - q}{e^{k}(1 - q)}\right).$$
(5)

When $q = q^{\dagger}$ the two equilibria collapse into one.

As under the friendly authority, the constituency only approves the reform when it observes report m = 1. However, the unfriendly authority (unlike the friendly one) only sends report m = 1 when the quality is high (and with probability $\mu_1 < 1$). Thus, the politician maximizes $e\mu_1 - c(e)$. Equation (4) determines the politician's best response to the authority's reporting strategy. Since $\mu_1 < 1$, the politician exerts a level of effort lower than e^* . Equation (5) determines the authority's best response to the politician's level of effort. The authority chooses a reporting strategy that makes the constituency indifferent between accepting or rejecting the reform after m = 0, that is, it sets $\pi(0) = q$.

Figure 2 depicts best response functions for the authority (in blue) and for the politician (in red). Figure 2a illustrates the case of two interior equilibria: a *low-effort equilibrium* (e^{ℓ}, μ^{ℓ}) and a *high-effort equilibrium* (e^{h}, μ^{h}) . This case occurs for lower values of the approval threshold ($q < q^{\dagger}$). Instead, Figure 2b illustrates the case of no interior equilibria, which occurs when the approval threshold is high ($q > q^{\dagger}$).

¹⁶If the equilibrium level of effort was positive and lower than q, the constituency would reject the reform regardless of the report. The politician would then choose zero effort. See the proof of Proposition 2 for details.

The potential multiplicity of equilibria arises because the response of the unfriendly authority to a higher level of effort amplifies the politician's incentives to work. When the politician exerts a higher level of effort, the unfriendly authority issues a lower share of false negatives (μ_1 increases). A higher μ_1 , in turn, increases the politician's marginal benefit of effort, and so it makes her work more (*e* increases). As Figure 2 illustrates, the best response functions of both the authority and the politician are increasing. This is unlike the case of the friendly authority, illustrated in Figure 1a. Finally, the assumption c'''(e) > 0 for all $e \in [0, 1]$ guarantees that the politician's best response is convex and thus there can be at most two interior equilibria.¹⁷ Even without this assumption on the third derivative, for low enough values of q, there would be at least two interior equilibria and all interior equilibria would satisfy equations (4) and (5). Moreover, the ranking over authorities discussed in Section 4 would still hold if we consider the highest effort equilibria under the unfriendly authority.





Figure 2: Multiple equilibria under the unfriendly authority

Notes: The red dash curve depicts the politician's best response. Blue thick curves represent best response functions for the authority. The cost of effort is $c(e) = \frac{2}{3}e^2$. The left panel corresponds to q = 0.15, while the right panel corresponds to q = 0.5.

¹⁷The authority's best response is strictly concave (whenever it chooses $\mu_1 > 0$). Then, the authority's and the politician's best responses cannot cross more than twice.

While e^{ℓ} increases with the approval threshold, e^{h} decreases with it. To see why, note that a higher q lowers the authority's best response (the blue line in Figure 2 shifts down), while keeping the politician's best response unchanged.

The authority's distortion is more pronounced in the low-effort equilibrium than in the high-effort one: $\mu_1^{\ell} < \mu_1^{h}$. Moreover, μ_1^{ℓ} increases with the approval threshold and μ_1^{h} decreases with it. Hence, the effects of changes in q on the politician's payoff and on the party's payoff (respectively, $e^k \mu_1^k - c(e^k)$ and $e^k \mu_1^k$) are straightforward. We summarize these results in Lemma 2.

LEMMA 2. COMPARATIVE STATICS WITH UNFRIENDLY AUTHORITY. Under the unfriendly authority, the effort level, the politician's payoff, and the party's payoff all increase with q in the low-effort equilibrium and decrease with q in the high-effort equilibrium.

We conclude the section characterizing the players' ranking over the low-effort and the high-effort equilibrium under the unfriendly authority.

LEMMA 3. RANKING OF EQUILIBRIA WITH UNFRIENDLY AUTHORITY. The politician, the party and the constituency all prefer the high-effort equilibrium to the low-effort equilibrium. The authority, instead, exhibits the reversed ranking.

Lemma 3 follows from the fact that the unfriendly authority issues a lower share of false negatives in the high-effort equilibrium than in the low-effort one. As for Lemma 1, we also take advantage of the results in Lemmas 2 and 3 when comparing the outcomes under different types of authorities in Section 4.

4 Ranking over authorities

In this section we present the agents' ranking over different types of authorities. We explain how a friendly authority can be detrimental to the party.

4.1 The politician and the party

The constituency approves the reform with positive probability whenever the equilibrium level of effort is positive. As a result, both the politician and the party strictly prefer any interior equilibrium to any equilibrium with zero effort. In what follows we thus focus on interior equilibria only.

The type of authority affects the approval probability of the reform through two channels. First, authorities distort their reports. The friendly authority issues some false positives: it sends report m = 1 not only when the reform is of high quality, but also with probability $\mu_0^F \in (0, 1)$ when the reform is of low quality. On the contrary, the unfriendly authority issues some false negatives: it sends report m = 0 not only when the reform is of low quality, but also with probability $1-\mu_1^k \in (0, 1)$ when the reform is of high quality. Holding the effort of the politician constant, these distortions imply that the approval probability is higher under the friendly authority. Second, because of the authorities' distortion, the politician works harder under the unfriendly authority than under the friendly one: $e^F < q < e^\ell < e^h$. Holding the authority's reporting strategy constant, this channel implies that the approval probability is higher under the unfriendly authority.

A revealed-preference argument shows that the politician prefers the friendly authority to the unfriendly one. The politician's payoff is $e^F + (1 - e^F)\mu_0^F - c(e^F)$ under the friendly authority and $e^k\mu_1^k - c(e^k)$ under the unfriendly one. Then for every $k \in \{\ell, h\}$:

$$e^{F} + (1 - e^{F})\mu_{0}^{F} - c(e^{F}) > e^{k} + (1 - e^{k})\mu_{0}^{F} - c(e^{k}) > e^{k}\mu_{1}^{k} - c(e^{k}).$$

The first inequality follows from the optimality of the politician's behavior and the second one from $\mu_1^k \in (0, 1)$. A similar argument shows that the politician ranks the truthful authority strictly between the friendly authority and the unfriendly one.

The party's ranking over authorities does not necessarily coincide with that of the politician. Although both the party and the politician care about the approval probability of the reform, the party, unlike the politician, does not factor in the cost of effort.

Figure 3 illustrates that the party may prefer the unfriendly authority over the friendly one. The cost function is $c(e) = \frac{1}{3} \left[\frac{3}{4}e + \frac{1}{4} \left(-\log(1-e) \right) \right]$, which is mildly convex.¹⁸ Thus, by Lemma 1, the party's payoff under the friendly authority (depicted with the blue thick curve) increases with q. Instead, by Lemma 2, the party's payoff under the unfriendly authority in the high-effort equilibrium (depicted with the red dashed curve) decreases with q. For low levels of the approval threshold, the party is strictly better off under the unfriendly authority than under the friendly authority. This corresponds to the range in Figure 3 where the red dashed curve lies above the blue thick curve.

¹⁸The cost function c is a weighted sum of $-\log(1-e)$ and a liner function. Hence, c is mildly convex, since $-\log(1-e)$ is convex relative to c.



Figure 3: Party's payoff under the friendly and unfriendly authority

Notes: The blue thick curve depicts the party's payoff under the friendly authority. The red dashed curve depicts the party's payoff in the high-effort equilibrium under the unfriendly authority. The cost function is $c(e) = \frac{1}{3} \left[\frac{3}{4}e + \frac{1}{4} \left(-\log(1-e) \right) \right]$, so $e^* = \frac{8}{9}$.

To see how this ranking over authorities may emerge, note that the friendly authority does not have much room for persuasion when the approval threshold q is close to one. The friendly authority must behave almost like the truthful one, so the politician's level of effort is close to e^* . The party thus obtains a payoff close to e^* . As the approval threshold *decreases*, the party's payoff under the friendly authority goes down: the negative effect from a lower level of effort dominates the positive effect from the friendly oversight. Note next that when the approval threshold q is close to zero the unfriendly authority behaves almost like the truthful one. This gives the party a payoff close to e^* in the high-effort equilibrium.

The party finds the friendly authority detrimental when the cost function is mildly convex and the approval threshold is low. A low approval threshold means that the unfriendly authority must keep the share of false negatives low, while the friendly authority can issue a large share of false positives. A mildly convex cost function, in turn, means that the level of effort is highly responsive to incentives: it decreases a lot when the oversight of the friendly authority becomes more lenient and it increases a lot when the oversight of the unfriendly authority becomes less severe. When the approval threshold is low (and the cost function is mildly convex), the party's payoff under the friendly authority is low because the (negative) effect on the level of effort offsets the (positive) effect of the favorable reporting strategy. At the same time, the party's payoff under the unfriendly authority is high thanks to the low share of false negatives and its associated high level of effort. The party's is thus better off with the unfriendly authority.

When the cost function is highly convex, the level of effort does not react much to incentives. Then, the party always prefers the friendly authority. Indeed, while the approval probability under the unfriendly authority is bounded above by e^* , this probability is bounded below by e^* under the friendly authority.

PROPOSITION 3. PARTY'S PREFERENCES OVER INTERIOR EQUILIBRIA. Whenever the cost function is mildly convex, there exists $\overline{q} \in [0, q^{\dagger}]$ such that the party prefers the high-effort equilibrium under the unfriendly authority to the interior equilibrium under the friendly one if and only if $q \leq \overline{q}$.

Whenever the cost function is highly convex the party always prefers the interior equilibrium under the friendly authority to the high-effort equilibrium under the unfriendly one.

The party may even prefer the low-effort equilibrium under the unfriendly authority to the interior equilibrium under the friendly one. We next discuss how the party ranks the low-effort equilibrium under different configurations of the primitives. First, the party unambiguously prefers the friendly authority whenever either the cost function is highly convex or $q > \overline{q}$.¹⁹ Second, when the cost function is mildly convex and q is close to zero, the party again prefers the friendly authority. The politician's level of effort converges to zero as q approaches zero both under the friendly authority and in the low-effort equilibrium under the unfriendly one.²⁰ Nonetheless, the politician benefits from the oversight of the friendly authority, so the party prefers the friendly authority to the low-effort equilibrium under the unfriendly one.²¹ Finally, for interior values of the approval threshold q, the party may rank the equilibrium under the friendly authority.

Figure 4 illustrates the party's ranking between the friendly authority and the low effort equilibrium under the unfriendly one with two examples. Figure 4a shows that with the cost function from Figure 3, the party always prefers the equilibrium

¹⁹Proposition 3 implies that the party prefers (e^F, μ^F) to (e^h, μ^h) and Lemma 3 implies that the party prefers (e^h, μ^h) to (e^ℓ, μ^ℓ) .

²⁰We show this for the low-effort equilibrium under the unfriendly authority in the proof of Proposition 3. Under the friendly authority, this limit follows from rewriting equation (A-1) as $1 - c'(e^F(q)) \left[1 - e^F(q)\right] = e^F(q)/q$. The left-hand side is bounded, so $\lim_{q\to 0} e^F(q) = 0$.

²¹The party's equilibrium payoff under the friendly authority is equal to e^{F}/q and converges to 1 - c'(0) > 0 as q approaches 0. This follows from applying L'Hôpital Rule and exploiting (A-3). Instead, the party's payoff in the low-effort equilibrium converges to 0 as q approaches 0.

under the friendly authority to the low-effort equilibrium under the unfriendly one. Figure 4b shows that with a similar cost function the order can be reversed: for an interior range of approval thresholds, the party prefers the low-effort equilibrium under the unfriendly authority to the equilibrium under the friendly one.



Figure 4: Party may prefer low-effort equilibrium under unfriendly authority to equilibrium under friendly authority.

Notes: The blue thick curve depicts the party's payoff under the friendly authority. The red dashed curve depicts the party's payoff in the high-effort equilibrium under the unfriendly authority. The black dotted curve depicts the party's payoff in the low-effort equilibrium under the unfriendly authority.

The degree of convexity of the cost function also guides the comparison between the friendly authority and the truthful one. When the cost function is mildly convex, the party prefers the truthful authority to the friendly one. The opposite is true when the cost function is highly convex. To see why, note first that the party's payoff under the friendly authority is equal to e^* when q = 1 and it is increasing (decreasing) in q if the cost function is mildly (highly) convex. Second, the party's payoff under the truthful authority is exactly e^* for all values of the approval threshold q (see Remark 2).

The level of effort e^* is also an upper bound for the party's payoff under the unfriendly authority. Thus, the party always rank the truthful authority above the unfriendly one. Finally, there are cost functions that are neither mildly convex nor highly convex. In Appendix B we provide sufficient conditions for an unambiguous ranking over authorities from the point of view of the party.

4.2 The constituency

The constituency is indifferent between any equilibrium with zero effort and the interior equilibrium under the friendly authority. Whenever the politician exerts zero effort, the constituency rejects the reform and obtains a payoff equal to the approval threshold q. In the interior equilibrium under the friendly authority, the constituency rejects the reform after receiving report m = 0 and obtains a payoff of q. When it instead receives report m = 1, the constituency is indifferent between approving the reform or rejecting it, so it also receives a payoff of q. Thus, the constituency obtains a payoff of q in the interior equilibrium under the friendly authority.

The payoff of the constituency in an interior equilibrium is larger under the unfriendly authority than under the friendly one. Under the unfriendly authority, the constituency approves the reform after receiving report m = 1. When instead it receives report m = 0, the constituency is indifferent between approving the reform or rejecting it. Thus, the constituency's equilibrium payoff is equal e^k , the payoff it would get by always approving the reform. As $e^k > q$ for both $k \in \{\ell, h\}$, the constituency prefers the unfriendly authority to the friendly one.

Finally, the constituency ranks the truthful authority over the unfriendly one (and thus also over the friendly one). Indeed, when the authority is truthful, the politician exerts a level of effort equal to e^* and the constituency observes a perfectly informative signal. Thus, the constituency obtains a payoff of $e^* + (1 - e^*)q$. Since $e^* > e^k$ for both $k \in \{\ell, h\}$, the constituency is better off under the truthful authority than under the unfriendly one.

REMARK 3. CONSTITUENCY'S PREFERENCES OVER INTERIOR EQUILIBRIA. *The constituency prefers the truthful authority to the unfriendly one and the unfriendly one to the friendly.*

The constituency's ranking over authorities is intuitive. The constituency wants to approve reforms of high quality. As such, it would like the politician to exert high effort and the reports to be informative. The truthful authority achieves both these goals. On the contrary, the friendly authority disincentives effort and can also result in the approval of low quality reforms. The unfriendly authority ranks in-between: although some high-quality reforms are rejected, the reforms that are approved are always of high quality.

5 Discussion

5.1 Observable effort

In our baseline model, the politician's level of effort is unobservable. In this subsection we discuss instead an environment without moral hazard: we assume that the level of effort is observable. We modify the timing of the game in a simple way. We assume that after the politician chooses her level of effort, both the authority and the constituency observe such effort. Afterwards, the authority chooses its reporting strategy. The remainder of the timing is as before.

When the level of effort is observable, the authority's reporting strategy is a function of the level of effort. Thus, the equilibrium is characterized by a pair (e, μ) , where μ represents a reporting strategy function $e \in [0, q] \mapsto \mu(e) = (\mu_0(e), \mu_1(e))$.

The friendly authority's reporting strategy function is $(\mu_0(e), \mu_1(e)) = \left(\frac{1-q}{q} \frac{e}{1-e}, 1\right)$ for $e \in [0, q]$.²² The function μ_0 is the friendly authority's best response shown in Figure 1a. The politician then solves $\max_e e + (1-e) \mu_0(e) - c(e)$. The only difference with the case of unobservable effort is that now μ_0 depends on the politician's level of effort. When the effort is observable, the politician has an extra incentive to work hard: a higher effort allows the friendly authority to issue a larger share of false positives (μ_0 increases with e). This is why in equilibrium the politician exerts a higher level of effort compared to the case of unobservable effort.

The unfriendly authority's reporting strategy function is $(\mu_0(e), \mu_1(e)) = \left(0, \frac{e-q}{e(1-q)}\right)$ for $e \in [q, 1]$.²³ The function μ_1 is the unfriendly authority's best response shown in Figure 2. The politician either exerts zero effort or chooses the (unique) effort level that solves $\max_{e \in [q,1]} e\mu_1(e) - c(e)$. Again, the politician has an extra incentive to work hard: a higher effort limits the unfriendly authority's ability to report false negatives $(\mu_1 \text{ increases with } e)$. Hence, the politician exerts a higher level of effort compared to the case of unobservable effort.

²²The authority's reporting strategy is uninformative when e > q. Thus, the politician never chooses an effort level strictly larger than q.

²³Whenever e < q, the authority's reporting strategy is uninformative and the reform is always rejected.

Does the party prefer the friendly or the unfriendly authority when effort is observable? The party always prefers the friendly authority when q is lower than 1/2.²⁴ This is unlike the case of unobservable effort. When q is higher than 1/2 we can build examples in which the party prefers either type of authority. Moral hazard plays a key role in the results of our baseline model. Under the friendly authority, moral hazard is particularly harmful for the approval probability of the reform when q is low (and the cost function is mildly convex). In this case, the friendly authority issues a high share of false positives. This gives incentives to the politician to free ride on this lenient oversight and exert little effort. In contrast, the observability of effort forces the friendly authority to become more severe when the politician works less. This, in turn, disciplines the politician to exert a higher level of effort.

5.2 Uncertainty about the preferences of the constituency

In our baseline model, the payoffs of the constituency are common knowledge. In some settings, however, the constituency may have private information about its payoffs. For example, the approval threshold q may be unknown to the authority, to the politician and to the party. In this case, q would represent the constituency's type.

How does private information on the approval threshold q affect our results? Although a complete analysis of this case is beyond the scope of our paper, the driving forces behind our results extend to it. We can find mildly convex cost functions such that, if q is low with probability sufficiently high, the party prefers the unfriendly authority to the friendly one. The opposite ranking arises when q is high with probability sufficiently high.

The study of the case with private information on q sheds light on a key aspect of our model. When the payoffs of the constituency are common knowledge, the constituency is indifferent between following the authority's advice or ignoring it. Then, the authority's message brings no value to the constituency. The constituency benefits from the presence of the authority only insofar as the authority provides incentives for the politician to work.

²⁴To see this, note first that under the friendly authority, the politician's equilibrium level of effort is either equal to q or lower than it. If it is equal to q the reform is always approved, so the party prefers the friendly authority. Suppose next that the equilibrium level of effort is strictly lower than q. This, together with q < 1/2, implies that under the unfriendly authority, the politician exerts zero effort and the reform is always rejected. Thus, the party again prefers the friendly authority.

When the approval threshold q is private information, some types of constituency may strictly benefit from the report of the authority. To see why, consider the case with a friendly authority. After a report m = 0, the expected quality of the reform is equal to 0. Instead, after a report m = 1, the expected quality of the reform is equal to some value $\rho \in (0, 1)$.²⁵ Constituencies with a type lower than ρ strictly prefer to approve the reform after report m = 1 and to reject it after report m = 0. For these types, the authority's report is valuable. In contrast, constituencies with a type above ρ reject the reform independently of the report. For those types, the authority's report is not valuable. The same logic also applies to the case of the unfriendly authority: some types may strictly benefit from the authority's report.

6 Conclusion

In many economic and political environments, experts with the technical know-how to evaluate intricate proposals inform decision makers about the quality of such proposals. Experts, however, may have their own agendas and can use their expertise to steer decisions towards their own interests. The interests of experts may be aligned or misaligned with those of the interest groups behind the proposals.

We study how such alignment of interests affects the likelihood of approval of proposals. To do so, we focus on a simple political environment. A politician brings forward a reform proposal. The quality of the reform depends positively on the politician's unobservable and costly effort. A constituency does not observe the quality of the reform directly. Instead, an expert authority provides a report about such quality. The authority may want the reform to be approved (friendly authority) or to be rejected (unfriendly authority).

We identify the main forces at play in this simple environment of information design, moral hazard and biased authorities. We find that the friendly authority is detrimental for the approval of the reform when two conditions hold. First, the alternative to the reform is not too attractive for the constituency. Second, incentivizing the politician's effort is not too difficult.

²⁵The optimality of reporting strategies with only two reports extends to this setting with uncertainty about the constituency's preferences.

Our results highlight how a friendly authority can be detrimental exactly when high-quality reforms are needed the most, namely when they can represent a significant improvement over the status quo.

Appendix

A Proofs

Proof of Proposition 1

We first show the auxiliary result that $e^F < q$. Suppose instead that there exists an equilibrium where the politician exerts a level of effort $e \ge q$. The authority can then guarantee the approval of the reform by choosing an uninformative reporting strategy, namely by setting $\mu_0 = \mu_1$. In such a case, the constituency approves the reform with certainty after either report. This in turn pushes the politician to choose e = 0, which violates the premise that $e \ge q$ (as q > 0).

We next obtain the authority's best response to an interior level of effort $e \in (0, q)$. Given the behavior of the constituency, standard results from the literature on persuasion imply that the authority's best response must be as follows. First, the authority reports m = 1 whenever the reform is of high quality. Second, the authority reports m = 0 with probability $\mu_0 \in (0, 1)$ when the reform is of low quality, where μ_0 solves $\frac{e}{e+(1-e)\mu_0} = q$. The authority's best response is thus the reporting strategy $\mu = \left(\frac{e(1-q)}{(1-e)q}, 1\right)$, which is well defined since e < q. This reporting strategy induces posterior beliefs $\pi(0) = 0$ and $\pi(1) = q$. The constituency then approves the reform after receiving report m = 1 and rejects it after receiving report m = 0.

We then solve for the politician's best response to an arbitrary reporting strategy $\mu = (\mu_0, \mu_1)$ with $\mu_1 \ge \mu_0$. The politician solves $\max_{e \in [0,1)} e\mu_1 + (1-e)\mu_0 - c(e)$, which is a strictly concave problem (given the properties of *c*). The politician's best response thus satisfies $\mu_1 - \mu_0 \le c'(e)$ with strict equality whenever e > 0.

The combination of the two best responses leads to equations (1) and (2) in the main text, which characterize an interior equilibrium. Moreover, equations (1) and

(2) together imply that the following equality must hold in equilibrium:

$$1 - \frac{1 - q}{q} \frac{e}{1 - e} = c'(e).$$
 (A-1)

The left-hand side of the previous expression decreases with e, while the right-hand side increases with e. Hence, the equality holds at most once. Moreover, the left hand side is larger than the right hand side at e = 0, while the right hand side is larger than the left hand side at e = q. Thus, the equilibrium exists and it is unique.

Proof of Lemma 1

Equations (1) and (2) together imply that in the unique interior equilibrium the following equality must hold:

$$c'(e^F)(1-e^F) - 1 + \frac{e^F}{q} = 0$$
 (A-2)

Because of the implicit function theorem, this implies that

$$\frac{de^F}{dq} = \frac{1 - c'(e^F)(1 - e^F)}{1 - c'(e^F)q + c''(e^F)q(1 - e^F)} > 0,$$
(A-3)

where the inequality follows from $c'(e^F) < 1$. Hence, the equilibrium level of effort increases with q.

Note next that the equilibrium payoff of the politician is equal to

$$e^{F}\mu_{1}^{F} + (1 - e^{F})\mu_{0}^{F} - c(e^{F}) = e^{F} + (1 - e^{F})\frac{e^{F}(1 - q)}{(1 - e^{F})q} - c(e^{F}) = \frac{e^{F}}{q} - c(e^{F}).$$

The derivative of the politician's payoff with respect to q is thus equal to

$$\frac{de^F}{dq}\frac{1-c'(e^F)q}{q}-\frac{e^F}{q^2}<\frac{1-c'(e^F)(1-e^F)}{q}-\frac{e^F}{q^2}=0,$$

where the inequality follows from inequality (A-3) and the equality follows from equation (A-2). Hence, the politician's equilibrium payoff decreases with q.

Finally, equation (3) in the main text shows that the derivative of the party's payoff with respect to q is equal to

$$\left[c'(e^F)\frac{de^F}{d\mu_0}+(1-e^F)\right]\frac{d\mu_0^F}{dq}$$

Since $d\mu_0^F/dq < 0$, the party's payoff strictly increases with q if $c'(e^F)\frac{de^F}{d\mu_0} + (1 - e^F) < 0$, and strictly decreases with q if the opposite strict inequality holds. Applying the implicit function theorem on equation (1), we get $\frac{de^F}{d\mu_0} = -(c''(e^F))^{-1}$. Hence, the party's payoff is strictly increasing in q if and only if

$$-\frac{c'(e^F)}{c''(e^F)} + (1 - e^F) < 0$$

Rearranging terms, we get the inequality in the statement of the proposition. \Box

Proof of Proposition 2

We first show the auxiliary result that $e^k > q$. Assume instead that there exists an equilibrium where the politician exerts a level of effort $e \in (0, q]$. The authority can then guarantee the rejection of the reform by adopting an uninformative reporting strategy, that is, by setting $\mu_0 = \mu_1$. In such a case the constituency rejects the reform with certainty after either report. This in turn pushes the politician to choose e = 0, which contradicts $e \in (0, q]$.

We next obtain the authority's best response to a level of effort $e \in (q, 1)$. Given the behavior of the constituency, standard results in the persuasion literature imply that the authority's best response must be as follows. First, the authority reports m = 0 when the reform is of low quality. Second, the authority reports m = 1 with probability $\mu_1 \in (0, 1)$ when the reform is of high quality, where μ_1 solves $\frac{e(1-\mu_1)}{e(1-\mu_1)+1-e} = q$. The authority's best response is thus the reporting strategy $\mu = \left(0, \frac{e-q}{e(1-q)}\right)$, which is well defined because e > q. This reporting strategy induces posterior beliefs $\pi(0) = q$ and $\pi(1) = 1$. The constituency approves the reform after receiving report m = 1 and rejects it after receiving report m = 0.

The politician's best response is identical to that under the friendly authority and satisfies $\mu_1 - \mu_0 \le c'(e)$, with strict equality whenever e > 0.

The combination of the two best responses leads to equations (4) and (5) in the main text, which characterize an interior equilibrium. Equation (4) directly implies $e^k < e^*$.

The previous steps show that in any interior equilibrium (e^k, μ^k) , the level of effort must satisfy $q < e^k < e^*$. Thus, whenever $q \ge e^*$ no interior equilibrium exists. We then assume that $q < e^*$ hereafter.

Next, for any $q \in (0, e^*)$ we define the auxiliary function $e \in [q, e^*] \mapsto f_q(e)$ as follows

$$f_q(e) := \frac{e - q}{e(1 - q)} - c'(e).$$
(A-4)

The properties of the cost function imply that for every $q \in (0, e^*)$ f_q is continuous and strictly concave. Furthermore, both $f_q(q) = -c'(q) < 0$ and $f_q(e^*) = \frac{e^*-q}{e^*-qe^*} - 1 < 0$. Then, fixing q, either (i) $f_q(e) < 0$ for all $e \in (q, e^*)$, or (ii) $f_q(e) \ge 0$ for some $e \in (q, e^*)$. An interior equilibrium level of effort e^k must satisfy the necessary sufficient condition $f_q(e^k) = 0$. In case (i), there are no interior equilibria. In case (ii), the concavity of f_q implies that exactly two solutions exists if $f_q(e) > 0$ for some $e \in (q, e^*)$ and only one solution exists otherwise. An interior equilibrium exists if and only if $f^*(q) := \max_{e \in [q, e^*]} f_q(e) \ge 0$, where the maximum is well defined. Let $e^{\dagger}(q) = \arg \max_{e \in [q, e^*]} f_q(e)$. The maximizer is unique because f_q is strictly concave. The function $f^*(q)$ is continuous because of the maximum theorem and satisfies $\lim_{q \to 0} f^*(q) > 0$ and $\lim_{q \to e^*} f^*(q) = -1 < 0$. Finally, $f^*(q)$ is strictly decreasing: for any q' > q

$$f^{*}(q) = f_{q}\left(e^{\dagger}(q)\right) \ge f_{q}\left(e^{\dagger}(q')\right) > f_{q'}\left(e^{\dagger}(q')\right) = f^{*}(q'),$$

where the strict inequality holds since, for any given e, $f_{q'}(e) < f_q(e)$. Thus, there exists a unique $q^{\dagger} \in (0,1)$ such that $f^*(q^{\dagger}) = 0$. Three cases are possible. First, if $q > q^{\dagger}$, then there is no interior equilibrium. Second, if $q = q^{\dagger}$, then there exists a unique interior equilibrium with level of effort $e^{\dagger}(q^{\dagger})$. Third, if $q < q^{\dagger}$, then there exist two interior equilibria (e^{ℓ}, μ^{ℓ}) and (e^{h}, μ^{h}) with $q < e^{\ell} < e^{\dagger}(q) < e^{h} < e^{*}$.

Proof of Lemma 2

Proposition 2 shows that for any $q < q^{\dagger}$ there exist exactly two interior equilibria, (e^{ℓ}, μ^{ℓ}) and (e^{h}, μ^{h}) . The equilibrium levels of effort e^{ℓ} and e^{h} are the two roots of $f_q(e) = 0$, where the function f_q is as defined in (A-4) and $e^{\ell} < e^{h}$. So for any $q < q^{\dagger}$, we let $e^{\ell}(q)$ and $e^{h}(q)$ be such that $e^{\ell}(q) < e^{h}(q)$ and $f_q(e^{k}(q)) = 0$ for $k \in \{\ell, h\}$. First,

note that $f_q(e) < f_q(e^{\ell}(q)) = 0$ for all $e < e^{\ell}(q)$. Second, note that if q < q', then $f_{q'}(e) < f_q(e)$ for any e. Then, whenever $q < q' \leq q^{\dagger}$,

$$f_{q'}(e) < f_q(e) \le f_q(e^\ell(q)) = 0 \qquad \forall e \in \left(0, e^\ell(q)\right)$$

Then, $e^{\ell}(q') > e^{\ell}(q)$. A similar argument shows that $e^{h}(q') < e^{h}(q)$.

In the interior equilibrium (e^k, μ^k) , the politician's payoff is equal to

$$e^k \mu_1^k - c(e^k) = e^k c'(e^k) - c(e^k).$$

The derivative of the politician's payoff with respect to q is thus equal to

$$e^{k}c''(e^{k})\frac{de^{k}}{dq} + \frac{de^{k}}{dq}c'(e^{k}) - c'(e^{k})\frac{de^{k}}{dq} = e^{k}c''(e^{k})\frac{de^{k}}{dq}$$

Hence, the politician's payoff increases with q in the low-effort equilibrium and decreases with q in the high-effort equilibrium.

Finally, consider the equilibrium payoff of the party. Under the unfriendly authority, this is equal to $e^k \mu_1^k$. The derivative of this payoff with respect to q is equal to

$$e^{k}c''(e^{k})\frac{de^{k}}{dq} + \frac{de^{k}}{dq}c'(e^{k}) = \left[e^{k}c''(e^{k}) + c'(e^{k})\right]\frac{de^{k}}{dq}.$$

Hence, once more, the party's payoff increases with q in the low-effort equilibrium and decreases with q in the high-effort equilibrium.

Proof of Lemma 3

The party's payoff under the unfriendly authority in an interior equilibrium (e^k, μ^k) , with $k \in \{\ell, h\}$, is equal to $e^k \mu_1^k$. The ranking of equilibria for the party follows from observing that $e^h > e^\ell$ and $\mu_1^h > \mu_1^\ell$, so $0 < e^\ell \mu_1^\ell < e^h \mu_1^h < 1$. The ranking for the authority is the opposite to that of the party.

Under the unfriendly authority, the constituency only approves reforms of high quality. Then, the constituency gets a payoff equal to

$$e^k \mu_1^k + (1 - e^k \mu_1^k)q = (1 - q)e^k \mu_1^k + q = e^k - q + q = e^k.$$

The constituency thus prefers the high-effort equilibrium to the low-effort one.

Finally, the politician's payoff is equal to $e^k \mu_1^k - c(e^k)$. The following inequalities show that the politician prefers the high-effort equilibrium to the low-effort one

$$e^h \mu_1^h - c(e^h) > e^\ell \mu_1^h - c(e^\ell) > e^\ell \mu_1^\ell - c(e^\ell),$$

where the first strict inequality follows from the optimality of the politician, and the second one from $\mu_1^h > \mu_1^\ell$.

Proof of Proposition 3

Under the friendly authority, the equilibrium level of effort e^F increases with q and it is bounded above by e^* . Thus,

$$\lim_{q \to 1} \mu_0^F(q) = \lim_{q \to 1} \frac{1-q}{q} \frac{e^F(q)}{1-e^F(q)} = 0,$$

and so, by equation (1), $\lim_{q\to 1} e^F(q) = e^*$. We let $v^F(q)$ denote the party's payoff under the friendly authority:

$$v^F(q) = e^F \mu_1^F + (1 - e^F) \mu_0^F = e^F + \frac{(1 - q)e^F}{q} = \frac{e^F}{q}.$$

Then, $\lim_{q\to 1} v^F(q) = e^*$.

In an interior equilibrium (e^k, μ^k) with $k \in \{\ell, h\}$ under the unfriendly authority, the party obtains a payoff $v^k(q) = e^k \mu_1^k$. The condition $f_q(e^k) = 0$ is necessary and sufficient for an interior equilibrium under the unfriendly authority. We rewrite this condition as

$$c'(e^k(q))e^k(q)(1-q) = e^k(q) - q.$$

Then $\lim_{q\to 0} e^{\ell}(q) = 0$ and $\lim_{q\to 0} e^{h}(q) = e^*$. This, together with equation (4) implies that $\lim_{q\to 0} \mu_1^{\ell}(q) = c'(0)$ and $\lim_{q\to 0} \mu_1^{h}(q) = 1$. Hence, $\lim_{q\to 0} v^{\ell}(q) = 0$ and $\lim_{q\to 0} v^{h}(q) = e^*$.

When the cost function is highly convex, the party ranks the friendly authority over the unfriendly one. To see why, note that the party's payoff under the friendly authority, $v^F(q)$, strictly decreases with q when the cost function is highly convex (see Lemma 1 and the definition of highly convex cost functions in terms of the Arrow-Pratt coefficient). Since $\lim_{q\to 1} v^F(q) = e^*$, the party's payoff is bounded below by e^* . On the contrary, the party's payoff $v^h(q)$ under the unfriendly authority is bounded above by e^* . Indeed, $v^h(q)$ strictly decreases with q (see Lemma 2) and $\lim_{q\to 0} v^h(q) = e^*$.

When the cost function is mildly convex, the ranking over authorities depends on the value of the approval probability q. To see why, note that $v^F(q)$ strictly increases with q when the cost function is mildly convex (see Lemma 1 and the definition of mildly convex cost functions in terms of the Arrow-Pratt coefficient). Moreover, $\lim_{q\to 1} v^F(q) = e^*$, so $v^F(q) < e^*$ for all $q \in (0,1)$. Instead, $\lim_{q\to 0} v^h(q) = e^*$. Then, $v^h(q) > v^F(q)$ for q close to zero. Finally, under the unfriendly authority, the party's payoff strictly decreases with q (see Lemma 2). Thus, the party ranks the high-effort equilibrium under the unfriendly authority over the interior equilibrium under the friendly one if and only if q is below a threshold \bar{q} .

B Sufficient conditions to rank equilibria for arbitrary cost functions

Proposition 3 characterizes the party's preferred equilibrium and authority when the cost function is either mildly convex or highly convex. Some cost functions are neither mildly convex, nor highly convex. For arbitrary cost functions satisfying the assumptions in Section 2, Proposition B.1 and Proposition B.2 below provide sufficient conditions for the existence of an unambiguous ranking over authorities.

Proposition B.1 defines a lower bound for the equilibrium level of effort under the friendly authority and an upper bound for the equilibrium level of effort under the unfriendly authority. These are z_1 and z_2 , respectively. These levels of effort provide bounds for the party's payoff under the two authorities.

PROPOSITION B.1. Fix $q \leq q^{\dagger}$ and let $z_1 = \frac{q[1-c'(q)]}{1-qc'(q)}$ and $z_2 = q[1+(1-q)c'(q)]$. If

$$z_2c'(z_2) > \frac{1 - c'(z_1)}{1 - qc'(z_1)},$$

then the party is better off in any interior equilibrium under the unfriendly authority than in the interior equilibrium under the friendly authority. **Proof.** The party's equilibrium payoff under the friendly authority is $e^F \mu_1^F + (1 - e^F)\mu_0^F = e^F/q$. By equation (A-1), we have:

$$\frac{e^F}{q} = [1 - c'(e^F)] \frac{1 - e^F}{1 - q}.$$
(B-1)

The equilibrium value of μ_0^F is such that $\pi(1) = q$. Hence

$$e^F = \frac{\mu_0^F q}{1 - q(1 - \mu_0^F)} = \frac{\mu_0^F q}{1 - qc'(e^F)},$$

where the second equality follows from equation (1). Using this expression into equation (B-1), we get

$$\frac{e^F}{q} = \frac{1 - c'(e^F)}{1 - qc'(e^F)}.$$

Note that $z_1 \in (0,q)$. Define the function v(x;q) = (1-x)/(1-qx) and observe that $v(\cdot;q)$ is decreasing for every q. If $e^F < z_1$, we have

$$\frac{e^F}{q} < \frac{z_1}{q} = v(c'(q);q) < v(c'(e^F);q) = \frac{1 - c'(e^F)}{1 - qc'(e^F)}.$$

This contradicts equation (B-1), hence Proposition 1. We conclude that $e^F \ge z_1$. This implies that, fixing q, the party's payoff in the interior equilibrium under the friendly authority is bounded above by $\overline{v}^F(q) \equiv [1 - c'(z_1)]/[1 - qc'(z_1)]$.

Consider the unfriendly authority and pick an arbitrary interior equilibrium (e^k, μ^k) . The party's equilibrium is $e^k \mu_1^k = e^k c'(e^k) = (e^k - q)/(1 - q)$. Note that $z_2 > q$. If $e^k \in (q, z_2)$, we have

$$\frac{e^k - q}{1 - q} < \frac{z_2 - q}{1 - q} = qc'(q) < e^k c'(e^k).$$

This contradicts $e^k \mu_1^k = (e^k - q)/(1 - q)$. We conclude $e^k \ge z_2$. The party's payoff is thus bounded below by $\underline{v}^N(q) \equiv z_2 c'(z_2)$.

The party is better off under the unfriendly authority than under the friendly one if $\underline{v}^{N}(q) > \overline{v}^{F}(q)$.

As in the previous result, also Proposition B.2 defines bounds on the equilibrium levels of effort under the two authorities. Differently from Proposition B.1, Proposition B.2 uses the bounds to show that the party's payoff under the friendly is higher than under the unfriendly one.

PROPOSITION B.2. Fix $q \le q^{\dagger}$ and define $z_3 = q + (1-q)\frac{1-c'(q)}{1-qc'(q)}$. If $z_3c'(z_3) > 1$, then the party is better off in the interior equilibrium under the friendly authority than in any interior equilibrium under the unfriendly authority.

Proof. Let $\chi \in (0, 1)$ be the solution of $\chi c'(\chi) = 1$. Since *c* is increasing and c'(1) > 1, this is well defined. By construction and by the convexity of *c*: $c'(1) > c'(\chi) = 1/\chi > 1$.

Pick an interior equilibrium under the unfriendly authority, (e^k, μ^k) . If $e^k \ge \chi$, equation (4) implies that the party's equilibrium payoff is $e^k \mu_1^k = e^k c'(e^k) > 1$. This establishes a contradiction because the approval probability is bounded above by 1. Hence, we must have $e^k < \chi$. We conclude that the party's equilibrium payoff $e^k \mu_1^k = (e^k - q)/(1 - q)$ is bounded above by $\overline{v}^N \equiv (\chi - q)/(1 - q)$.

Now consider the interior equilibrium under the friendly authority. Recall that $e^F < q$. Following the same steps of the proof of Proposition B.1, we can conclude that the party's payoff are bounded below by $\underline{v}^F(q) \equiv \frac{1-c'(q)}{1-qc'(q)}$.

The party prefers the interior equilibrium under the friendly authority to any of the equilibria under the unfriendly authority if $\underline{v}^F(q) > \overline{v}^N(q)$ or equivalently if $z_3 > \chi$. The function $x \mapsto xc'(x)$ is strictly increasing in $x, z_3 > k$ if and only if $z_3c'(z_3) > 1$.

References

- Albano, Gian Luigi and Alessandro Lizzeri, "Strategic Certification and Provision of Quality," *International Economic Review*, 2001, 42 (1), 267–283.
- Alesina, Alberto and Guido Tabellini, "Bureaucrats or Politicians? Part I: A Single Policy Task," *American Economic Review*, 2007, 97 (1), 169–179.
- _ and _ , "Bureaucrats or Politicians? Part II: Multiple Policy Tasks," *Journal of Public Economics*, 2008, 92 (3-4), 426–447.
- Alonso, Ricardo and Odilon Câmara, "Persuading Voters," American Economic Review, November 2016, 106 (11), 3590–3605.
- Ashworth, Scott, "Electoral Accountability: Recent Theoretical and Empirical Work," Annual Review of Political Science, 2012, 15, 183–201.

- Austen-Smith, David, Wioletta Dziuda, Bård Harstad, and Antoine Loeper, "Gridlock and Inefficient Policy Instruments," *Theoretical Economics*, 2019, 14 (4), 1483– 1534.
- Banks, Jeffrey S, "Agency Budgets, Cost Information, and Auditing," American Journal of Political Science, 1989, 33 (3), 670–699.
- Bardhi, Arjada and Yingni Guo, "Modes of Persuasion toward Unanimous Consent," *Theoretical Economics*, 2018, *13* (3), 1111–1149.
- **Bizzotto, Jacopo and Bård Harstad**, "The Choice of Certifier in Endogenous Markets," *Available at SSRN 3730508*, 2020.
- __, Eduardo Perez-Richet, and Adrien Vigier, "Communication via Third Parties," Available at SSRN 3530804, 2021.
- **Boleslavsky, Raphael and Christopher Cotton**, "Grading Standards and Education Quality," *American Economic Journal: Microeconomics*, 2015, 7 (2), 248–79.
- and Kyungmin Kim, "Bayesian Persuasion and Moral Hazard," Working paper, 2021.
- **Bueno de Mesquita, Ethan and Matthew C Stephenson**, "Regulatory Quality under Imperfect Oversight," *The American Political Science Review*, 2007, *101* (3), 605–620.
- **Chan, Jimmy, Seher Gupta, Fei Li, and Yun Wang**, "Pivotal Persuasion," *Journal of Economic Theory*, 2019, *180*, 178–202.
- **Dragu, Tiberiu, Xiaochen Fan, and James H Kuklinski**, "Designing Checks and Balances," *Quarterly Journal of Political Science*, 2014, 9 (1), 45–86.
- Duggan, John and César Martinelli, "The Political Economy of Dynamic Elections: Accountability, Commitment, and Responsiveness," *Journal of Economic Literature*, 2017, 55 (3), 916–84.
- Feng, Xin and Jingfeng Lu, "The Optimal Disclosure Policy in Contests with Stochastic Entry: A Bayesian Persuasion Perspective," *Economics Letters*, 2016, 147, 103–107.
- Gailmard, Sean and John W Patty, "Formal Models of Bureaucracy," Annual Review of Political Science, 2012, 15 (1), 353–377.

- Georgiadis, George and Balazs Szentes, "Optimal Monitoring Design," *Econometrica*, 2020, *88* (5), 2075–2107.
- **Gulotty, Robert and Zhaotian Luo**, "Fire Alarm Fatigue: How Politicians Evade Accountability," *Available at SSRN 3815391*, 2021.
- **Heese, Carl and Stephan Lauermann**, "Persuasion and Information Aggregation in Large Elections," *Working paper*, 2021.
- **Hölmstrom, Bengt**, "Moral Hazard and Observability," *The Bell Journal of Economics*, 1979, pp. 74–91.
- Huber, John D, Charles R Shipan, and Madelaine Pfahler, "Legislatures and Statutory Control of Bureaucracy," *American Journal of Political Science*, 2001, 45 (2), 330– 345.
- **Iaryczower, Matias, Garrett Lewis, and Matthew Shum**, "To Elect or to Appoint? Bias, Information, and Responsiveness of Bureaucrats and Politicians," *Journal of Public Economics*, 2013, 97, 230–244.
- Kamenica, Emir, "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 2019, *11* (1), 249–272.
- and Matthew Gentzkow, "Bayesian Persuasion," American Economic Review, October 2011, 101 (6), 2590–2615.
- Krehbiel, Keith, Pivotal Politics: A Theory of US Lawmaking, University of Chicago Press, 1998.
- **Lizzeri, Alessandro**, "Information Revelation and Certification Intermediaries," *The RAND Journal of Economics*, 1999, pp. 214–231.
- Mas-Colell, Andreu, Michael D Whinston, and Jerry R Green, Microeconomic Theory, Oxford University Press, 1995.
- Maskin, Eric and Jean Tirole, "The Politician and the Judge: Accountability in Government," *American Economic Review*, September 2004, 94 (4), 1034–1054.
- **Miklós-Thal, Jeanine and Heiner Schumacher**, "The Value of Recommendations," *Games and Economic Behavior*, 2013, 79, 132–147.

- Ortner, Juan, "A Theory of Political Gridlock," *Theoretical Economics*, 2017, 12 (2), 555–586.
- **Prato, Carlo and Ian R Turner**, "Institutional Foundations of the Power to Persuade," *Working paper*, May 2022.
- Rodina, David, "Information Design and Career Concerns," Working paper, 2020.
- _ and John Farragut, "Inducing Effort Through Grades," *Working paper*, 2020.
- **Taylor, Curtis R. and Huseyin Yildirim**, "Subjective Performance and the Value of Blind Evaluation," *The Review of Economic Studies*, 02 2011, *78* (2), 762–794.
- Zapechelnyuk, Andriy, "Optimal Quality Certification," *American Economic Review: Insights*, June 2020, 2 (2), 161–76.
- Zhang, Jun and Junjie Zhou, "Information Disclosure in Contests: A Bayesian Persuasion Approach," *The Economic Journal*, 2016, 126 (597), 2197–2217.