Horizontal reputation and strategic audience management

Matthieu Bouvard*, Raphaël Levy†

August 2017

Abstract

We study how a decision maker uses his reputation to simultaneously influence the actions of multiple receivers with heterogenous biases. The reputational payoff is single-peaked around a bliss reputation at which the incentives of the average receiver are perfectly aligned. We evidence two equilibria characterized by repositioning towards this bliss reputation that only differ through a multiplier capturing the efficiency of reputational incentives. Repositioning is moderate in the more efficient equilibrium, but the less efficient equilibrium features overreactions, and welfare may then get lower than in the no-reputation case. Finally, we highlight how strategic audience management (e.g., delegation to third parties with dissenting objectives, centralization) alleviates inefficient reputational incentives, and how multiple organizational or institutional structures may arise in equilibrium as a result.

*McGill University, Desautels Faculty of Management. E-mail: matthieu.bouvard@mcgill.ca.
†HEC Paris. E-mail: levyr@hec.fr.

1 Introduction

The literature has carefully discussed how reputation provides implicit incentives in the absence of formal commitment, and how these incentives may either improve or worsen welfare.[1] However, it has almost exclusively focused on environments where the reputation-concerned party faces an homogenous audience with monotone preferences over his type and actions.[2] In many situations, though, reputation is used to influence audiences composed of heterogenous receivers. For instance, policy-makers devise policies so as to induce efficient behavior (e.g., correct externalities) from a large population of agents with a wide array of preferences and vested interests. Similarly, managers need to get different business units with diverging objectives to work towards the common good of the organization. With such heterogenous audiences, the preferences of each constituent can be captured by his preferred location on an axis along which the reputation-concerned party tries to position himself: reputation is horizontal. There, as a monopolist optimally locates in the middle of the Hotelling segment, the value of reputation is highest at some moderate reputation. The contribution of the paper is two-fold: first, we introduce a tractable infinite horizon framework of horizontal reputation, and show the existence of two equilibria with distinct efficiency properties. Second, we introduce strategic audience management as a natural remedy to the inefficiency of reputation in the presence of heterogenous audiences. In particular, we show how organizational or institutional design may alter the modalities of interaction with the audience and improve welfare. As a result, the multiplicity of reputational equilibria endogenously translates into multiple organizational forms.

We build a model in which a decision maker (e.g., organization leader, policy-maker) tries to influence the investment decisions of heterogenous receivers. Receivers differ in the magnitude of a bias that distorts their investment decisions from the efficient action that the decision maker wants to reach. This discrepancy provides a rationale for policy interventions aiming at realigning incentives. Specifically, the decision maker can affect the environment by uniformly shifting the marginal benefit of investment for all receivers. Because the decision maker lacks commitment power, this intervention is driven by the

[1] For a detailed account on the literature on reputation, see Mailath and Samuelson (2006), or Bar-Isaac and Tadelis (2008).

[2] For instance, in Holmström (1999), the market rewards more managers perceived as more productive.

desire to build a reputation. Reputation is horizontal, in that the reputational payoff is quadratic and reaches a maximum at a bliss reputation at which the incentives of the average receiver in the audience are perfectly aligned.

In equilibrium, the decision maker's actions then aim at reducing his reputational deficit, that is, the distance between his current reputation and this bliss reputation. We derive the existence of two linear equilibria that only differ through a multiplier measuring the responsiveness to this reputational deficit: while responsiveness is moderate in one equilibrium, it is excessive in the other. In the moderate equilibrium, the aggregate investment level is more efficient when the decision maker has reputational concerns than none, that is, reputation provides a welfare-enhancing (though imperfect) substitute to commitment (reputation is "good"). However, in the high-responsiveness equilibrium, welfare is strictly lower than in the moderate equilibrium, and possibly even lower than in the infinitely repeated static game ("bad reputation"). This equilibrium multiplicity arises from intertemporal complementarities between the *current responsiveness* to the reputational deficit and the *efficiency of future responses*. In the high-responsiveness equilibrium, the reactivity to future reputational deficits is inefficiently strong, which makes those deficits more costly to withstand. This in turn raises the current benefit from reaching a better reputation, hence a high current responsiveness. By the same logic, a moderate future responsiveness makes future adjustments more efficient, which justifies current moderation.

In a second stage, we take advantage of the closed-form solution we obtain for the equilibrium payoffs to draw implications for audience management. First, we allow the decision maker to freely choose the composition of his audience, that is, to exclude receivers he prefers out. For instance, a politician can choose the coalition of interests he wants to serve, or an organization can choose its portfolio of activities. Because influencing receivers whose preferences are far from the decision maker's is more costly when reputation is less efficient, endogenous audience selection results in narrower, less diverse and more congruent audiences the less efficient the equilibrium is. We then consider alternative audience management strategies when the decision maker cannot exclude receivers. First, we show that delegating control to an otherwise identical decision maker may be beneficial if the delegate targets a different audience. In particular, the decision maker optimally requests the delegate to target a different receiver from the one he himself targets. This dissent in their objectives allows to correct for the inefficient incentives that reputational concerns induce, hence increases with the inefficiency of reputation. When the decision maker cannot control how the delegate composes his audience, their preferences over the best target audience become endogenously misaligned, which lowers the value of delegation. However, delegation always remains optimal when reputation is bad: when reputational deficits are very costly to withstand, both the DM and the delegate agree on the necessity to target a congruent audience. Such an equilibrium featuring delegation then coexists with one better equilibrium in which reputation is more efficient and delegation accordingly undesirable. Our audience-based motive for delegation provides a rationale for *narrow mandates*, i.e., the requirement for an agency (e.g., a central bank) to pursue restrictive objectives (e.g., price stability) which insulate them from certain audiences (e.g., political pressure). It is also consistent with the tendency of policy-makers to delegate to independent bodies policies which benefit special interests and to later blame them for being insufficiently representative of the electorate's interests ("blame-shifting").

Finally, we consider how the decision maker can improve the impact of his intervention by treating different receivers differentially. One such strategy pertains to the choice between centralization or decentralization in organizations or politics. For instance, in organizations, the top management can either centralize decision making and then uniformly impact all workers, or delegate to division managers the care of aligning incentives of workers in their own business units. We show that decentralization dominates if and only if reputation is good. Intuitively, decentralization allows to tailor division managers' interventions to the idiosyncrasies of their local units, which may improve the efficiency of these interventions. However, under decentralization, managers who target receivers with extreme biases have larger reputational deficits than when targeting the average receiver, as under centralization. As a result, when reputation is bad and decision makers overreact to their reputational deficits, decentralization exacerbates reputational costs, and centralization becomes dominant. This suggests that multiple organizational forms may coexist in equilibrium, with centralization arising in the worse equilibrium as a by-product of the inefficiency of reputation. This also provides a rationale for the fact that organizations often switch back and forth between a centralized and a decentralized structure (Eccles, Nohria, and Berkley, 1992; Nickerson and Zenger, 2002). Another strategy is to grant exemptions, whereby the decision maker insulates a fraction of his audience from the impact of his intervention. In this case, the optimal exemption strategy consists of exempting the extreme receivers which the decision maker is relatively less able to influence to boost his credibility, hence his impact, with respect to receivers at the other extreme.

Our paper builds on the seminal model of "career concerns" by Holmström (1999), where an agent jams the market's inference about his type by exerting costly unobservable effort. The key difference is that the reputational payoff is linear in Holmström, and the equilibrium strategy is accordingly independent of the reputation. By contrast, in our setting with single-peaked reputational concerns, the equilibrium strategies always depend on the reputation, and the concavity of the payoff function generates multiple equilibria and possibly inefficient reputation-building. Second, our paper relates to Cisternas (2017), who studies signal-jamming in continuous time and derives general conditions under which his equilibrium strategy can be characterized by a first-order approach. While Cisternas focuses on the dynamics of incentives within an equilibrium, that is, how shocks to the agent's reputation change his future incentives (the "ratchet effect"), we show how intertemporal complementarities emerge to generate multiple equilibria. In addition, we consider a tractable quadratic specification where the actions of the decision maker have an impact on efficiency (i.e., do not serve the sole purpose of jamming the market's inference). In this context, we can show the existence of multiple equilibria in closed form and establish their distinctive efficiency properties as compared to the no-reputation case. This multiplicity in turn results in multiple organizational forms.

The paper also relates to a recent literature on multi-audience reputation. A stream of papers has analyzed how the presence of multiple audiences may generate non-monotone reputational payoffs. Bar-Isaac and Deb (2014b) shows that a monopolist discriminating horizontally differentiated market segments may derive a profit non-monotonic in his reputation; Bouvard and Levy (2017) establish that a certifier who needs to attract sellers and buyers reaches his maximum profit when his reputation for accuracy is interior. In Shapiro and Skeie (2015), a bank regulator faces ambiguous reputational incentives: a stronger tendency to bail out distressed institutions reassures depositors but induces banks to take excessive risk. Bar-Isaac and Deb (2014a) also consider the impact of the environment on reputational incentives by contrasting reputation building with two audiences under common or separate observation of actions, and show that separate observation may cause reputation to lower welfare. As we do, these papers obtain repositioning towards the bliss reputation, but all of these papers consider two-period environments only. Instead, our infinite horizon analysis allows to establish that (a) multiple equilibria coexist, while the equilibrium is unique in any finite version of the game, (b) dwelling on implications drawn from the two-period case is misguided: for instance, increasing the quality of monitoring always improves welfare in the stationary case, but may decrease it in the two-period game. In addition, none of these papers considers audience management.

Finally, our focus on whether reputation improves or worsens welfare relates us to models of "bad reputation" (Morris, 2001; Ely and Välimäki, 2003; Ely, Fudenberg, and Levine, 2008), in which an "honest" type ends up taking actions detrimental to the audience to separate from biased types. By contrast, such separating strategies are impossible in our model, as information remains symmetric in equilibrium. Accordingly, bad reputation does not stem from the fact that reputation provides incentives to act in the wrong direction, as in those papers. Instead, reputational incentives always go in the right direction, but sometimes lead the decision-maker to go too far in that direction ("overshooting"), which, given the single-peakedness of his payoff function, impairs welfare.

The remainder of the paper is structured as follows. We introduce the model in Section 2. In Section 3, we analyze reputation-building, derive the existence of multiple equilibria, and examine their welfare and comparative statics properties. In Section 4, we examine strategic audience management and its implications for organizational design. Section 5 concludes.

2 The model

2.1 Setup

We consider a long-lived decision maker (later "DM") who interacts at every period t with a mass one of short-lived receivers. In period t, each receiver takes an action (investment, effort) $y_t \in \mathbb{R}$ that generates a payoff $y_t - \frac{y_t^2}{2}$ to the DM.

While it would be socially efficient to play $y_t = 1$, receivers have a preferred action $y_t = 1 - b$ that deviates from efficiency by an idiosyncratic bias b distributed according to some c.d.f. $F(b)$ on a support $\mathcal{B}$. The DM, who maximizes social surplus, has one instrument at hand which he uses to correct this misalignment of incentives. The impact of his intervention is captured by a variable x_t which shifts incentives of all receivers in a uniform way.[3]

Specifically, the private surplus of Receiver b given the DM's intervention x_t reads

$$(1 - b + x_t)y_t - \frac{y_t^2}{2}. \tag{1}$$

The impact of the intervention x_t is stochastic and only partially controlled by the DM. x_t is decomposed as follows:

$$x_t \equiv \theta_t + a_t + \varepsilon_t,$$

where $\theta_t \in \mathbb{R}$ is the DM's type, $a_t \in \mathbb{R}$ is an action the DM takes at a private cost $\gamma \frac{a_t^2}{2} \geq 0$, and ε_t is an i.i.d. shock.

A critical assumption is that neither a_t nor x_t can be observed by receivers when they choose their actions $y_t(b)$. Receiver b then maximizes his expected surplus, hence chooses

$$y_t(b) = 1 - b + \mathbb{E}_t(x_t),$$

where $\mathbb{E}_t(x_t)$ denotes receivers' expectation of x_t given their information at date t.

The DM's payoff in any period t is equal to the expected social surplus:

$$\int_{b \in \mathcal{B}} [y_t(b) - \frac{1}{2} y_t(b)^2] \, dF(b)$$

$$= \frac{1 - \mathbb{V}(y_t)}{2} - \frac{1}{2} \left(1 - \mathbb{E}(y_t)\right)^2. \tag{2}$$

Given $y_t(b) = 1 - b + \mathbb{E}_t(x_t)$, (2) becomes

$$\frac{1 - \mathbb{V}(b)}{2} - \frac{1}{2} [\mathbb{E}_t(x_t) - \bar{b}]^2, \tag{3}$$

[3] We deliberately make this assumption in a first stage to examine how the DM tries to simultaneously influence several audiences. In Section 4, we allow the DM to design the environment in such a way to differentiate the impact of his intervention across receivers.

where $\bar{b} \equiv \mathbb{E}(b) = \int_{b \in \mathcal{B}} b \, dF(b)$.

From (3), the maximal surplus is then attained when receivers' expectation of the DM's intervention $\mathbb{E}_t(x_t)$ perfectly adjusts the incentives of the receiver with the average bias $\bar{b}$, that is, when investment is on average efficient ($\mathbb{E}(y_t) = 1$). However, even in this ideal case, the DM's payoff deviates from the maximal social surplus attainable by a term proportional to the dispersion of the receivers' biases, $\mathbb{V}(b)$, reflecting that the DM's uniform impact is imperfectly tailored to each receiver's idiosyncratic bias. At the end of this section, we discuss the possible interpretations of this setup (Section 2.2), as well as alternative specifications that would generate a similar reputational payoff (Section 2.3).

The expression in (3) makes transparent how the DM's payoff depends on receivers' expectation about his intervention rather than his actual intervention.[4] This creates scope for reputation-building, as the DM would like to influence receivers and have them believe that his intervention exactly offsets the average bias $\bar{b}$.

We build on the "career concerns" setup pioneered by Holmström (1999) and assume that the DM and receivers are symmetrically informed about the DM's type. The initial type of the DM, θ_1, is drawn from a normal distribution with mean m_1 and precision (i.e., inverse variance) h_1. Besides, his type θ_t is subject to repeated shocks, but exhibits persistence: for all $t \geq 1$, $\theta_{t+1} = \theta_t + \eta_t$, where η_t are i.i.d. normal variables with zero mean and precision h_η. Finally, ε_t are i.i.d. normally distributed with mean 0 and precision h_ε. The variables θ_1, ε_t and η_t are mutually independent for all t.

Reputation-building is possible because there is ex post learning on the DM's past interventions. We assume that receivers' actions and payoffs are publicly observed once realized, so that $x_t = \theta_t + a_t + \varepsilon_t$ can be inferred from (1). Despite the action a_t being privately observed by the DM, receivers can update their beliefs on the DM's type for any given action a_t^e they might expect. Given the normality and independence assumptions, the dynamics of beliefs is simple to characterize: the conditional distributions of the DM's type at any date t is Normal with mean m_t and precision h_t. For a given action a_t^e that

[4] It would be possible to enrich the model by allowing the DM's payoff to directly depend on θ_t and/or a_t, at the cost of more complexity, but this would not generate qualitatively different results.

receivers expect the DM to play, the motions of m_t and h_t are given by:

$$m_{t+1} = \frac{h_t}{h_t + h_\varepsilon} m_t + \frac{h_\varepsilon}{h_t + h_\varepsilon}[x_t - a_t^e], \tag{4}$$

and

$$h_{t+1} = \frac{(h_t + h_\varepsilon)h_\eta}{h_t + h_\varepsilon + h_\eta}. \tag{5}$$

Since the motion of the variance of beliefs (5) is exogenous, hence does not affect the DM's problem, the critical state variable we focus on is m_t, which we call the DM's *reputation* at date t.

2.2 Interpretation

We interpret the DM as a planner (organization, policy-maker) who uses a policy instrument to influence his environment and achieve a better alignment between individual and social incentives. For instance, policy-makers (or public agencies mandated by them, e.g., central banks, regulators) need to design policies such that agents (at least partly) internalize the externality that their actions (e.g., labor supply, R&D investments, savings, location choice...) inflict on others. Similarly, organizations strive to provide the right effort incentives (e.g., foster synergies, encourage information acquisition...) to workers with different preferences or skills. In either case, reputation is instrumental. In politics, it is impossible for politicians to write contracts with all their potential stakeholders, hence the importance of maintaining a reputation for fear of alienating some key players (e.g., international lenders, large corporations, top taxpayers...). In organizations, although formal contracts are widely used, contractual frictions often result in imperfectly aligned incentives. Alternatively, misaligned incentives could endogenously arise as an optimal contractual form in organizations facing a tradeoff between the scale benefits of centralized processes and the cost of imperfect adaptation to local conditions (Dessein, Garicano, and Gertner, 2010). In any case, corporate reputation ("corporate culture") is a key complement to explicit contracts, and plays a critical role in enabling coordination (Kreps, 1990).

In this context, the DM's reputation captures how much receivers expect their investment or effort to be rewarded. The DM's type may accordingly capture his intrinsic ability to create a propitious environment. For instance, it could account for a policy maker's ability to design a policy (e.g., fiscal, monetary, trade policy) that provides the right incentives to invest, or a manager's talent at identifying the contribution of different workers to a common project and appropriately reward it. Alternatively, in the spirit of Carrillo and Gromb (1999)'s view of corporate culture as a production technology, θ could also capture a firm's technological structure, in particular the type of skills that generate the largest match value with the firm's (tangible or intangible) assets.[5] Finally, the type θ could account for some unknown state of the world which governs the returns on investment in the economy or the organization (e.g., the level of inflation, or the severity of technological, financial or institutional constraints). The DM can undertake costly actions to increase or decrease the return on investment beyond its intrinsic level. The cost of these actions may be monetary (e.g., subsidies), account for the disutility of the effort required to distort the impact of one's intervention away from its intrinsic level, or capture any indirect costs that affect the DM through other channels than receivers' investments (e.g., rents resulting from moral hazard in the implementation of the policy, bargaining or political costs, etc). The assumption that the cost of a_t is symmetric in the positive and negative ranges, while ensuring tractability, allows more flexible interpretations.[6]

For the DM, these actions serve the purpose of jamming the audience's inference about his type in order to get closer to his bliss reputation. But they also do change the true returns on investment for the receivers, hence have an impact on its aggregate efficiency. Therefore, the role of reputation is two-fold. First, it has a direct influence on receivers, as their beliefs about the DM's type m_t affect their actions. Second, reputation provides the DM with commitment power to take actions a_t that also influence receivers' decisions. For instance, investment and trading decisions depend about inflation expectations, but also about central banks' interventions, e.g., on foreign exchange markets or bond markets, to bring inflation closer to its target level. In turn, central banks adjust the magnitude of these interventions as a function of their current reputations, and to how much they are willing to maintain or improve it.[7]

[5] Implicit here is the view that skills may be horizontally differentiated, that is, two agents with different skills may each be more productive in two different firms.

[6] For instance, if y_t is an investment which generates pollution, a positive a_t may be interpreted as an implicit subsidy to the industry, and a negative a_t as a reward for eco-friendly investments. Likewise, according to the interpretation, a_t may capture subsidies to consumption or savings.

[7] In line with our modeling, these interventions are typically imperfectly observed, and they affect

2.3 Alternative specifications

Our analysis more generally applies to any specification where the DM's payoff is quadratic in his audience's anticipation about his positioning x_t. For instance, consider a setup in which receivers want to coordinate their actions with the DM's, adjusted for a bias b, i.e., have a payoff

$$-(y_t - x_t + b)^2.$$

If the DM wants receivers' actions to match a target normalized to 0, i.e., has a payoff

$$\frac{1 - \int_{b\in\mathcal{B}} y_t^2(b)\, dF(b)}{2},$$

then the DM's payoff also equals (3).

The DM may also want to maximize the monetary payments he extracts from the receivers. For instance, consider a media organization raising revenues from selling advertising slots, and let x_t capture journalistic integrity. Suppose that viewers attach a positive value to integrity, while advertisers care about reaching viewers but otherwise dislike integrity, as they value the ability to sway the editorial line and avoid content that is damaging to their interests. Formally, if the mass of viewers is equal to perceived integrity $\mathbb{E}_t(x_t)$ and the willingness to pay per viewer of advertiser b is $b - \mathbb{E}_t(x_t)$, the medium obtains a quadratic profit $\int_{b\in\mathcal{B}} \mathbb{E}_t(x_t)[b - \mathbb{E}_t(x_t)]\, dF(b)$, so that his ideal (perceived) integrity is again $\bar{b}$. Alternatively, a politician raising campaign financing from lobbies that derive utility $-(x_t - b)^2$ from the future policy x_t obtains a quadratic profit maximized at $E_t(x_t) = \bar{b}$ when able to fully extract the surplus from each lobby.

Overall, our setup is meant to capture any kind of situation where the DM's reputational concerns are "horizontal," in that the need to accommodate several receivers with heterogenous or conflicting preferences over the DM's type and actions leads to a preference for an intermediate reputation. We now analyze how these horizontal reputational incentives shape the DM's actions.

heterogeneous agents in an undifferentiated way (Farhi and Tirole, 2012).

3 Equilibrium analysis

3.1 The commitment benchmark

Before we go through the analysis of reputation building, let us first derive the optimal profile of actions under full commitment. Let

$$\pi(x) \equiv \frac{1-\mathbb{V}(b)}{2} - \frac{1}{2}(x-\bar{b})^2 \tag{6}$$

denote the gross surplus function of the DM. His total payoff in period t then reads

$$\pi(m_t + a_t) - \gamma\frac{a_t^2}{2} = \frac{1-\mathbb{V}(b)}{2} - \frac{(m_t + a_t - \bar{b})^2}{2} - \gamma\frac{a_t^2}{2}, \tag{7}$$

which is maximized at $a_t = a^{FB}(m_t) \equiv \frac{1}{1+\gamma}(\bar{b} - m_t)$.

In the first best, the DM tries to correct for the intrinsic impact of his intervention m_t to push it closer to his ideal impact $\bar{b}$. Accordingly, $\mathbb{E}(x_t) = \frac{1}{1+\gamma}\bar{b} + \frac{\gamma}{1+\gamma}m_t$ is a weighted average between m_t and $\bar{b}$, with weights depending on the cost for the DM to steer away from his intrinsic impact.

3.2 The two-period case

To provide a first intuition, we begin with the analysis of the two-period game. In period 2, since his payoff only depends on the expected but not on the actual x_t, and since the DM has no reputational concerns, he optimally selects $a_2^* = 0$ no matter his reputation m_2. Therefore, his total payoff in period 2 is $\pi(m_2 + a_2^*) - \gamma\frac{a_2^{*2}}{2} = \pi(m_2)$. Denoting δ the discount factor of the DM, and using (4), the equilibrium action in period 1 a_1^* satisfies

$$a_1^* \in \underset{a_1}{\operatorname{argmax}}\, \delta\mathbb{E}\pi\left\{\frac{h_1}{h_1+h_\varepsilon}m_1 + \frac{h_\varepsilon}{h_1+h_\varepsilon}[\theta_1 + \varepsilon_1 + a_1 - a_1^*)]\right\} - \gamma\frac{a_1^2}{2}$$

Since π is concave, a_1^* is the unique solution to

$$\delta\frac{h_\varepsilon}{h_1+h_\varepsilon}\mathbb{E}\pi'\left\{\frac{h_1}{h_1+h_\varepsilon}m_1 + \frac{h_\varepsilon}{h_1+h_\varepsilon}[\theta_1+\varepsilon_1]\right\} - \gamma a_1^* = 0$$

$$\Leftrightarrow a_1^* = \frac{\delta h_\varepsilon}{\gamma(h_1+h_\varepsilon)}\left(\bar{b} - m_1\right).$$

Proposition 1. *The two-period game admits a unique equilibrium:* $a_1^* = k_1(\bar{b} - m_1)$, *where* $k_1 \equiv \frac{\delta h_\varepsilon}{\gamma(h_1+h_\varepsilon)}$.

Let us single out two important features of Proposition 1, which will prove relevant to the understanding of the stationary case. First, the equilibrium is unique. Second, the DM's equilibrium action depends on his current reputation m_1; more precisely, it aims at correcting his reputational deficit $\bar{b} - m_1$, that is, how far m_1 falls away from his bliss reputation.[8] If the DM is perceived as being overly rewarding investment, receivers should rationally expect him to take an action which lowers the marginal benefit of investment, and conversely. The magnitude of this correction depends on a multiplier k_1 that captures the strength of reputational concerns. Notice that a_1^* and a_1^{FB} have the same sign: reputational concerns provide incentives to reach reputations closer to $\bar{b}$, which is achieved by distorting x_t in the direction of $\bar{b}$, as in the first best. However, $a_1^* \neq a_1^{FB}$ generically, and the equilibrium may feature both underreactions ($|a_1^*| < |a_1^{FB}|$) and overreactions ($|a_1^*| > |a_1^{FB}|$) as compared to the first best. This inefficiency will play a critical role in the construction of equilibria in the stationary case, which we now turn to.

3.3 The stationary case

In this section, we analyze the asymptotic state of the infinite horizon game, where the precision of receivers' information about the DM's type h_t is constant across periods. The dynamics of h_t is driven by two opposite forces. On the one hand, players learn about θ_t upon observing past values of x. Since there is persistence in the DM's type, this increases the precision of beliefs on θ_{t+1}. On the other hand, because θ_t changes according to unobservable shocks η_t, each period brings additional uncertainty. The precision always converges to a steady state value at which these two effects exactly offset:[9]

$$h_t \underset{t\to+\infty}{\to} h \text{ with } h = \frac{(h+h_\varepsilon)h_\eta}{h+h_\varepsilon+h_\eta} \Leftrightarrow h = \frac{\sqrt{h_\varepsilon^2 + 4h_\eta h_\varepsilon} - h_\varepsilon}{2} \tag{8}$$

In what follows, we focus on this steady state, and assume that $h_1 = h$, i.e., the variance of the distribution of types never changes. This simplifies the analysis, as beliefs

[8] This notably contrasts with Holmström (1999), where equilibrium actions are independent of the reputation.

[9] Since a_t has no impact on the motion of the precision, this holds independently of the DM's actions.

on θ_t given any history of the game are fully characterized by the mean of the posterior distribution. However, since deviations are observed by the DM but not by receivers, we still need to keep track of two state variables: (a) the receivers' beliefs about the mean of θ_t, which we denote m_t and call the DM's public reputation, and (b) the DM's private beliefs about his type, which we denote m_t^{DM} and call the DM's private reputation.

In the stationary case, (4) becomes

$$m_{t+1}(a_t, a_t^e) = \lambda m_t + (1-\lambda)[\theta_t + \varepsilon_t + a_t - a_t^e], \tag{9}$$

where $\lambda \equiv \dfrac{h}{h+h_\varepsilon}$.

Instead, the motion of the private reputation never depends on the profile of actions:

$$m_{t+1}^{DM} = \lambda m_t^{DM} + (1-\lambda)(\theta_t + \varepsilon_t). \tag{10}$$

We restrict attention to Markovian strategies $a(m_t^{DM}, m_t)$ that are functions of those two state variables only. Since deviations are not detectable and players start with a common prior, this implies that, if $a(m_t^{DM}, m_t)$ is an equilibrium strategy, the audience must believe that the DM plays $a_t^e = a(m_t, m_t)$ in equilibrium.

Let $V(m_t^{DM}, m_t)$ denote the expected discounted payoff of the DM when his private reputation is m_t^{DM} and his public reputation is m_t. An equilibrium features a value function $V(.,.)$ and a strategy $a(.,.)$ such that for any pair (m_t^{DM}, m_t) :

i) given $V(.,.)$ and receivers' expectations about his action a_t^e, the DM chooses the period-t action optimally:

$$a(m_t^{DM}, m_t) \in \underset{a_t}{\operatorname{argmax}}\, \delta \mathbb{E} V\left(m_{t+1}^{DM}, m_{t+1}[a_t, a_t^e]\right) - \gamma \frac{a_t^2}{2}, \tag{11}$$

ii) receivers have rational expectations:

$$a_t^e = a(m_t, m_t), \tag{12}$$

iii) $V(.,.)$ satisfies a Bellman optimality condition:

$$\begin{aligned} V(m_t^{DM}, m_t) &= \pi[m_t + a(m_t, m_t)] - \gamma \frac{a(m_t^{DM}, m_t)^2}{2} \\ &+ \delta \mathbb{E} V\left(m_{t+1}^{DM}, m_{t+1}[a(m_t^{DM}, m_t), a(m_t, m_t)]\right), \end{aligned} \tag{13}$$

iv) $V(.,.)$ satisfies a transversality condition:

$$\lim_{t \to +\infty} \delta^t \mathbb{E}_1 V\left(m_t^{DM}, m_t(\hat{a})\right) = 0, \tag{14}$$

where $m_t(\hat{a})$ is the public reputation when receivers expect the DM to follow the equilibrium strategy $a(.,.)$, but he follows an arbitrary strategy $\hat{a}$ between 0 and t instead.[10]

Note that strategies describe the DM's behavior both on and off-path. In particular, condition (11) states that the DM's action is optimal even following an undetected deviation (i.e., if $m_t^{DM} \neq m_t$).[11]

Proposition 2. *There exist two Markovian equilibria in linear strategies of the form* $a^*(m_t^{DM}, m_t) = \beta_1 m_t^{DM} + \beta_2 m_t + \beta_3$. *On the equilibrium path,*

- $m_t^{DM} = m_t$ *and the DM plays an action*

$$a^*(m_t) = k(\bar{b} - m_t), \quad \text{with } k \in \left\{\underline{k}, \overline{k}\right\} \text{ and } 0 \leq \underline{k} \leq \overline{k}.$$

- *The DM's value function in the equilibrium with multiplier k reads*

$$V^k(m_t) = \frac{1}{2(1-\delta)} \left(1 - \mathbb{V}(b) - K(\bar{b} - m_t)^2 - K\Sigma\right), \tag{15}$$

where $K \equiv (1-k)^2 + \gamma k^2$ and Σ is a constant.

- *The equilibrium with multiplier $\underline{k}$ yields a higher value:*

$$V^{\underline{k}}(m_t) \geq V^{\overline{k}}(m_t) \text{ for all } m_t.$$

[10] More precisely, we require that, if (14) does not hold for an admissible $\hat{a}$, then $\hat{a}$ is dominated by a strategy that satisfies (14) (see Appendix).

[11] Notice in this respect that, as long as receivers expect the DM to play Markovian linear strategies, this is a best response for the DM to do so, both on and off the equilibrium path.

In the DM's value function (15), $K = (1-k)^2 + \gamma k^2$ captures the efficiency of the DM's action. When the DM plays the equilibrium strategy $a_t = k(\bar{b} - m_t)$, his payoff in t actually reads

$$\pi[m_t + k(\bar{b} - m_t)] - \gamma \frac{k^2(\bar{b} - m_t)^2}{2} = \frac{1 - \mathbb{V}(b)}{2} - \frac{K}{2}(\bar{b} - m_t)^2. \tag{16}$$

As Figure 1 shows, K is U-shaped in k, which evidences the two effects of the DM's actions described in Section 2.2: while the second term γk^2 captures the fact that a higher responsiveness k is more costly (while still failing to manipulate receiver beliefs), the first term $(1-k)^2$ measures the impact of the DM's intervention on the efficiency of average investment $\mathbb{E}(y_t) = 1 - (1-k)(\bar{b} - m_1)$, the perfect alignment of incentives being reached at $k = 1$. These two effects alternatively dominate, and the highest efficiency is reached at $k^{FB} \equiv \frac{1}{1+\gamma}$, i.e., at the (first best) level of responsiveness the DM would like to commit to. K accordingly measures how good a substitute for commitment reputation is (the lower K the more efficient reputation).

To understand where the forms of the equilibrium actions come from, consider the impact of a marginal deviation from the equilibrium behavior $a_t = k(\bar{b} - m_t)$ in period t. Such a deviation has two consequences: first, it affects the distribution of all future public beliefs m_τ for $\tau > t$, hence all the DM's future payoffs; second it creates an information asymmetry between the DM and receivers, i.e., m_τ and m_τ^{DM} cease to coincide, which, in turn, may affect the DM's future optimal strategy. A necessary equilibrium condition is that such a deviation be non-profitable even if the DM ignores this second effect and continues to play as if private and public beliefs still coincided.[12]

Viewed from period t, the expected payoff in $t+i$ is

$$\frac{1 - \mathbb{V}(b)}{2} - \frac{K}{2}[\bar{b} - \mathbb{E}_t(m_{t+i})]^2 - \frac{K}{2}\mathbb{V}_t(m_{t+i}). \tag{17}$$

The impact of a_t on m_{t+1} is linear and corresponds to the weight receivers put on the period-t signal when updating beliefs, $1 - \lambda$. In turn, m_{t+1} has a persistent effect of magnitude λ^{i-1} on m_{t+i} (see (9)).[13] Overall, the marginal effect of a_t on the payoff in

[12]We refer the reader to the appendix for the sufficiency part of the argument and a full-blown derivation of strategies on and off-path.

[13]Notice that the impact of a_t on m_{t+i} is deterministic, meaning that the variance of future reputations

period $t+i$ is

$$(1-\lambda)\lambda^{i-1}K[\bar{b}-\mathbb{E}_t(m_{t+i})].$$

Summing up across periods and using the martingale property of beliefs, the benefit of a marginal deviation from $a_t = k(\bar{b}-m_t)$, that is, its impact on the discounted sum of future payoffs reads

$$(1-\lambda)\sum_{i=1}^{+\infty}\delta^i\lambda^{i-1}K\left[\bar{b}-\mathbb{E}_t(m_{t+i})\right] = \frac{\delta(1-\lambda)}{1-\delta\lambda}K(\bar{b}-m_t), \tag{18}$$

while its marginal cost reads

$$\gamma k(\bar{b}-m_t).^{14} \tag{19}$$

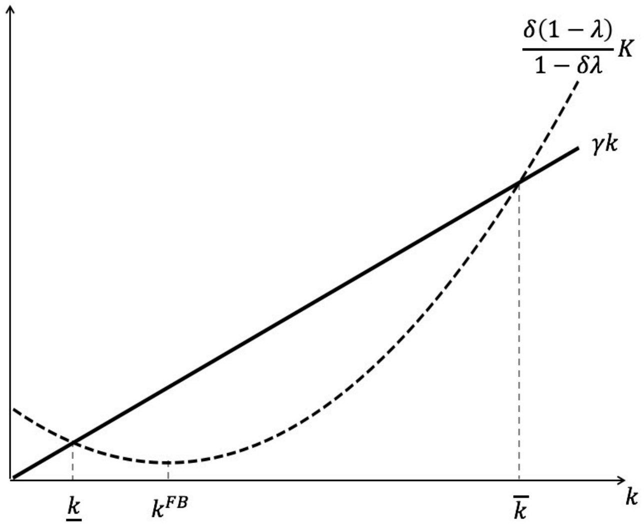

Figure 1: Marginal Cost (solid) and Marginal Benefit (dashed) of reactivity k.

In a stationary equilibrium, the multiplier k must be the same in every period, meaning that k must satisfy a fixed point condition given by the equality of (18) and (19). As illustrated in Figure 1, there are two fixed points, corresponding to two equilibria: one low-responsiveness equilibrium $\underline{k}$ where the DM underreacts ($\underline{k} \leq k^{FB}$) and one high-responsiveness equilibrium $\bar{k}$ where the DM overreacts ($\bar{k} \geq k^{FB}$). In the latter equilib-

does not depend on the DM's actions. From (17), one sees that the DM cares about the risk that future reputations m_{t+i} end up far away from $\bar{b}$, which, given the curvature of π, is costly to him.

[14]Notice the critical role played by the martingale property of beliefs, which allows to express the marginal benefit of a_t, which depends on the expected future reputations, as a function of the current reputation, i.e., in the same "unit" as the marginal cost. This is why we make the important assumption that the interaction between the DM and the audience is long-standing and that the DM can never exit the market, even following large shocks to his reputation.

rium, the stronger reactivity is (relatively) less efficient. This, in turn, makes it more costly for the DM to see his future reputation move far away from the bliss reputation $\bar{b}$, and raises the marginal benefit from reacting today (18), hence a high responsiveness. Conversely, the anticipation of more efficient (moderate) future reactions sustains a moderate current reaction.

We close this section with a discussion on the source of equilibrium multiplicity. Notice first that the two-period game features a unique equilibrium, as would any finite-horizon version of the model.[15] Indeed, intertemporal complementarities arise because different expectations about future actions generate different current incentives. With a finite horizon, the last period action is uniquely determined, and a backward-induction argument implies, in turn, a unique equilibrium action in every previous period.[16]

Second, the complementarity between current and future responsiveness is driven by the concavity of the payoff function π. When receivers expect the DM to be highly responsive, they discount more aggressively the signal x_t when updating their beliefs. If the DM's action does not match expectations, his reputation is then likely to be pushed in a region far from the bliss reputation where the payoff function is very steep. Accordingly, concavity raises the marginal benefit of the DM's action and helps sustain the high-reactivity equilibrium. Conversely, in the low-responsiveness equilibrium, even if the DM were not to match receivers' (moderate) expectations, changes in his reputation would likely be relatively smaller, hence more affordable given the curvature of the payoff function.

Finally, intertemporal complementarities are reinforced by the impact of the DM's action on receivers' payoffs. Intuitively, in the low-responsiveness equilibrium, the marginal benefit of the action is low not only because the DM is expected to expend little in the future to correct his reputational deficit, but also because his future actions are relatively efficient at correcting the average bias. This, in turn, provides lower incentives for the DM to try and adjust his reputation. Conversely, in the high-responsiveness equilibrium, both effects combine to make it more costly for the DM to let his reputation slip away from $\bar{b}$. If we were to consider a setup where the DM cares about receivers' expectations

[15] This contrasts with Dewatripont, Jewitt, and Tirole (1999), where multiple equilibria resulting from complementarities in the technology of learning arise even in the two-period case.

[16] In the T-period game, the unique equilibrium strategy converges to $\underline{k}$ as $T \to \infty$.

about his type but not about his action, equilibrium existence would require an additional restriction that the DM's action is costly enough relative to the discount factor.[17]

3.4 Welfare: Good and bad reputation

As discussed above, the impact of reputation-building on welfare is two-fold: on the one hand, it lowers welfare because the actions a_t are costly and the attempts to manipulate the beliefs of the audience vain; on the other hand, reputation provides some commitment power to take actions closer to efficient. The total impact of reputation on welfare is therefore potentially ambiguous.[18] Before investigating the welfare properties of each equilibrium, let us introduce the following definition:

Definition *We say that* reputation is good (resp. bad) *when the DM obtains a equilibrium payoff larger (smaller) than in the infinitely repeated static game.*

We then refer to "bad reputation" to describe situations where the DM would like to commit not to build a reputation.[19]

Proposition 3. *In the low-responsiveness equilibrium $\underline{k}$, reputation is good for any reputation level m_t. In the high-responsiveness equilibrium $\overline{k}$, reputation is good (for any reputation m_t) if $\overline{k} \leq \frac{2}{1+\gamma} \Leftrightarrow \gamma \leq \frac{\delta(1-\lambda)}{(1-\delta\lambda)+(1-\delta)}$. Otherwise, it is bad.*

In the moderate equilibrium, one has $0 \leq \underline{k} \leq \frac{1}{1+\gamma}$: this equilibrium exhibits the familiar pattern that reputation alleviates moral hazard in helping the DM commit to take more efficient actions than in the no-reputation case, but are generically insufficient to reach efficiency.[20] On the contrary, the equilibrium $\overline{k}$ features excessive responsiveness: $\overline{k} \geq 1 \geq \frac{1}{1+\gamma}$. Actually, when $\overline{k} > \frac{2}{1+\gamma}$, the DM not only overreacts to his reputational deficit compared to the first best, but the overreaction is so large that he ends up being

[17] In his quadratic (continuous-time) specification where the DM's payoff depends on expectations about his type, but not his action, Cisternas (2017) derives a similar condition.

[18] By welfare, we mean here the expected discounted payoff of the DM. In the applications we consider, we implicitly have in mind benevolent planners, but the DM's payoff need not coincide with social welfare.

[19] While this term was coined by Ely and Välimäki (2003) to illustrate that reputation may shut down gains from trade, the mechanics of bad reputation in their paper largely differs from ours (see Section 3.5 below).

[20] See Holmström (1999).

worse off than in the no-reputation case. Notice that the result that reputation decreases welfare may hold even if one abstracts from the costs borne by the DM to build a reputation. Indeed, the gross payoff of the DM in period t

$$\pi(m_t + a_t^*) = \frac{1}{2}\left(1 - \mathbb{V}(b) - (k-1)^2(m_t - \bar{b})^2\right) \tag{20}$$

is lower when the DM has reputational concerns than when he has none ($k = 0$) provided

$$(k-1)^2 > 1 \Leftrightarrow k > 2.$$

Accordingly, overshooting may be as large as to decrease the average efficiency of investment. Finally, since $\overline{k} > 1$, the average level of investment $1-\bar{b}+\mathbb{E}(x_t) = 1-(1-\overline{k})(\bar{b}-m_t)$ is decreasing in m_t in the overreaction equilibrium. Therefore, overshooting also results in reversals, in that the average investment becomes negatively correlated with the DM's intrinsic ability to reward investment.[21]

3.5 Comparative statics

In both equilibria, the DM tries to reposition in the direction of the bliss reputation $\bar{b}$. In this section, we examine how the magnitude of this repositioning depends on the key parameters of the model.

Proposition 4. *An increase in δ or h_ε, or a decrease in h_η causes the DM to be more responsive in the low-responsiveness equilibrium ($\underline{k}$ increases), and less responsive in the high-responsiveness equilibrium ($\overline{k}$ decreases).*

Proof In the Appendix.

Figure 1 helps understand the result. There, we have plotted the marginal benefit and marginal cost of increasing the responsiveness k (given by (18) and (19)). As one sees, the slope of the marginal benefit curve is smaller than the slope of the marginal cost line at $\underline{k}$

[21] Such a reversal is reminiscent of the "It takes a Nixon to go to China" effect (Cukierman and Tommasi, 1998), whereby politicians with a reputation on one side of the political spectrum become more likely to implement policies preferred by voters of the other side than politicians of the other camp themselves. In a similar vein, Kartik and Van Weelden (2014) show that voters may prefer a less congruent politician over some range of beliefs in which more congruent politicians are more eager to build a reputation, hence indulge more in pandering, which ultimately hurts voters.

and larger at $\overline{k}$. Therefore, any parameter change which raises the marginal benefit, i.e., shifts the dashed curve upwards (e.g., an increase δ, h_ε or $-h_\eta$) should be compensated by an increase in the cost (i.e., a higher k) at $\underline{k}$, and a decrease in the cost (in k) in the overreaction equilibrium. Given that K is decreasing in k at $\underline{k}$ and increasing at $\overline{k}$, we immediately derive the following corollary:

Corollary 1. *Any equilibrium is more efficient when δ and h_ε increase, and when h_η decreases.*

By more efficient, we mean that, for any realization of m_t, the net surplus of the DM in period t is larger.[22] A common feature of both equilibria is that more salient reputational concerns help the DM realign his course of action with the efficient one, i.e., the one he would like to commit to. This result stands in contrast with the comparative statics of the two-period equilibrium. There, a increase in, say, the quality of monitoring h_ε makes reputation more salient, hence increases the responsiveness k_1. This in turn lowers welfare if $k_1 > k^{FB}$. Meanwhile, an increase in h_ε always increases welfare in the stationary equilibrium. This shows that dwelling on the two-period model to derive policy implication can be misguided.[23]

Corollary 1 implies that the equilibrium is more efficient when δ increases. In particular, one easily shows that $\underline{k}$ tends to k^{FB} as δ goes to 1, a result reminiscent of folk theorems in repeated games. However, the fact that the inefficient equilibrium also becomes less inefficient when δ increases notably contrasts with the results derived in the literature on bad reputation (Morris, 2001; Ely and Välimäki, 2003; Ely, Fudenberg, and Levine, 2008). There, the very desire of the DM to build a reputation results in strategic behavior which ultimately impairs welfare. The DM takes less efficient actions when he cares more about the future, as his reputation is then more salient.[24] On the contrary, the adverse welfare impact of reputation is not driven here by heightened reputational concerns: when the DM cares more about his reputation, the inefficiency actually diminishes.[25] Accordingly, the reason why reputation depresses welfare is essentially different,

[22]Notice that one might also care about dynamic efficiency, that is, how the strength of reputational concerns affects the variance of future reputations captured by the constant Σ in (15).

[23]Another illustration is the impact of the cost parameter γ: when γ tends to 0 the action becomes infinitely inefficient in the 2-period equilibrium ($k_1 \to \infty$), while k tends to $k^{FB} = 1$ in both equilibria of the stationary game.

[24]In Ely and Välimäki (2003), the no-trade result arises in the limit case where $\delta \to 1$.

[25]One may find surprising that a higher δ increases welfare after we have stressed that the DM could

and actually stems from the single-peakedness of the DM's payoff function combined with him overshooting in the bad equilibrium to build a reputation.

4 Strategic audience management

In the previous section, we have analyzed how the DM builds a reputation taking his audience as given, and shown that he obtains an expected discounted surplus in period 1 proportional to

$$1 - \mathbb{V}(b) - K(\overline{b} - m_1)^2 - K\Sigma. \tag{21}$$

This value function evidences two types of losses for the DM. First, his intervention is imperfectly tailored to the idiosyncratic bias of each receiver: this is measured by the variance of investment $\mathbb{V}(b)$. Second, there is a loss stemming from the mismatch between his initial reputation m_1 and the average bias $\overline{b}$ (his bliss reputation), which depends on the (in)efficiency of his attempts to reduce his reputational deficit, measured by K. Both these losses depend on the composition of his audience, which suggests a rationale for an organizational design optimizing audience composition both along the mean and variance dimensions. In this section, we take advantage of the simple form of the value function (21) to examine the interplay between strategic audience management and the efficiency of reputation.

4.1 Preliminary steps

Before turning to the various audience management strategies we consider, we describe how our baseline model can accommodate flexible audience composition. We first assume that the cost of a_t is proportional to the size of the audience: if the audience has a mass μ, the cost function becomes $\mu\gamma\frac{a_t^2}{2}$. This rules out technological effects driven by economies of scale and focuses the analysis on the role of reputation only. To simplify matters, we also assume that receivers' biases b are uniformly distributed on $\mathcal{B} = [-A, A]$, where A is a positive parameter.[26] Finally, we interpret audience design as a permanent organizational

be better off in the game where he behaves myopically than in the high-responsiveness equilibrium. This is due to the fact that the equilibrium payoff of the DM in the high-responsiveness equilibrium is not continuous at $\delta = 0$: $\lim_{\delta \to 0} \overline{k} = \infty$, while the unique equilibrium is $k^{static} = 0$ when $\delta = 0$.

[26]This assumption is made for analytical convenience, but is inconsequential.

or institutional decision made at $t = 1$, which, once taken, cannot be adjusted when future reputations are realized.

The contribution of receivers belonging to a subset $\mathcal{I} \subset \mathcal{B}$ to the DM's payoff in period t is

$$\int_{b\in\mathcal{I}} [1 - b + \mathbb{E}_t(x_t) - \frac{1}{2}(1 - b + \mathbb{E}_t(x_t))^2 - \gamma\frac{a_t^2}{2}]\, dF(b)$$

$$= P(\mathcal{I})\left(\frac{1}{2} - \frac{1}{2}\mathbb{V}(b|b\in\mathcal{I}) - \frac{1}{2}(m_t + a_t - \mathbb{E}(b|b\in\mathcal{I}))^2 - \gamma\frac{a_t^2}{2}\right), \tag{22}$$

where $P(\mathcal{I}) \equiv \int_{b\in\mathcal{I}} dF(b)$ denotes the mass of receivers in $\mathcal{I}$.

Note that facing a subset of the initial audience does not qualitatively change the DM's reputational incentives, and that he will follow an equilibrium strategy

$$a^*(m_t) = k\left(\mathbb{E}(b|b\in\mathcal{I}) - m_t\right),$$

where $k \in \left\{\underline{k}, \overline{k}\right\}$. We derive that the contribution of receivers in $\mathcal{I}$ to the expected discounted payoff of the DM at date 1 is

$$\Pi(\mathcal{I}) \equiv \frac{1}{2(1-\delta)}P(\mathcal{I})\left(1 - \mathbb{V}(b|b\in\mathcal{I}) - K[\mathbb{E}(b|b\in\mathcal{I}) - m_1]^2 - K\Sigma\right). \tag{23}$$

4.2 Optimal audience composition

We first study the case where the DM can choose the optimal composition of his audience, i.e., select from the segment $[-A, A]$ the set of receivers whose actions affect his payoff, or equivalently, exclude those whose actions then become irrelevant to him. For instance, an organization can set its boundaries, i.e., decide which markets to target, which activity or product to develop, or in which region to invest. Any of these decisions involves selecting agents with particular preferences to be part of the organization. That mix in turn affects the reputational incentives of the top management. Politicians also have some leeway when choosing their constituency: they can decide in which district to run, what political affiliation to carry, or the lobbies they raise money from. These choices affect the pool of stakeholders they care about when they actually take policy decisions once in office.

Formally, the DM's problem is to find the subset $\mathcal{I} \subset \mathcal{B}$ of receivers which maximizes $\Pi(\mathcal{I})$. The expression in (23) shows how tailoring $\mathcal{I}$ affects the DM's surplus. First, expanding $\mathcal{I}$ increases the mass of receivers $P(\mathcal{I})$ that affect the DM's surplus. The highest-value receiver is the one whose bias b matches the DM's reputation m_1 because his investment decision can be realigned at no cost. Expanding $\mathcal{I}$ around this point implies adding receivers whose biases ever more deviate from m_1, hence whose marginal contributions to the DM's payoff decline. This is captured by the term $\mathbb{V}(b|b \in \mathcal{I})$. In addition, expanding $\mathcal{I}$ may also affect the mismatch between the DM's reputation m_1 and the average bias in his audience $\mathbb{E}(b|b \in I)$. A more severe mismatch is costly because it requires higher efforts to realign the average investment decision in $\mathcal{I}$, and it is all the more costly as these efforts are less efficient (i.e., when K increases).

Before characterizing $\mathcal{I}^*$, let us assume that $K^{FB}\Sigma < 1$. If this does not hold, $\Pi(\mathcal{I}) < 0$ for all $\mathcal{I}$, which implies $\mathcal{I}^* = \varnothing$ (or $\mathcal{I}^*$ has zero mass) for any m_1 and K.

Proposition 5. *$\mathcal{I}^*$ is an interval. Both $P(\mathcal{I}^*)$ and $\left(\mathbb{E}(b|b \in \mathcal{I}^*) - m_1\right)^2$ decrease in K: the DM selects a wider and less congruent audience when the equilibrium is more efficient.*

Proof In the Appendix.

Intuitively, if $\mathcal{I}^*$ is not an interval, one can always reshuffle some mass from the extremes to the center, without changing neither the total mass of $\mathcal{I}^*$ nor the conditional mean in $\mathcal{I}^*$. Such a change decreases the mass of receivers with large biases whose actions are less efficient. Inspecting (23), the impact of the efficiency parameter K on the optimal audience $\mathcal{I}^*$ is twofold. First, the marginal value of an additional receiver is lower when reputation is less efficient, hence *narrower* audiences. Second, when K is higher, the mismatch between the DM and the average receiver has a stronger (negative) impact on the DM's payoff. Therefore, more inefficient equilibria are associated with *more congruent* audiences. Combining the results of Propositions 3, 4 and 5, we derive the following corollary:

Corollary 2. *In both equilibria, the audience is narrower and more congruent than in the first best. The audience is narrower and more congruent in the high-responsiveness equilibrium than in the low-responsiveness equilibrium, and than in the no-reputation case when $K > 1$.*

Corollary 2 suggests that the cost of being trapped in an inefficient equilibrium takes the form not only of costly overreactions, but also of excessively narrow and like-minded audiences. For instance, in political regimes where policy-making is less efficient, politicians should rely on smaller constituencies. In line with this prediction, dictatorial or authoritarian regimes are indeed characterized by a narrow and highly patronized base of support (Acemoglu, Robinson, and Verdier, 2004; Acemoglu and Robinson, 2008).[27]

The previous analysis makes clear that the DM increases his payoff by excluding a subset of receivers. In many applications, however, there are limits to the DM's ability to select which receivers he cares about. For instance, a planner who cares about aggregate macroeconomic performance cannot possibly ignore the impact of his policy on the behavior of some specific agents. Alternatively, the DM's audience can be interpreted as an attribute of his enduring preferences, for instance, the set of agents a policy maker intrinsically cares about. Even when this is not the case, e.g, if a politician can strategically choose to rely on particular constituencies, the office typically prescribes that he should act in the general interest of the whole electorate.[28] In organizations, technology and competition impose limits to firms' ability to design their portfolio of activities or markets.[29] From now on, we therefore rule out the possibility that the DM can exclude receivers.

4.3 Delegation

In this section, we examine how delegation can be used to redesign the audience. In order to insulate the audience-based rationale for delegation, we assume that the DM can delegate decision rights to a delegate identical in terms of prior reputation m_1 and efficiency K who can only differ from the DM through the composition of his audience.[30] In addition, to rule out that delegation be motivated by the prospect of cost optimization, we assume that the DM internalizes the cost of the delegate's intervention on the entire

[27] Interpreting h_η as the rate of renewal of politicians or ideas, and h_ε as a measure of transparency (how well political outcomes are observed), Corollary 1 clearly predicts these regimes to be inefficient.

[28] Although this requirement is not formally implementable, a long-term commitment to ignore a subset of voters would jeopardize reelection prospects.

[29] For instance, in two-sided markets (e.g., matching platforms), the existence of participation externalities between both sides of the market makes it impossible to exclude one audience while keeping the other.

[30] Delegation is obviously optimal if the delegate intrinsically does better than the DM, either because he is intrinsically more efficient (i.e., has a smaller K), or better suited to the audience (has a reputation m_1 closer to $\bar{b}$). In particular, a DM trapped in an equilibrium with bad reputation should optimally delegate control to a party with no career concerns who always implements the static action $a^{static} = 0$. See for instance Maskin and Tirole (2004).

audience.

We contrast two variants of delegation:

1. *Delegation with mandatory target:* the DM can assign a specific mission to the delegate, i.e., require him to target a specific audience.

2. *Delegation with no mandatory target:* the DM cannot assign a specific mission. The delegate composes his audience, then chooses his target following his own interest.

4.3.1 Delegation with mandatory target

Let us start with the first variant and suppose that the DM can assign a target $\tilde{b}$ to the delegate.[31] The next Proposition shows how this target should be optimally set.

Proposition 6. *The DM assigns to the delegate a target* $\tilde{b}^* = \left(1 - \frac{1}{(1+\gamma)k}\right) m_1$. *The amount of dissent* $|\tilde{b}^*|$ *increases with* K, *i.e., when reputation becomes less efficient. The benefit from delegation is nonnegative and increasing in* K.

Proof: In the Appendix.

Notice that, in the first best case where $k^{FB} = \frac{1}{1+\gamma}$, the profile of actions is efficient, and delegation brings no value: $\tilde{b}^* = \overline{b} = 0$, and the delegates would simply replicate the actions the DM himself takes. The value of delegation stems from the ability of the DM to assign to the delegate a dissenting target to correct for the fact that equilibrium actions are suboptimal. Indeed, assigning a target $\tilde{b}^*$ induces the optimal action in period 1:[32]

$$a_1^* = -\frac{1}{1+\gamma} m_1.$$

As K increases, the equilibrium actions of the delegate become less efficient, so that the level of dissent necessary to correct for this inefficiency increases. Relatedly, the DM is less efficient himself, which increases the value of delegation.

Finally, notice that, in the low-responsiveness equilibrium, as the DM underreacts ($\underline{k} < \frac{1}{1+\gamma}$), the optimal target is negatively correlated with reputation ($m_1 \tilde{b}^* < 0$). Conversely, it

[31] Since only the average bias in the audience determines the equilibrium actions, there is no loss of generality in assigning to the delegate a target $\tilde{b}$ rather than a subset of the audience. Notice that one could even in principle assign a target $\tilde{b} \notin [-A, A]$.

[32] However, actions in subsequent periods are no longer optimal, as the delegate's reputation moves away from m_1, while $\tilde{b}^*$ is set once for all.

is positively correlated with reputation in the high-responsiveness equilibrium ($m_1\tilde{b}^* > 0$) because of the DM's overreaction ($\overline{k} > \frac{1}{1+\gamma}$). It follows that a less efficient equilibrium not only results in more dissent between the delegate and the DM, but also in more congruence between the delegate and his (assigned) audience: the mismatch between the delegate and his target is smaller.

The benefit of dissent in delegation relationships is illustrated by the use of *narrow mandates*. For instance, the main stated objective of the European Central bank is price stability, even though monetary policy has broader effects on economic activity. Even beyond the focus on price stability, several central banks are assigned an explicit target for the inflation level (inflation targeting). In the same vein, public agencies are often assigned narrow and well-identified missions although the public authorities that specify their mandates have much broader objectives (Wilson, 1989).

More generally, our result sheds light on the role that audience management can play in delegation relationships when implicit contracts (e.g., career concerns) play an important role. Actually, a principal who wants to provide the right reputational incentives should strategically choose the audience relevant to the agent. For instance, a researcher cares both about his reputation within his own institution, which directly affects his career opportunities, and the external perception of his peers, which affects his outside options. Research institutions can shape reputational incentives by varying the weight assigned to each audience in researchers' evaluation process. For instance, tenure decisions may more or less rely on evaluations from outside referees, or on performance at tasks more difficult to evaluate for an external observer (e.g., administrative work, mentoring). Note that, in this case, while the principal may manipulate the agent's reputational incentives, he does not fully control them: the agent may choose to prioritize his inside or outside reputation depending on his own preferences and skills. This is the problem we now turn to.

4.3.2 Delegation with no mandatory target

As Proposition 6 shows, delegation is always profitable if the DM could assign a mandatory target to the delegate. However, there are instances in which the DM has no control over how the delegate shapes his audience. For instance, by delegating decision rights to a supranational institution that has primacy over their own authority, politicians not only relinquish political control over decisions that affect their constituents, but also become vulnerable to the risk of being under-represented in the political agenda set by the supranational authority. We examine under what conditions delegation is still profitable in such a situation. We know from Section 4.2 that if the DM were able to exclude some receivers from his audience, he would likely do so. However, the fact that the DM's objective remains fixed while the delegate can freely compose his audience creates an ex post misalignment in their preferences in spite of them being ex ante identical: the delegate's endogenous selection will lead him to ignore segments of the audience the DM does care about. Therefore, the benefits of dissent underlined in the above section have to be weighted against the cost that the extent of dissent is at the discretion of the delegate.

Proposition 7. *When the delegate can compose his audience, the DM always delegates control if reputation is bad ($K \geq 1$). He always retains control if reputation is good ($K \leq 1$) provided $|m_1| < A - \sqrt{1 - K^{FB}\Sigma}$.*

Proof: In the Appendix.

To understand this result, consider the two polar levels of efficiency $K = +\infty$ and $K = K^{FB}$. In the former case, there is no conflict between the DM and the delegate regarding their preferred targets: the DM would assign a target $\tilde{b}^* = m_1$ to the delegate, but the delegate also chooses a perfectly congruent audience himself ($\tilde{b} = m_1$) to minimize the costs of reputation.[33] Delegation is accordingly always profitable. In the latter (first best) case, delegation brings no value even when the DM can assign a specific target to the delegate. A fortiori, it is dominated when the delegate composes his audience. As K decreases between these two extremes, the benefit from being able to assign a target decreases (Proposition 6), and the misalignment between the target which the DM would assign to the delegate and the target the delegate himself chooses increases. Therefore, delegation becomes relatively less beneficial. When the initial mismatch with the average receiver in the population is not too large, delegation is optimal if and only if the equilibrium involves bad reputation. When it exists, such an equilibrium with delegation then coexists with one in which reputation is more efficient and delegation accordingly undesirable. In line with our audience-based approach of delegation, politicians that delegate

[33] See the proof of Proposition 5 in the Appendix for details.

authority to supranational institutions (e.g., central banks, the EU) often later shift the blame onto them, in particular for inadequately representing their own constituents, that is, the very rationale for delegating in the first place.[34] For instance, politicians of the euro area have consistently made public statements to deplore that the ECB does not pay enough attention to job creation.[35]

4.4 Differential treatment

As we have just seen, delegation, by changing the identity of the average receiver, improves the efficiency of the DM's intervention, hence the average efficiency of investment, but because the impact of the intervention is uniform across receivers, as in our baseline model, delegation has no impact on the variance of investment. We now consider alternative audience management strategies aiming at lowering this variance by treating receivers differentially.

4.4.1 Centralization versus decentralization

We first consider the choice between two modes of intervention that differ through their impacts on the variance of investment, but not on average investment: centralization and decentralization. Under centralization, the DM chooses an action which uniformly impacts all receivers, as in the baseline model. Under decentralization, the DM empowers local decision-makers, and asks each of them to take an action that only impacts receivers in their local environments. While the definition of centralization is clear, there are several possible ways to decentralize decision making. An extreme form is decentralization at the individual level, in which each receiver type is assigned a local decision maker. But there are also intermediate forms where local decision makers each exert influence on a group of heterogenous receivers.[36] We now investigate the optimal institutional (hierarchical)

[34]A usual interpretation of such *blame shifting* is that broad delegation mechanisms allow to enact policies in favor of special interests at the expense of the whole electorate. See for instance Schoenbrod (2008). Pei (2015) also considers a blame-shifting theory of delegation motivated by reputational concerns, but where delegation is used because it changes the way voters infer information from policies, and not because it allows to redesign the target audience.

[35]For instance, Nicolas Sarkozy declared in July 2008: *"I have the right as president of the French republic to wonder if it is reasonable to raise the European rates to 4.25 percent while the Americans have rates of 2.0 percent."*. See https://euobserver.com/economic/26451.

[36]For instance, organizations can design their hierarchical structure, and decide the allocation of managerial authority at different layers of the hierarchy. Similarly, a politician may design institutions with

structure. Formally, the DM can arbitrarily partition the support of types $\mathcal{B}$ into different segments and assign a different local decision-maker to each segment. To address this optimal segmentation problem, let us assume that all the local decision-makers are ex ante identical, in particular have same initial reputation.[37] We also assume away learning across segments, that is, neither local DMs nor receivers in a given segment can learn from the outcomes in other segments. This may be because they do not observe the payoffs of players in other segments, or because the shocks underlying these payoffs are orthogonal across segments.[38] Finally, we assume no agency friction in that each local DM simply maximizes the surplus which receivers on his own segment generate to the "central" DM (net of the reputation costs he incurs).

Proposition 8. *The optimal segmentation strategy is bang bang: the DM either chooses centralization or decentralization at the receiver level. He chooses centralization if and only if reputation is bad ($K \geq 1$).*

To understand the intuition, note that the per-period welfare can be written as

$$\frac{1-\mathbb{V}(y_t)}{2} - \frac{1}{2}\left(1-\mathbb{E}(y_t)\right)^2 - \frac{1}{2}\gamma\left([\mathbb{E}(a_t)]^2 + \mathbb{V}(a_t)\right).$$

In this expression, decentralization only affect the variance terms.[39] First, the variance of the actions $\mathbb{V}(a_t)$ is nil under centralization (there is only one action), but equal to $k^2\mathbb{V}(b)$ under decentralization. Intuitively, decentralization requires some DMs to tailor their actions to receivers with extreme biases, which is on average costlier than targeting the average receiver. Hence, the total cost incurred under decentralization inflates when there are more extreme receivers ($\mathbb{V}(b)$ is large) and when reputational incentives generate a high responsiveness (k is large). However, decentralization allows the DM (or his surrogates) to tailor their actions to the specific bias of receivers, which may improve

more or less degree of adaptation to local preferences or competences. In a decentralized structure, decision rights can besides be allocated at different levels (e.g., city, district, state).

[37] Considering DMs with different prior reputations would twist the results towards more decentralization, as it would then be possible to allocate DMs to a segment where they are a priori better aligned with receivers.

[38] Allowing cross-learning would be akin to raising the quality of monitoring h_ε under decentralized making, which given Corollary 1, always benefits the DM. It would then clearly twist the choice of the DM in favor of more decentralization.

[39] It is easy to check that decentralization leaves both the average investment $\mathbb{E}(y_t) = 1-(1-k)(\bar{b}-m_t)$ and the average action $\mathbb{E}(a_t) = k(\bar{b}-m_t)$ unchanged.

the efficiency of their investments. This is captured by the term $\mathbb{V}(y_t)$, which equals $(1-k)^2\mathbb{V}(b)$ under decentralization and $\mathbb{V}(b)$ under centralization. Intuitively, if the reactivity is not too strong, decentralization brings investment decisions on average closer to the optimum $y_t = 1$, i.e., $\mathbb{V}(y_t)$ is lower under decentralization. Whether the benefit from adaptation dominates the extra reputation costs depends on the efficiency of the equilibrium. Actually, decentralization is akin to creating new intervention channels between decision makers and receivers.[40] When reputation is good, opening new channels is efficient and decentralization accordingly dominates. As responsiveness increases and reputation gets bad, centralization becomes dominant. Moreover, not only does decentralization compound reputation costs but it may at some point push investment decisions even further away from efficiency than centralization (i.e., when $k > 2$, $\mathbb{V}(y_t)$ also becomes larger under decentralization).

The intuition is reminiscent of the merits of public versus private communication in cheap talk games with multiple audiences (Farrell and Gibbons, 1989). Centralization generates mutual discipline in that dealing with a large unique audience provides some commitment not to pander to every single receiver, which lowers reputation-building costs. However, it comes at the cost of a weaker impact, as interventions are less tailored to the idiosyncratic bias of each receiver: the relationship between the DM and a given receiver is subverted by the presence of other receivers in the audience.

Notice that because the optimal segmentation pattern is bang bang one can ignore the question of which receivers the DM pools together. This comes from the ability of the DM to partition the set of receivers in an arbitrarily fine way. Suppose instead that there is finite supply of local decision makers and that the partition cannot consist of more than N segments. In this case, the identity of receivers belonging to each segment matters when $K < 1$. When b is uniformly distributed, as we assume, it is actually simple to derive that the DM maximizes his payoff by having N intervals of equal size.

Overall, Proposition 8 suggests the coexistence of two organizational structures in equilibrium. Accordingly, similar organizations or countries could endogenously choose different hierarchical or institutional forms as a response to the different reputation costs they face. In the less efficient equilibrium, centralization is a byproduct (not a cause) of

[40] A feature of the optimal segmentation strategy is that the DM de facto chooses *how many reputations* to pursue.

inefficiencies in policy-making. In addition, since the optimal organization mode does not depend on the reputation m_t, the DM would have no incentive to change the organizational structure as his reputation changes even if he could. However, an unanticipated switch from the equilibrium with good reputation to one with bad reputation would cause the organization to change its structure and switch to centralization. In this respect, Proposition 8 provides a rationale for the alternation between centralization and decentralization in organizations, a stylized fact that has been widely documented.[41]

4.4.2 Exemption

Finally, the last audience management strategy we consider is exemption, whereby the decision maker insulates a fraction of his audience from the impact of his intervention. Exemptions (or exceptions) are widely used in fiscal policy, but the question of asymmetric treatment is also critical in some ongoing debates on regulation practices (e.g., net neutrality).[42] Under exemption, the DM strategically decides the set $\mathcal{I}$ of receivers he influences, and the set $\mathcal{E}$ of receivers he exempts. Since receivers in $\mathcal{E}$ choose $y_t(b) = 1 - b$, the DM's expected discounted payoff in period 1 equals

$$\Pi(\mathcal{I}) + \frac{1}{1-\delta}\int_{b\in\mathcal{E}} [1 - b - \frac{1}{2}(1-b)^2]\, dF(b)$$

$$\propto\ 1 - \mathbb{V}(b) + P(\mathcal{I})\left(\mathbb{E}(b|b\in\mathcal{I})^2 - K[\mathbb{E}(b|b\in\mathcal{I}) - m_1]^2 - K\Sigma\right) \tag{24}$$

Proposition 9. *In equilibrium, both $\mathcal{I}$ and $\mathcal{E}$ are intervals. In addition, $\mathcal{E} \neq \varnothing$: the DM always exempts some receivers. If $\mathcal{I} \neq \varnothing$, one has*

- $m_1 > 0 \Rightarrow A \in \mathcal{I}$ *and* $-A \in \mathcal{E}$
- $m_1 < 0 \Rightarrow -A \in \mathcal{I}$ *and* $A \in \mathcal{E}$

In words, the optimal exemption policy consists of dividing the set of receivers into two intervals and exempting those belonging to the interval which the DM is a priori

[41] See for instance Eccles, Nohria, and Berkley (1992); Nickerson and Zenger (2002). Eccles, Nohria, and Berkley (1992) describe this phenomenon as the "time-honored cycle between centralization and decentralization."

[42] For instance, while the FCC has an asymmetric treatment of fixed and mobile networks, the European Union has a uniform regulatory approach. See also Choi, Jeon, and Kim (2014), who analyze the effects of net neutrality regulation on innovation incentives of major content providers.

intrinsically less able to influence. The intuition is as follows: Extreme types generate the strongest externality, and are accordingly those the DM needs to direct his intervention at. However, it is impossible to simultaneously impact receivers at both extremes of the spectrum for if the average bias in the DM's audience is too moderate extreme receivers have no reason to expect actions which correct their incentives. In order to tilt his target audience towards extreme receivers of one side, the DM has to exempt extreme receivers of the other side. The DM therefore has to choose his side, and he optimally chooses the one he is a priori more efficient at influencing.

While the intuition for this result is extremely natural with no reputation (e.g., when the DM repeatedly plays $a^{static} = 0$), it is not completely immediate with reputational concerns. First, as under exclusion, a DM with a moderate reputation could be tempted to include receivers from both extremes so as to lower his initial reputational deficit. Second, the fact that the bad equilibrium features overreactions could lead the DM to target an audience he is a priori less efficient at influencing (reversal). Proposition 9 shows that reputation-building, no matter how inefficient, is not strong enough to qualitatively overturn the static intuition. However, as K gets too large, the DM may end up exempting the whole audience, as reputational concerns are then too costly to sustain.

More generally, it is impossible to establish generic comparative statics results on how the set of exempted receivers varies with K. Actually, the DM cares both about having congruent (to save on reputational costs) and extreme (to increase impact) receivers. These two objectives are more or less antagonistic according to the value of m_1, that is, how intrinsically moderate the DM is. However, we are able to derive a comparative statics result in the case where $m_1 = 0$, that is, when the tension between these two objectives is maximum.

Proposition 10. *Consider the case $m_1 = 0$. If $K > 1 - \frac{\Sigma}{A^2}$, then $\mathcal{I} = \varnothing$. If $K \leq 1 - \frac{\Sigma}{A^2}$, then $\mathcal{I} \neq \varnothing$, and one then has $m_1 \in \mathcal{E}$. In addition, $P(\mathcal{I})$ and $-\left(\mathbb{E}(b|b \in \mathcal{I}) - m_1\right)^2$ decrease in K : the DM selects a wider and more congruent audience when the equilibrium is more efficient.*

As under exclusion, the audience gets narrower as the inefficiency of reputation increases. In particular, when reputation is bad, all receivers are exempted, meaning that the DM does not intervene, as building a reputation would be prohibitively costly.[43] However, there is an important difference in the way the set of included receivers is constructed. Under exclusion, the DM constructs his audience starting from the receiver generating most surplus to him, that is, the one for whom correcting the bias is costless ($b = m_1$), and then expands his audience up to the point where the marginal receiver becomes too costly to include. Under exemption, including the most congruent type ($b = m_1 = 0$) has no value, as this receiver behaves in a way which maximizes the DM's payoff even without his intervention, and distracts him from his target, that is, extreme audiences. On the contrary, the DM expands the audience starting from the extreme receiver which the DM is more efficient at influencing. This explains why the endogenous mismatch between the DM and his target increases with efficiency under exclusion and decreases with efficiency under exemption.

5 Conclusion

We consider a decision maker who uses his reputation to influence an audience composed of heterogenous receivers. Reputation is horizontal in that the decision maker has a preference for being perceived as moderate, which is captured by the value of reputation being single-peaked. We show that intertemporal strategic complementarities endogenously arise that sustain two equilibria with distinct welfare properties. The DM's reaction to reputational incentives is excessive in one equilibrium, and reputation may accordingly depress welfare. A possible remedy to this inefficiency is to redesign his audience. This can be done directly by selecting receivers, which leads to narrower and more congruent audiences, or through organizational or institutional design. We show that multiple organizational forms can arise in equilibrium as a response to the costs of reputation-building. For instance, delegation or centralization arise when reputation is bad, but never when it is good. More generally, our approach suggests that strategic audience management can be a key lever in settings where reputation matters.

[43]This suggests that laissez-faire policies that are not efficient per se may arise from the inability to commit to good policies.

References

Acemoglu, D., and J. A. Robinson, 2008, "Persistence of Power, Elites, and Institutions," *American Economic Review*, 98(1), 267–293.

Acemoglu, D., J. A. Robinson, and T. Verdier, 2004, "Kleptocracy and Divide-and-Rule: A Model of Personal Rule. The Alfred Marshall Lecture.," *Journal of the European Economic Association*, 2(2/3), 162–192.

Bar-Isaac, H., and J. Deb, 2014a, "(Good and Bad) Reputation for a Servant of Two Masters," *American Economic Journal: Microeconomics*, 6(4), 293–325.

——— , 2014b, "What is a good reputation? Career concerns with heterogeneous audiences," *International Journal of Industrial Organization*, 34, 44–50.

Bar-Isaac, H., and S. Tadelis, 2008, "Seller Reputation," *Foundations and Trends in Microeconomics*, 4(4), 273–351.

Bouvard, M., and R. Levy, 2017, "Two-sided reputation in certification markets," *Management Science*.

Carrillo, J. D., and D. Gromb, 1999, "On the strength of corporate cultures," *European Economic Review*, 43(4), 1021–1037.

Choi, J. P., D.-S. Jeon, and B.-C. Kim, 2014, "Asymmetric neutrality regulation and innovation at the edges: fixed vs. mobile networks," working paper.

Cisternas, G., 2017, "Two-Sided Learning and the Ratchet Principle," *The Review of Economic Studies*.

Cukierman, A., and M. Tommasi, 1998, "When Does It Take a Nixon to Go to China?," *American Economic Review*, 88(1), 180–197.

Dessein, W., L. Garicano, and R. Gertner, 2010, "Organizing for synergies," *American Economic Journal: Microeconomics*, 2(4), 77–114.

Dewatripont, M., I. Jewitt, and J. Tirole, 1999, "The Economics of Career Concerns: Part 1," *Review of Economic Studies*, 66, 183–98.

Eccles, R., N. Nohria, and J. Berkley, 1992, "Beyond the hype: rediscovering the essence of management," *Harvard Business School Press, Boston.*

Ely, J., D. Fudenberg, and D. Levine, 2008, "When is reputation bad?," *Games and Economic Behavior*, 63(2), 498–526.

Ely, J., and J. Välimäki, 2003, "Bad Reputation," *Quarterly Journal of Economics*, 118(3), 785–814.

Farhi, E., and J. Tirole, 2012, "Collective moral hazard, maturity mismatch, and systemic bailouts," *The American Economic Review*, 102(1), 60–93.

Farrell, J., and R. Gibbons, 1989, "Cheap Talk with Two Audiences," *The American Economic Review*, 79(5), 1214–1223.

Holmström, B., 1999, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182.

Kartik, N., and R. Van Weelden, 2014, "Informative Cheap Talk in Elections," *Unpublished mansucript.*

Kreps, D. M., 1990, "Corporate culture and economic theory," in *Perspectives on Positive Political Economy*, ed. by J. E. Alt, and K. A. Shepsle. Cambridge University Press, Cambridge, pp. 91–143.

Mailath, G. J., and L. Samuelson, 2006, *Repeated games and reputations.* Oxford University Press.

Maskin, E., and J. Tirole, 2004, "The Politician and the Judge: Accountability in Government," *American Economic Review*, 94(4), 1034–1054.

Miao, J., 2014, *Economic dynamics in discrete time.* MIT press.

Morris, S., 2001, "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.

Nickerson, J. A., and T. R. Zenger, 2002, "Being efficiently fickle: A dynamic theory of organizational choice," *Organization Science*, 13(5), 547–566.

Pei, H. D., 2015, "A Blame Shifting Theory of Delegation," working paper.

Schoenbrod, D., 2008, *Power without responsibility: How Congress abuses the people through delegation*. Yale University Press.

Shapiro, J., and D. Skeie, 2015, "Information management in banking crises," *Review of Financial Studies*, 28(8), 2322–2363.

Wilson, J., 1989, *Bureaucracy: What government agencies do and why they do it*. Basic Books.

Appendix

Proof of Proposition 2 Suppose $V(m_t^{DM}, m_t)$ is quadratic, that is,

$$V(m_t^{DM}, m_t) = \alpha_1 (m_t^{DM})^2 + \alpha_2 m_t^2 + \alpha_3 m_t m_t^{DM} + \alpha_4 m_t^{DM} + \alpha_5 m_t + \alpha_6.$$

In order for the optimization problem (11) to be convex, we need to make sure that

$$2\alpha_2(1-\lambda) - \frac{\gamma}{\delta(1-\lambda)} < 0, \tag{25}$$

which will be checked *ex post* to be verified. The first-order condition writes

$$\begin{aligned} & \delta(1-\lambda)[2\alpha_2 \mathbb{E}(m_{t+1}) + \alpha_3 \mathbb{E}(m_{t+1}^{DM}) + \alpha_5] = \gamma a_t \\ \Leftrightarrow \quad & 2\alpha_2 \left\{ \lambda m_t + (1-\lambda)[m_t^{DM} + a_t - a_t^e] \right\} + \alpha_3 m_t^{DM} + \alpha_5 = \tfrac{\gamma}{\delta(1-\lambda)} a_t \end{aligned}$$

In order to satisfy the equilibrium conditions (11) and (12), the following condition must hold for any pair (m_t^{DM}, m_t):

$$2\alpha_2 \left\{ \lambda m_t + (1-\lambda)[m_t^{DM} - a(m_t, m_t)] \right\} + \alpha_3 m_t^{DM} + \alpha_5 = \left[\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda) \right] a_t(m_t^{DM}, m_t). \tag{26}$$

Given $V(.,.)$, there exists a unique linear strategy, $a_t(m_t^{DM}, m_t) = \beta_1 m_t^{DM} + \beta_2 m_t + \beta_3$, which satisfies (26). For all (m_t^{DM}, m_t), $(\beta_1, \beta_2, \beta_3)$ must satisfy

$$2\alpha_2 \left\{ [\lambda - (1-\lambda)(\beta_1+\beta_2)] m_t + (1-\lambda) m_t^{DM} - (1-\lambda)\beta_3 \right\} + \alpha_3 m_t^{DM} + \alpha_5$$

$$= \left[\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda)\right](\beta_1 m_t^{DM} + \beta_2 m_t + \beta_3).$$

This gives

$$\beta_1 = \frac{\alpha_3 + 2\alpha_2(1-\lambda)}{\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda)}. \tag{27}$$

$$\beta_2 = \frac{\delta(1-\lambda)}{\gamma}(2\alpha_2 + \alpha_3) - \beta_1 = \frac{\delta(1-\lambda)}{\gamma}2\alpha_2[\lambda - (1-\lambda)\beta_1] \tag{28}$$

$$\beta_3 = \frac{\delta(1-\lambda)}{\gamma}\alpha_5 \tag{29}$$

Note that, from $\lambda = \dfrac{h}{h+h_\varepsilon}$, and $h = \frac{\sqrt{h_\varepsilon^2 + 4h_\eta h_\varepsilon} - h_\varepsilon}{2}$, we derive $h = (1-\lambda)^2 h_\eta$ and $h_\varepsilon = \frac{1-\lambda}{\lambda}h_\eta$. This implies

$$\mathbb{V}(m_{t+1}^{DM}) = \mathbb{V}(m_{t+1}) = (1-\lambda)^2\mathbb{V}(\theta_t + \varepsilon_t) = (1-\lambda)^2\left(\frac{1}{h} + \frac{1}{h_\varepsilon}\right) = \frac{1}{h_\eta}.$$

We will also make use of the following expectations, derived using (9) and (10), where (9) is rewritten as $m_{t+1} = \lambda m_t + (1-\lambda)\left[\theta_t + \varepsilon_t + \beta_1(m_t^{DM} - m_t)\right]$.

$$\begin{aligned}
\mathbb{E}(m_{t+1}^{DM}) &= m_t^{DM} \\
\mathbb{E}(m_{t+1}) &= [\lambda - (1-\lambda)\beta_1]m_t + (1-\lambda)(1+\beta_1)m_t^{DM} \\
\mathbb{E}[(m_{t+1}^{DM})^2] &= (m_t^{DM})^2 + \frac{1}{h_\eta} \\
\mathbb{E}(m_{t+1}^2) &= [\lambda - (1-\lambda)\beta_1]^2 m_t^2 + (1-\lambda)^2(1+\beta_1)^2(m_t^{DM})^2 \\
&\quad +2[\lambda - (1-\lambda)\beta_1](1-\lambda)(1+\beta_1)m_t^{DM}m_t + \frac{1}{h_\eta} \\
\mathbb{E}(m_{t+1}^{DM}m_{t+1}) &= (1-\lambda)(1+\beta_1)(m_t^{DM})^2 + [\lambda - (1-\lambda)\beta_1]m_t^{DM}m_t + \frac{1}{h_\eta}
\end{aligned}$$

Since all the previous terms are quadratic in (m_t^{DM}, m_t), $\pi(.)$ and the cost of a_t are also quadratic, and a_t is linear in (m_t^{DM}, m_t), we derive that

$$\pi[m_t + a_t(m_t, m_t)] + \delta\mathbb{E}V\{m_{t+1}^{DM}, m_{t+1}[a_t(m_t^{DM}, m_t), a_t(m_t, m_t)]\} - \gamma\frac{a_t(m_t^{DM}, m_t)^2}{2}$$

is quadratic in (m_t^{DM}, m_t).

In order to identify the coefficients, we first write:

$$\begin{aligned}\pi[m_t + a_t(m_t, m_t)] &= \frac{1-\mathbb{V}(b)}{2} - \frac{1}{2}[(1+\beta_1+\beta_2)m_t + \beta_3 - \bar{b}]^2 \\ &= \frac{1-\mathbb{V}(b)}{2} - \frac{1}{2}(1+\beta_1+\beta_2)^2 m_t^2 + (\bar{b}-\beta_3)(1+\beta_1+\beta_2)m_t - \frac{1}{2}(\bar{b}-\beta_3)^2\end{aligned}$$

$$\begin{aligned}\gamma\frac{a_t^2(m_t^{DM}, m_t)}{2} &= \gamma\frac{(\beta_1 m_t^{DM} + \beta_2 m_t + \beta_3)^2}{2} \\ &= \frac{\gamma}{2}[\beta_1^2(m_t^{DM})^2 + \beta_2^2 m_t^2 + 2\beta_1\beta_2 m_t^{DM} m_t + 2\beta_1\beta_3 m_t^{DM} + 2\beta_2\beta_3 m_t + \beta_3^2]\end{aligned}$$

We can now identify coefficients, using all the previous equations:

$$\alpha_1 = \delta[\alpha_1 + \alpha_2(1-\lambda)^2(1+\beta_1)^2 + \alpha_3(1-\lambda)(1+\beta_1)] - \frac{\gamma}{2}\beta_1^2 \tag{30}$$

$$\alpha_2 = -\frac{1}{2}(1+\beta_1+\beta_2)^2 + \delta\alpha_2[\lambda - (1-\lambda)\beta_1]^2 - \frac{\gamma}{2}\beta_2^2 \tag{31}$$

$$\alpha_3 = \delta\{2\alpha_2[\lambda-(1-\lambda)\beta_1](1-\lambda)(1+\beta_1) + \alpha_3[\lambda-(1-\lambda)\beta_1]\} - \gamma\beta_1\beta_2 \tag{32}$$

$$\alpha_4 = \delta\alpha_4 + \delta\alpha_5(1-\lambda)(1+\beta_1) - \gamma\beta_1\beta_3 \tag{33}$$

$$\alpha_5 = (\bar{b}-\beta_3)(1+\beta_1+\beta_2) + \delta\alpha_5[\lambda-(1-\lambda)\beta_1] - \gamma\beta_2\beta_3 \tag{34}$$

$$\alpha_6 = \frac{1-\mathbb{V}(b)}{2} - \frac{1}{2}(\bar{b}-\beta_3)^2 + \delta\left[(\alpha_1+\alpha_2+\alpha_3)\frac{1}{h_\eta} + \alpha_6\right] - \frac{\gamma}{2}\beta_3^2 \tag{35}$$

Notice also that the relations in (27) and (28) can be rewritten as

$$\alpha_2 = \frac{\frac{\gamma}{\delta(1-\lambda)}\beta_2}{2[\lambda - \beta_1(1-\lambda)]} \qquad \text{and} \qquad \alpha_3 = \frac{\gamma}{\delta(1-\lambda)}\left[\beta_1 - \beta_2\frac{(1-\lambda)(1+\beta_1)}{\lambda-\beta_1(1-\lambda)}\right] \tag{36}$$

Using (36) to substitute α_2 and α_3 in the RHS of (31) and (32),

$$\frac{1}{\gamma}\alpha_2 = -\frac{1}{2\gamma}(1+\beta_1+\beta_2)^2 + \frac{1}{2}\beta_2\left[\frac{\lambda}{1-\lambda} - \beta_1\right] - \frac{1}{2}\beta_2^2 \tag{37}$$

$$\frac{1}{\gamma}\alpha_3 = \frac{\lambda}{(1-\lambda)}\beta_1 - \beta_1^2 - \beta_2\beta_1 \tag{38}$$

$2 \times (37) + (38)$ yields, using (28),

$$\begin{aligned}\frac{1}{\delta(1-\lambda)}(\beta_1+\beta_2) &= -\frac{1}{\gamma}(1+\beta_1+\beta_2)^2+\beta_2\left[\frac{\lambda}{1-\lambda}-\beta_1\right]-\beta_2^2+\frac{\lambda}{(1-\lambda)}\beta_1-\beta_1^2-\beta_2\beta_1\\ &= -\frac{1}{\gamma}(1+\beta_1+\beta_2)^2+\frac{\lambda}{1-\lambda}(\beta_1+\beta_2)-(\beta_1+\beta_2)^2\end{aligned}$$

Let $k \equiv -(\beta_1+\beta_2)$.

$$\begin{aligned}-\frac{1}{\delta(1-\lambda)}k &= -\frac{1}{\gamma}(1-k)^2-\frac{\lambda}{(1-\lambda)}k-k^2\\ \Leftrightarrow \quad \varphi(k) &\equiv (1+\gamma)\,k^2-\left[\frac{\gamma(1-\delta\lambda)}{\delta(1-\lambda)}+2\right]k+1=0.\end{aligned} \tag{39}$$

It is easy to see that φ is convex in k. In addition, denoting $z \equiv \frac{1-\delta\lambda}{\delta(1-\lambda)} \geq 1$, one remarks $\varphi(0)>0, \varphi'(0)<0$, $\varphi(z)\geq 0, \varphi'(z)\geq 0$ $\varphi(\frac{1}{1+\gamma})\leq 0, \varphi'(\frac{1}{1+\gamma})\leq 0$, and $\varphi(1)\leq 0$.

This implies that (39) admits two solutions $\underline{k}$ and $\overline{k}$ such that

$$0 \leq \underline{k} \leq \frac{1}{1+\gamma} \leq 1 \leq \overline{k} \leq \frac{1-\delta\lambda}{\delta(1-\lambda)}$$

Let us now check that there exist β_1 and β_2 solutions to (37) and (38). Rearranging (37),

$$\begin{aligned}&\frac{\beta_2}{\delta(1-\lambda)[\lambda-\beta_1(1-\lambda)]} = -\frac{1}{\gamma}(1+\beta_1+\beta_2)^2+\beta_2\left[\frac{\lambda}{1-\lambda}-\beta_1\right]-\beta_2^2\\ \Leftrightarrow \quad &\frac{\beta_2}{\delta(1-\lambda)^2} = \left[-\frac{1}{\gamma}(1-k)^2+\beta_2\left(\frac{\lambda}{1-\lambda}+k\right)\right]\left(\frac{\lambda}{1-\lambda}+k+\beta_2\right)\\ \Leftrightarrow \quad &\left(\frac{\lambda}{1-\lambda}+k\right)\beta_2^2+\left[-\frac{1}{\gamma}(1-k)^2+\left(\frac{\lambda}{1-\lambda}+k\right)^2-\frac{1}{\delta(1-\lambda)^2}\right]\beta_2\\ &-\frac{1}{\gamma}(1-k)^2\left(\frac{\lambda}{1-\lambda}+k\right)=0\end{aligned}$$

Letting $G(\beta_2)$ denote the polynomial in the last line and remembering that $k>0$, we derive that $G(.)$ has two roots of opposite signs. Consider the positive root first. $k>0$ and $\beta_2>0$ implies $\beta_1<0$. Using this and (36), (25) is equivalent to $-(1-\lambda)k<\lambda$ which is always true. Turn now to the negative root of $G(.)$. $G[-k-\lambda/(1-\lambda)]>0$ implies $-k-\lambda/(1-\lambda)<\beta_2$ and therefore $\lambda-(1-\lambda)\beta_1>0$. This implies in turn, from (36), that $\alpha_2<0$ so that (25) holds. In conclusion, for any k solution to (39) there exist two pairs (β_1, β_2), such that (25), (36), (37) and (38) hold.

In order to fully characterize equilibrium strategies, it only remains to derive β_3. From (34),

$$\begin{aligned}
\alpha_5\{1-\delta[\lambda-(1-\lambda)\beta_1]\} &= (\bar{b}-\beta_3)(1+\beta_1+\beta_2)-\gamma\beta_2\beta_3 \\
\Leftrightarrow \frac{\gamma-\delta\gamma[\lambda-(1-\lambda)\beta_1]}{\delta(1-\lambda)}\beta_3 &= (\bar{b}-\beta_3)(1+\beta_1+\beta_2)-\gamma\beta_2\beta_3 \\
\Leftrightarrow \left[\gamma\frac{1-\delta\lambda}{\delta(1-\lambda)}+1-(1+\gamma)k\right]\beta_3 &= \bar{b}(1-k) \\
\Leftrightarrow \beta_3 &= k\bar{b}
\end{aligned}$$

where the last equality makes uses of (39).

We therefore conclude that the strategy $a_t(m_t, m_t) = a_t^*(m_t) = k(\bar{b} - m_t)$, where $k \in \{\underline{k}, \overline{k}\}$, is an equilibrium strategy provided that it satisfies the transversality condition, which we check below (see separate proof).

When $a_t^*(m_t) = k(\bar{b} - m_t)$, the payoff of the DM in each period t reads

$$\begin{aligned}
&\pi[m_t + a_t^*(m_t)] - \gamma\frac{a_t^*(m_t)^2}{2} \\
&= \frac{1-\mathbb{V}(b)}{2} - \frac{1}{2}[m_t + k(\bar{b}-m_t) - \bar{b}]^2 - \gamma\frac{k^2(\bar{b}-m_t)^2}{2} \\
&= \frac{1-\mathbb{V}(b)}{2} - \frac{1}{2}(\bar{b}-m_t)^2\left[(k-1)^2+\gamma k^2\right]
\end{aligned}$$

Therefore, one derives the expected discounted payoff the DM date t in an equilibrium k :

$$V^k(m_t) = -\frac{1}{2}\left[(k-1)^2+\gamma k^2\right]\sum_{i=t}^{+\infty}\delta^{i-t}\mathbb{E}_t(\bar{b}-m_i)^2 + \frac{1-\mathbb{V}(b)}{2(1-\delta)}.$$

Using $K = (k-1)^2 + \gamma k^2$, one derives

$$V^k(m_t) = \frac{1-\mathbb{V}(b)}{2(1-\delta)} - \frac{1}{2}K\sum_{s=0}^{+\infty}\delta^s\left(\mathbb{E}_t(\bar{b}-m_{t+s})\right)^2 - \frac{1}{2}K\sum_{s=0}^{+\infty}\delta^s\mathbb{V}_t(\bar{b}-m_{t+s}),$$

where $\mathbb{E}_t$ and $\mathbb{V}_t$ refer to the expectation and variance of m_{t+s} viewed from period t.

One easily shows by induction that, for all $s \geq 1$,

$$m_{t+s} = \lambda^s m_t + (1-\lambda)\sum_{i=0}^{s-1} \lambda^{s-1-i}(\theta_{t+i} + \epsilon_{t+i}) \tag{40}$$

in equilibrium.

It is clear that $\mathbb{E}_t(m_{t+s}) = m_t$ (martingale property). In addition, we derive that

$$\mathbb{V}_t(m_{t+s}) = (1-\lambda)^2 \left\{ \frac{1}{h}\left(\sum_{i=0}^{s-1}\lambda^i\right)^2 + \frac{1}{h_\varepsilon}\sum_{i=0}^{s-1}\lambda^{2i} + \frac{1}{h_\eta}\sum_{i=1}^{s-1}\left(\sum_{j=0}^{i-1}\lambda^j\right)^2 \right\} \tag{41}$$

Recalling $h = (1-\lambda)h_\eta$ and $h_\varepsilon = \frac{1-\lambda}{\lambda}h_\eta$, and using simple algebra, one can simplify (41) as

$$\mathbb{V}_t(m_{t+s}) = \frac{s}{h_\eta} \text{ for all } s \geq 0. \tag{42}$$

This implies that

$$\sum_{s=0}^{+\infty} \delta^s \mathbb{V}_t(\bar{b} - m_{t+s}) = \frac{\delta}{(1-\delta)^2}\frac{1}{h_\eta} \tag{43}$$

Finally, denoting $\Sigma = (1-\delta)\sum_{s=0}^{+\infty}\delta^s \mathbb{V}_t(\bar{b} - m_{t+s}) = \frac{\delta}{1-\delta}\frac{1}{h_\eta}$, we derive

$$V^k(m_t) = \frac{1}{2(1-\delta)}\left(1 - \mathbb{V}(b) - K(\bar{b} - m_t)^2 - K\Sigma\right).$$

Finally, one has

$$\begin{aligned}
(\overline{\kappa}-1)^2 + \gamma\overline{\kappa}^2 - (\underline{\kappa}-1)^2 - \gamma\underline{\kappa}^2 &= (\overline{\kappa}-\underline{\kappa})\left[(\overline{\kappa}+\underline{\kappa})(1+\gamma) - 2\right] \\
&= \gamma\frac{1-\delta\lambda}{\delta(1-\lambda)}(\overline{\kappa}-\underline{\kappa}) \\
&> 0,
\end{aligned}$$

using (39).

It is then immediate that $V^{\underline{k}}(m_t) \geq V^{\overline{k}}(m_t)$ for any m_t.

Transversality Condition (TC) We now check that our equilibria satisfy the transversality condition. Our proof uses Theorem 7.1.2 in Miao (2014): i) we show that TC holds for the equilibrium strategy ; ii) we show that any admissible strategy either satisfies TC or is dominated by a strategy that satisfies TC, namely the equilibrium strategy.

i) The equilibrium strategy satisfies TC

We want to show that for any couple m_1 and m_1^{DM}

$$\lim_{t\to+\infty}\delta^t\mathbb{E}_1[V^k(m_t^{DM},m_t)|m_1,m_1^{DM},a^*(.,.)]=0. \tag{44}$$

In words, the discounted sum of expected payoffs in period 1 given that DM plays the equilibrium strategy tends to 0 (on or off the equilibrium path).

$\mathbb{E}_1(m_{t+1}^{DM})=m_1^{DM}$ and we know from (42)

$$\mathbb{V}_1(m_{t+1}^{DM})=\frac{t}{h_\eta},$$

which implies

$$\mathbb{E}_1[(m_{t+1}^{DM})^2]=\frac{t}{h_\eta}+(m_1^{DM})^2.$$

It follows that

$$\lim_{t\to+\infty}\delta^t\mathbb{E}_1(m_{t+1}^{DM})=\lim_{t\to+\infty}\delta^t\mathbb{E}_1[(m_{t+1}^{DM})^2]=\lim_{t\to+\infty}\delta^t\mathbb{V}_1[(m_{t+1}^{DM})^2]=0.$$

Since DM plays the equilibrium strategy,

$$\begin{aligned}
m_{t+1}-m_{t+1}^{DM} &= \lambda\left(m_t-m_t^{DM}\right)+(1-\lambda)(a_t-a_t^e)\\
&= \lambda\left(m_t-m_t^{DM}\right)+(1-\lambda)(\beta_1 m_t^{DM}+\beta_2 m_t+\beta_3-\beta_1 m_t-\beta_2 m_t-\beta_3)\\
&= [\lambda-(1-\lambda)\beta_1]\left(m_t-m_t^{DM}\right)
\end{aligned}$$

which implies

$$m_{t+1}=m_{t+1}^{DM}+[\lambda-(1-\lambda)\beta_1]^t\left(m_1-m_1^{DM}\right) \tag{45}$$

Hence,

$$
\begin{aligned}
\mathbb{E}_1(m_{t+1}^2) &= \mathbb{V}_1(m_{t+1}^{DM}) + 2m_1^{DM}[\lambda - (1-\lambda)\beta_1]^t \left(m_1 - m_1^{DM}\right) \\
&\quad + [\lambda - (1-\lambda)\beta_1]^{2t} \left(m_1 - m_1^{DM}\right)^2 \\
&\quad + Constant \\
\mathbb{E}_1(m_{t+1} m_{t+1}^{DM}) &= \mathbb{V}_1(m_{t+1}^{DM}) + m_1^{DM}[\lambda - (1-\lambda)\beta_1]^t \left(m_1 - m_1^{DM}\right) + Constant \\
\mathbb{E}_1(m_{t+1}) &= [\lambda - (1-\lambda)\beta_1]^t \left(m_1 - m_1^{DM}\right) + Constant
\end{aligned}
$$

Therefore a necessary and sufficient condition for (44) to hold is

$$
\delta[\lambda - (1-\lambda)\beta_1]^2 < 1 \tag{46}
$$

We have shown that for $k \in \{\overline{k}, \underline{k}\}$, there exists a solution to equations (27) to (35) such that $\alpha_2 < 0$. Then (31) can be rewritten as

$$
\alpha_2[1 - \delta\alpha_2[\lambda - (1-\lambda)\beta_1]^2] = -\frac{1}{2}(1 + \beta_1 + \beta_2)^2 - \frac{\gamma}{2}\beta_2^2
$$

which implies

$$
1 - \delta\alpha_2[\lambda - (1-\lambda)\beta_1]^2 > 0.
$$

Hence (46) is true and the transversality condition is verified for the equilibrium strategy. This shows in particular that the DM's value function coincide with the discounted sum of his expected payoffs from playing the equilibrium strategy ("no bubble").

ii) Any admissible strategy satisfies TC or is dominated by the equilibrium strategy.

Consider the equilibrium associated with multiplier k and let $m_t(a)$ denote the DM's public reputation when agents believe that the DM follows the equilibrium strategy, but DM follows strategy $a \equiv \{a_t\}_{t>0}$ instead.

An adapted strategy $a \equiv \{a_t\}_{t>0}$ is admissible if

$$
J(a) \equiv \sum_{t=1}^{+\infty} \delta^t \mathbb{E}_1 \left\{ \frac{1 - \mathbb{V}(b)}{2} - \frac{[m_t(a) - \overline{b}]^2}{2} - \gamma \frac{a_t^2}{2} \right\}
$$

exists, i.e., is either a finite number or ∞. $J(a)$ represents the DM's expected utility from deviating from the equilibrium strategy to a when receivers believe that he follows the equilibrium strategy.

Note first that $J(a)$ is bounded above, so that if $J(a)$ is not finite, then $J(a) = -\infty$ which is dominated by the equilibrium strategy. Hence, we can restrict attention to strategies a such that $J(a)$ is finite. It follows that

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 \left\{ \frac{1-\mathbb{V}(b)}{2} - \frac{[m_t(a)-\overline{b}]^2}{2} - \gamma \frac{a_t^2}{2} \right\} = 0. \tag{47}$$

Since

$$\lim_{t\to+\infty} \delta^t \frac{1-\mathbb{V}(b)}{2} = 0,$$

and

$$\frac{[m_t(a)-\overline{b}]^2}{2} > 0 \qquad \text{and} \qquad \gamma \frac{a_t^2}{2} > 0,$$

(47) implies

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 [m_t(a)-\overline{b}]^2 = 0 \tag{48}$$

Using $|m_t(a)-\overline{b}| < 1 + [m_t(a)-\overline{b}]^2$,

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 |m_t(a)-\overline{b}| \leq \lim_{t\to+\infty} \delta^t \mathbb{E}_1 \{1 + [m_t(a)-\overline{b}]^2\} = 0$$

Hence,

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 |m_t(a)| < \lim_{t\to+\infty} \delta^t \mathbb{E}_1 \left\{ |m_t(a)-\overline{b}| + |\overline{b}| \right\} = 0$$

and therefore,

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 m_t(a) = 0. \tag{49}$$

Hence (48) that can be written as

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 [\overline{b}^2 - 2\overline{b} m_t(a) + m_t(a)^2] = 0$$

implies

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 m_t^2(a) = 0. \tag{50}$$

Finally, using Cauchy-Schwartz inequality

$$\left\{\delta^t \mathbb{E}_1[m_t(a) m_t^{DM}]\right\}^2 \leq \delta^t \mathbb{E}_1(m_t^2(a)) \times \delta^t \mathbb{E}_1[(m_t^{DM})^2] = \delta^t \mathbb{E}_1(m_t^2(a)) \times \delta^t [\frac{t}{h_\eta} + m_0^2].$$

It follows that

$$\lim_{t \to +\infty} \delta^t \mathbb{E}_1[m_t(a) m_t^{DM}] = 0 \tag{51}$$

Combining (49), (50) and (51)

$$\begin{aligned} & \lim_{t \to +\infty} \delta^t \mathbb{E}_1 V^k [m_t(a), m_t^{DM})] \\ = & \lim_{t \to +\infty} \delta^t \mathbb{E}_1 \left[\alpha_1 (m_t^{DM})^2 + \alpha_2 m_t^2(a) + \alpha_3 m_t(a) m_t^{DM} + \alpha_4 m_t^{DM} + \alpha_5 m_t(a) + \alpha_6\right] \\ = & \ 0 \end{aligned}$$

□

Proof of Proposition 3 Extending the notation, let us denote by V^0 the expected discounted payoff of the DM in the infinitely repeated static game. Since the DM then chooses $a^{static} = 0$ in each period, this payoff corresponds to the value function V^k taken for $k = 0$, that is, $K = 1$:

$$V^0(m_t) = \frac{1}{2(1-\delta)} \left(1 - \mathbb{V}(b) - (\bar{b} - m_t)^2 - \Sigma\right) \tag{52}$$

Since the path of m_t does not depend on k, it is easy to compare the equilibrium payoff of the DM in any equilibrium to his payoff in the infinitely repeated stage game.

$V^k - V^0$ has the sign of $1 - (k-1)^2 - \gamma k^2 = -k[(1+\gamma)k - 2]$.

Since $\underline{k} \leq \frac{1}{1+\gamma} < \frac{2}{1+\gamma}$, one always has $V^{\underline{k}} \geq V^0$.

It is easy to check that $\overline{k} < \frac{2}{1+\gamma} \Leftrightarrow \varphi(\frac{2}{1+\gamma}) > 0$ and $\varphi'(\frac{2}{1+\gamma}) > 0 \Leftrightarrow \gamma < \frac{\delta(1-\lambda)}{(1-\delta\lambda)+(1-\delta)}$.

Therefore, we conclude $V^{\overline{k}} > V^0 \Leftrightarrow \gamma < \frac{\delta(1-\lambda)}{(1-\delta\lambda)+(1-\delta)}$.

Proof of Proposition 4 Recalling $z = \frac{1-\delta\lambda}{\delta(1-\lambda)}$, one rewrites (39) as

$$\tilde{\varphi}(k, z) = (1+\gamma)k^2 - (2 + \gamma z)k + 1 = 0 \tag{39'}$$

It is easy to see that $\tilde{\varphi}_z \leq 0$. In addition, $\tilde{\varphi}_k(\underline{k}, z) < 0$ and $\tilde{\varphi}_k(\overline{k}, z) > 0$.

Using (8), λ increases in h_η and decreases in h_ε. Since z decreases in δ and increases in λ, we derive, using the implicit function theorem:

$\frac{\partial \underline{k}}{\partial \delta} \geq 0$, $\frac{\partial \overline{k}}{\partial \delta} \leq 0$, $\frac{\partial \underline{k}}{\partial h_\eta} \leq 0$, $\frac{\partial \overline{k}}{\partial h_\eta} \geq 0$, $\frac{\partial \underline{k}}{\partial h_\varepsilon} \geq 0$, $\frac{\partial \overline{k}}{\partial h_\varepsilon} \leq 0$.

Proof of Proposition 5 Let us first prove that $\mathcal{I}^*$ is an interval. Suppose that $\mathcal{I}^*$ has positive mass but is not convex. Consider the alternative interval

$$\mathcal{I}' \equiv (\mathbb{E}(b|b \in \mathcal{I}^*) - \frac{P(\mathcal{I}^*)}{2}, \mathbb{E}(b|b \in \mathcal{I}^*) + \frac{P(\mathcal{I}^*)}{2}).$$

By construction, $\mathbb{E}(b|b \in \mathcal{I}') = \mathbb{E}(b|b \in \mathcal{I}^*)$ and $P(\mathcal{I}') = P(\mathcal{I}^*)$, but $\mathbb{V}(b|b \in \mathcal{I}') < \mathbb{V}(b|b \in \mathcal{I}^*)$. This implies $\Pi(\mathcal{I}') - \Pi(\mathcal{I}^*) = P(\mathcal{I}^*)\left[\mathbb{V}(b|b \in \mathcal{I}^*) - \mathbb{V}(b|b \in \mathcal{I}')\right] > 0$. Therefore, the DM is strictly better off choosing $\mathcal{I}'$ rather than $\mathcal{I}^*$, and $\mathcal{I}^*$ cannot be the solution of the DM's problem.

We can then write $\mathcal{I}^* = [\underline{a}, \overline{a}]$ and

$$\begin{aligned} P(\mathcal{I}^*) &= \frac{\overline{a} - \underline{a}}{2A} \\ \mathbb{E}(b|b \in \mathcal{I}^*) &= \frac{\overline{a} + \underline{a}}{2} \\ \mathbb{V}(b|b \in \mathcal{I}^*) &= \frac{(\overline{a} - \underline{a})^2}{12} \end{aligned}$$

For convenience of notation, let us write $P \equiv \frac{\overline{a}-\underline{a}}{2A}$ and $\tilde{b} \equiv \frac{\overline{a}+\underline{a}}{2}$. One remarks that $\mathbb{V}(b|b \in \mathcal{I}^*) = \frac{(\overline{a}-\underline{a})^2}{12} = \frac{A^2P^2}{3}$. Let us also denote $\rho(K) \equiv 1 - K\Sigma$.

Instead of maximizing over $\overline{a}$ and $\underline{a}$, one may equivalently maximize (23) over P and $\tilde{b}$:

$$\max_{P \in [0,1],\, \tilde{b} \in [-A(1-P), A(1-P)]} P\left(\rho(K) - \frac{A^2P^2}{3} - K(\tilde{b} - m_1)^2\right) \quad (53)$$

Let us first fix $P \in [0,1]$ and maximize (53) w.r.t. $\tilde{b} \in [-A(1-P), A(1-P)]$.

This gives:

$\tilde{b} = m_1$ if $-A(1-P) \leq m_1 \leq A(1-P)$

$\tilde{b} = A(1-P)$ if $m_1 > A(1-P)$

$\tilde{b} = -A(1-P)$ if $m_1 < -A(1-P)$.

There are four cases:

- $m_1 > A$: then, for all $P \in [0,1]$, we have $A(1-P) < m_1$, so $\tilde{b} = A(1-P)$
- $m_1 \in [0, A]$: Then $\tilde{b} = m_1$ if $0 \leq P \leq 1 - \frac{m_1}{A}$ and $\tilde{b} = A(1-P)$ if $1 - \frac{m_1}{A} \leq P \leq 1$
- $m_1 \in [-A, 0]$: Then $\tilde{b} = m_1$ if $0 \leq P \leq 1 + \frac{m_1}{A}$ and $\tilde{b} = -A(1-P)$ if $1 + \frac{m_1}{A} \leq P \leq 1$
- $m_1 < A$: then, for all $P \in [0,1]$, we have $m_1 < -A(1-P)$, so $\tilde{b} = -A(1-P)$

One remarks that as long as $m_1 \in \mathcal{B}$, then $\mathcal{I}^*$ must include m_1.

We now maximize over P. Let us focus on the first two cases (the other two are symmetric), and start with the case $m_1 > A$.

One then maximizes $g(P) \equiv \rho(K)P - \frac{A^2P^3}{3} - K[A(1-P) - m_1]^2 P$ on $[0,1]$.

It is easy to check that g is concave on $[0,1]$ when $m_1 > A$, and that g decreases in K. We conclude that the solution P^* is an nonincreasing function of K.

Notice that the DM chooses not to participate if $P^* = 0$, which happens when $g'(0) < 0 \Leftrightarrow \rho(K) - K(A - m_1)^2 < 0$, i.e., when K is large enough, or the DM is too far away from even the closest agent in the potential audience. If $P^* > 0$, we have $\tilde{b} = A(1 - P^*) < m_1$, which implies that $|\tilde{b} - m_1| = m_1 - A(1 + P^*)$ decreases in K.

Let us now consider the case $m_1 \in [0, A]$. Let $h(P) \equiv \rho(K)P - \frac{A^2P^3}{3}$.

One maximizes a function equal to $h(P)$ on $[0, 1 - \frac{m_1}{A}]$ and $g(P)$ on $[1 - \frac{m_1}{A}, 1]$.

h is concave on $[0,1]$, and nonincreasing in K. g is concave on $[1 - \frac{m_1}{A}, 1]$ when $m_1 \in [0, A]$. It is also easy to see that $g'(P) \leq h'(P)$ on $[1 - \frac{m_1}{A}, 1]$, with equality at $1 - \frac{m_1}{A}$, and $g'(1 - \frac{m_1}{A}) = h'(1 - \frac{m_1}{A})$.

We conclude that the solution of the problem is

- $P^* = 0$ if $h'(0) < 0 \Leftrightarrow \rho(K) < 0$
- $P^* \in [0, 1 - \frac{m_1}{A}]$ if $h'(1 - \frac{m_1}{A}) \leq 0 \leq h'(0) \Leftrightarrow 0 < \rho(K) < (A - m_1)^2$
- $P^* \in [1 - \frac{m_1}{A}, 1]$ if $g'(1) \leq 0 \leq g'(1 - \frac{m_1}{A}) \Leftrightarrow (A - m_1)^2 < \rho(K) < A^2 + Km_1^2 + 2KAm_1$
- $P^* = 1$ if $0 < g'(1) \Leftrightarrow \rho(K) > A^2 + Km_1^2 + 2KAm_1$

From the fact that $\rho(K)$ decreases in K, it is easy to conclude in any case that P^* is nonincreasing in K. In addition, one has $\tilde{b} = m_1$ as long as $P^* \leq 1 - \frac{m_1}{A}$, and $|\tilde{b} - m_1|$ decreasing in K otherwise, for the same reason as in the case $m_1 > A$. We can conclude that $|\tilde{b} - m_1|$ is nonincreasing in K. $\square$

Proof of Proposition 6 Given that the delegate optimally chooses $a_t^* = k(\tilde{b} - m_t)$, and since $\overline{b} = 0$, the DM's welfare in period t reads:[44]

$$\frac{1}{2} - \frac{1}{2}\mathbb{V}(b) - \frac{1}{2}(m_t + k(\tilde{b} - m_t))^2 - \frac{1}{2}\gamma k^2(\tilde{b} - m_t)^2$$

The present value in period 1 is therefore (proportional to)

$$1 - \mathbb{V}(b) - Km_1^2 - K\Sigma - (1+\gamma)k^2\tilde{b}^2 - 2\left(k(1-k) - \gamma k^2\right)\tilde{b}m_1$$

This function reaches a maximum at

$$\tilde{b}^* = \left(1 - \frac{1}{(1+\gamma)k}\right) m_1$$

Therefore, the net benefit from delegation is (proportional to)

$$-(1+\gamma)k^2\tilde{b}^{*2} - 2\left(k(1-k) - \gamma k^2\right)\tilde{b}^* m_1$$

Replacing $\tilde{b}^*$ and after some algebra, this reads

$$\frac{1}{1+\gamma} + (1+\gamma)k^2 - 2k$$

This function is decreasing on $[0, \frac{1}{1+\gamma}]$ and increasing on $[\frac{1}{1+\gamma}, +\infty)$, and equal to 0 at $k = k^{FB} = \frac{1}{1+\gamma}$.

Therefore, the value of delegation is always positive, and increasing in K. □

Proof of Proposition 7 The net benefit from delegating when the delegate targets an average audience $\tilde{b}$ reads

$$-(1+\gamma)k^2\tilde{b}^2 - 2\left(k(1-k) - \gamma k^2\right)\tilde{b}m_t \tag{54}$$

From the proof of Proposition 5, one sees that $\tilde{b}$ and m_1 have the same sign and that $|\tilde{b}| \leq |m_1|$. This implies $\tilde{b}^2 \leq \tilde{b}m_1$ for all m_1.

When $k > \frac{2}{1+\gamma}$, one has $0 < (1+\gamma)k^2 < 2\left((1+\gamma)k^2 - k\right)$. Using $0 \leq \tilde{b}^2 \leq \tilde{b}m_1$, one

[44] The uniform assumption is without loss of generality here.

derives that $(1+\gamma)k^2\tilde{b}^2 < 2\left((1+\gamma)k^2 - k\right)\tilde{b}m_1$, which is equivalent to (54) > 0.

Let us now consider the case $k \leq \frac{2}{1+\gamma}$. Since $\frac{\gamma}{1+\gamma}$ is the minimum of K and since $\rho(K)$ is decreasing, the assumption that $|m_1| < A - \sqrt{1 - K^{FB}\Sigma}$ ensures that $\rho(K) < (A - |m_1|)^2$ for all K, which implies $\tilde{b} = m_1$ in the optimal audience (see the proof of Proposition 5). Then the benefit from delegating becomes

$$-\left((1+\gamma)k^2 + 2\left(k(1-k) - \gamma k^2\right)\right)m_t^2. \tag{55}$$

Given that $k \leq \frac{2}{1+\gamma}$, this function is nonpositive. □

Proof of Proposition 8 Let $\mathcal{P}$ denote the partition chosen by the DM, and let $\mathcal{I}$ denote one element of the partition. Given the informational independence between segments, we derive that the expected value which the DM derives from a partition $\mathcal{P}$ is

$$\frac{1}{2(1-\delta)}\int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I})\left\{1 - \mathbb{V}(b|b\in\mathcal{I}) - K[\mathbb{E}(b|b\in\mathcal{I}) - m_1]^2 - K\Sigma\right\} \tag{56}$$

The DM should pick the partition which maximizes (56). Since $\mathcal{P}$ is a partition, one has

$$\begin{cases} \int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I}) = 1 & \text{(57a)} \\ \int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I})\mathbb{E}(b|b\in\mathcal{I}) = \mathbb{E}(b) & \text{(57b)} \end{cases}$$

(57a) reflects the fact that the total audience has mass 1, while (57b) is the Law of Iterated Expectations. After simplification, using (57a) and (57b), the DM maximizes

$$1 - K\left(\mathbb{E}(b) - m_1\right)^2 - K\mathbb{V}(b) - K\Sigma - (1-K)\int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I})\mathbb{V}(b|b\in\mathcal{I})$$

If $K \neq 1$, the choice of the partition only affects the DM's profit through $\int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I})\mathbb{V}(b|b\in\mathcal{I}) = \mathbb{E}[\mathbb{V}(b|b\in\mathcal{I})]$. It is easy to see, using the law of total variance, that $0 \leq \mathbb{E}[\mathbb{V}(b|b\in\mathcal{I})] \leq \mathbb{V}(b)$, with $\mathbb{E}[\mathbb{V}(b|b\in\mathcal{I})] = \mathbb{V}(b)$ when the partition consists of a single element $\mathcal{B}$, and $E[\mathbb{V}(b|b\in\mathcal{I})] = 0$ when each element of the partition is a singleton. If $K > 1$, the DM should maximize $\mathbb{E}[\mathbb{V}(b|b\in\mathcal{I})]$, and then selects centralization. On the contrary, if $K < 1$, he should minimize $\mathbb{E}[\mathbb{V}(b|b\in\mathcal{I})]$ and then builds individualized reputations with each of the receivers. □

Proof of Proposition 9 A first remark is that it is always optimal to exempt some receivers. Indeed, $P = 1$ and $\tilde{b} = \bar{b} = 0$ yields a negative payoff, so it is dominated by $P = 0$. Therefore, $\mathcal{E} \neq \varnothing$.

For P to be positive in equilibrium, one needs that there exists at least one value of $\tilde{b}$ such that $\tilde{b}^2 - K(\tilde{b} - m_1)^2 - K\Sigma > 0$. If $K < 1$, the function $\tilde{b}^2 - K(\tilde{b} - m_1)^2 - K\Sigma$ is convex, so is maximum either at $\tilde{b} = -A$ or at $\tilde{b} = A$. One also remarks that the maximal value is attained at $\tilde{b} = A$ if $m_1 \geq 0$ and at $\tilde{b} = -A$ if $m_1 \leq 0$.

Therefore, a necessary condition to have $P > 0$ is $A^2 - K(A - |m_1|)^2 - K\Sigma > 0$. Let us assume this is the case from now on.

In the case where $K > 1$, the function $\tilde{b}^2 - K(\tilde{b} - m_1)^2 - K\Sigma$ is concave. In this case, the function is not always negative on $[-A, A]$ if $m_1^2 > (K-1)\Sigma$ when $Km_1 < (K-1)A$ and $(1-K)A^2 + 2Km_1 - Km_1^2 - K\Sigma > 0$ otherwise.

Given a fixed $\tilde{b}$ such that $\tilde{b}^2 - K(\tilde{b} - m_1)^2 - K\Sigma > 0$, the DM wants to maximize P. It is easy to see that the maximum P compatible with a conditional expectation $\tilde{b}$ is $\frac{A-\tilde{b}}{A}$ if $\tilde{b} \geq 0$ and $\frac{A+\tilde{b}}{A}$ if $\tilde{b} \leq 0$.

Therefore we are interested in the maximum of

$$f(\tilde{b}) \equiv (A - \tilde{b})\left(\tilde{b}^2 - K(\tilde{b} - m_1)^2 - K\Sigma\right)$$

on $[0, A]$ and

$$g(\tilde{b}) \equiv (A + \tilde{b})\left(\tilde{b}^2 - K(\tilde{b} - m_1)^2 - K\Sigma\right)$$

on $[-A, 0]$.

Then the DM picks among the solutions of each maximization problem the one which yields the higher value. It is easy to see that if $m_1 > 0$ (resp. $m_1 < 0$), the overall solution must be positive (resp. negative). Indeed, suppose that $m_1 > 0$ and consider the value $b_0 \leq 0$ which maximizes g on $[-A, 0]$. One easily check that $f(-b_0) - g(b_0) = -4K(A + b_0)m_1 b_0 > 0$. Therefore, the global maximum $\tilde{b}$ attainable to the DM must have the same sign as m_1. □

Proof of Proposition 10 When $m_1 = 0$, we know from the above result that there are two equivalent solutions generating the same payoff, one on $[-A, 0]$, the other on $[0, A]$. Let us focus on the latter one. It maximizes $(A - \tilde{b})\left((1-K)\tilde{b}^2 - K\Sigma\right)$. First, it is easy to see that this function is nonpositive for any $K \geq 1$. In this case, the DM exempts everyone. If $K < 1$, the solution to this problem is $\tilde{b}(0) = \frac{1}{3}\left(A + \sqrt{A^2 + \frac{3K}{K-1}\Sigma}\right)$. It is easy to see that $\tilde{b}(0)$ is nondecreasing in K and that $\tilde{b}(0) > A/2$, which implies that $0 \in \mathcal{E}$. □