Horizontal reputation and strategic audience management

Matthieu Bouvard*, Raphaël Levy†

July 2018

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 be lower than in the no-reputation case. We highlight how strategic audience management (e.g., centralization, delegation to third parties with dissenting objectives) 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 mostly 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 coexistence of two equilibria with distinct efficiency properties. While reputational concerns always improve welfare in the more efficient equilibrium (reputation is "good"), they may depress welfare in the less efficient one (reputation is "bad"). Second, we introduce strategic audience management as a natural remedy to the inefficiency of reputational incentives in the presence of heterogenous audiences, and show how multiple institutional designs endogenously emerge as optimal responses to good and bad reputation.

We build a model in which a decision maker (e.g., 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 preferred by the decision maker. This distortion provides a rationale for policy interventions aiming at realigning incentives. Specifically, the decision maker affects receivers' incentives by uniformly shift-

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

[2]There are a few exceptions though: a recent stream of literature has begun to investigate aspects of reputation with multiple audiences (Bar-Isaac and Deb, 2014a,b; Frenkel, 2015; Shapiro and Skeie, 2015; Bouvard and Levy, forthcoming), following an older literature on cheap talk or signaling with two audiences (Farrell and Gibbons, 1989; Gertner, Gibbons, and Scharfstein, 1988).

ing their marginal benefit of investment. Because the decision maker lacks commitment power, his intervention is driven by the 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, and 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 lower than in the infinitely repeated static game where the decision maker always plays the myopic action: reputation may be "bad". 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. We introduce the possibility for the decision maker to ex ante design the institutional framework which governs the modalities of interaction with the audience. First, we allow the decision maker to decentralize decision making to benevolent surrogates, and show that decentralization dominates if and only if reputation is good. On the one hand, decentralization allows to tailor interventions to the idiosyncratic bias of each receiver. On the other hand, individualized interventions imply catering to extreme receivers, which generates large costs when reputation leads to overreactions. In that case, centralization provides a valuable commitment to target the average receiver only. Equilibrium multiplicity in the reputation game then translates into multiple institutional forms, with centralization arising in the worse equilibrium as a by-product of bad reputation. Next, we allow the decision maker to mandate a delegate whose actions impact the whole audience, but who targets a narrower audience ascribed by the decision maker. We show that such dissent is a way to correct for the inefficiency of reputational incentives. When the decision maker cannot impose a target audience to the delegate, the latter chooses his own audience to maximize the benefits he privately derives from exerting influence. The target chosen by the delegate now endogenously diverges from the one the decision maker would ascribe to him, which sometimes makes delegation undesirable. However, delegation always dominates when reputation is bad: when reputational deficits are more costly to withstand, both the DM and the delegate agree on the necessity to target an audience better aligned, on average, with their reputation. An equilibrium with delegation then coexists with one better equilibrium in which reputation is more efficient and delegation accordingly undesirable. Finally, we consider exemptions whereby the decision maker insulates a fraction of his audience from the impact of his intervention. The decision maker primarily needs to influence receivers with more extreme biases, who generate the highest welfare losses. However, exerting influence on receivers at one extreme requires exempting receivers at the other. As a result, the decision maker optimally picks one side, the one he is intrinsically more efficient at influencing, and exempts the other one. Here too, audience design is responsive to the efficiency of reputation: in the less efficient equilibrium, the scope of exemptions widens.

The leading application we consider throughout the paper is that of politics. Politics is a natural application, as politicians not only need to sustain a reputation vis-à-vis several stakeholders with different preferences, but can also somewhat shape the framework in which this reputation is being built. Our results shed light on the rationale behind strategies of delegation to independent or supranational institutions: public agencies (e.g., independent central banks) are often assigned restrictive objectives (e.g., price stability) in order to protect them from political pressure which politicians themselves can hardly commit to resist. We interpret such a recourse to *narrow mandates* as an institutional response aiming at insulating decision making from certain audiences whose relevance to the decision maker would be ultimately harmful. This audience-based motive for delegation is also consistent with the tendency of policy-makers to delegate control to independent bodies and to later blame them for being insufficiently representative of the electorate's interests ("blame-shifting"). The transfer of the prudential oversight of European banks from the local to the EU level in 2012 provides an illustration of our results on the relative merits of centralization and decentralization in an environment where agents' behaviors are particularly sensitive to expectations about government intervention. Beyond politics, our setup applies to any situation where a planner uses his reputation to simultaneously exert influence on multiple agents and achieve a better alignment between individual and social incentives. For instance, organizations strive to provide the right effort incentives (e.g., foster synergies, encourage information acquisition...) to workers with different preferences or skills, and often rely on corporate reputation ("corporate culture") as a key complement to explicit contracts in enabling coordination (Kreps, 1990). There, strategic audience management can be interpreted as an organizational choice (boundaries of the firm, hierarchical organization...) that responds to the (in)efficiency of reputational concerns. In such a context, the coexistence of an equilibrium with centralization and one with decentralization is consistent with the well documented empirical pattern that organizations often switch back and forth between a centralized and a decentralized structure (Eccles, Nohria, and Berkley, 1992; Nickerson and Zenger, 2002).

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 reputational payoff is linear in Holmström, and the equilibrium strategy is accordingly independent of the reputation. By contrast, in our setting, the equilibrium strategies always depend on the reputation, and the concavity of the payoff function generates multiple equilibria and the possibility of inefficient reputation-building. Second, our paper relates to Cisternas (2017), who studies signal-jamming in continuous time and derives general conditions under which the 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 we derive multiple equilibria in closed form and establish their distinctive efficiency properties as compared to the no-reputation case. This multiplicity, together with bad reputation, 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 (forthcoming) establish that a certifier who needs to attract sellers of different qualities 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. As we do, these papers obtain repositioning towards the bliss reputation, but 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 informativeness of past outcomes 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. Our audience design approach is related to Bar-Isaac and Deb (2014a), who consider a sender who picks two actions that each affect a different (fixed) audience, and contrast reputation building with common and separate observations of actions. By contrast, when studying decentralization, we purposefully abstract from learning across audiences, and focus on the trade-off between a uniform treatment of all receivers, and a differential treatment of endogenous subsets of the initial audience.

Finally, our focus on the welfare implications of reputation building 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. Such a separation motive is absent in our model where information remains symmetric in equilibrium. Rather, reputation becomes inefficient when the decision maker overreacts to his reputation.

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 surplus $y_t - \frac{y_t^2}{2}$.

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 a 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

[3] We deliberately make this assumption in a first stage to examine how the DM tries to simultaneously influence several audiences. In Section 4.2, we allow the DM to design the environment in such a way to differentiate the impact of his intervention across receivers.

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

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

where $\bar{b} \equiv \mathbb{E}(b)$.

From (3), the maximal surplus is then attained when receivers expect the DM's intervention to perfectly correct 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. 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 will exactly offset 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

[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.

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 action a_t^e they might expect. Given the normality and independence assumptions, the dynamics of beliefs is simple to characterize: the conditional distribution 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 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 (e.g., policy-maker) who uses a policy instrument to influence his environment and achieve a better alignment between individual and social incentives. Policy-makers (or public agencies mandated by them, e.g., central banks, regulators) need to design policies such that agents better internalize the externalities that their actions (e.g., labor supply, R&D investments, savings, location choice...) inflict on others. In this context, the DM's intervention affects how much receivers expect their investments to be rewarded, and his type θ_t captures his intrinsic ability to create environments that reward certain types of investment through the appropriate intervention (e.g., fiscal, monetary, trade policy). Alternatively, the type could account for an unknown state of the world which governs the returns on investment in the economy (e.g., the level of inflation, or the severity of technological, financial or institutional constraints).

We interpret the cost $\gamma\frac{a_t^2}{2}$ of the DM's intervention as a deadweight loss that captures resources spent to design and implement a policy, or frictions which the DM's intervention creates (e.g., rents related to moral hazard or adverse selection, distortionary taxation, bargaining costs, etc). The symmetry of the cost function in the positive and negative ranges, while ensuring tractability, allows more flexible interpretations.[5]

From the DM's perspective, actions are motivated by the desire to jam the audience's inference about his type in the hope of bringing his reputation closer to the bliss point $\bar{b}$. But actions also affect the true returns on investment for the receivers, hence have an impact on aggregate efficiency. Therefore, the role of reputation is two-fold. First, it has a direct influence on receivers' actions through beliefs; 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 reputation, and to how much they are willing to maintain or improve it.[6]

3 Equilibrium analysis

3.1 The commitment benchmark

Before going through the analysis of reputation building, we derive the optimal profile of actions under full commitment. Let

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

denote the function capturing the gross surplus. The net surplus in period t then reads

$$\pi(m_t + a_t) - \gamma\frac{a_t^2}{2} = \frac{1-\mathbb{V}(b)}{2} - \frac{(m_t + a_t - \bar{b})^2}{2} - \gamma\frac{a_t^2}{2}, \tag{6}$$

[5]For instance, if y_t is an investment which generates pollution, a positive a_t may be interpreted as an implicit subsidy to the polluting 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.

[6]In line with our modeling, these interventions are typically imperfectly observed, and they affect heterogeneous agents in an undifferentiated way (Farhi and Tirole, 2012).

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 he has no reputational concerns, the DM 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* $k_1 \equiv \frac{\delta h_\varepsilon}{\gamma(h_1+h_\varepsilon)}$.

Two features of Proposition 1 are particularly relevant to the understanding of the stationary case. First, the equilibrium is unique. Second, the DM's equilibrium action aims at correcting his reputational deficit $\bar{b} - m_1$, that is, how far m_1 falls away from his bliss reputation. 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:[7]

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

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 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

$$m_{t+1}(a_t, a_t^e) = \lambda m_t + (1-\lambda)[\theta_t + \varepsilon_t + a_t - a_t^e], \tag{7}$$

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

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

$$m_{t+1}^{DM} = \lambda m_t^{DM} + (1-\lambda)(\theta_t + \varepsilon_t). \tag{8}$$

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

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:

$$a(m_t^{DM}, m_t) \in \underset{a_t}{\operatorname{argmax}}\, \delta \mathbb{E} V\left(m_{t+1}^{DM}, m_{t+1}[a_t, a_t^e]\right) - \gamma \frac{a_t^2}{2}, \tag{9}$$

ii) receivers have rational expectations:

$$a_t^e = a(m_t, m_t), \tag{10}$$

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

$$\begin{aligned} V(m_t^{DM}, m_t) &= \pi[m_t + a(m_t, m_t)] - \gamma \frac{a(m_t^{DM}, m_t)^2}{2} \\ &+ \delta \mathbb{E} V\left(m_{t+1}^{DM}, m_{t+1}[a(m_t^{DM}, m_t), a(m_t, m_t)]\right), \end{aligned} \tag{11}$$

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

$$\lim_{t \to +\infty} \delta^t \mathbb{E}_1 V\left(m_t^{DM}, m_t(\hat{a})\right) = 0, \tag{12}$$

where $m_t(\hat{a})$ is the public reputation when receivers expect the DM to follow the equilibrium strategy $a(.,.)$, but the DM instead follows an arbitrary strategy $\hat{a}$ between 0 and t.[8]

Note that strategies describe the DM's behavior both on and off-path. In particular, condition (9) states that the DM's action is optimal even following an undetected deviation

[8]More precisely, we require that, if (12) does not hold for an admissible $\hat{a}$, then $\hat{a}$ is dominated by a strategy that satisfies (12) (see Appendix).

(i.e., if $m_t^{DM} \neq m_t$).[9]

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*

$$V^k(m_t) = \frac{1}{2(1-\delta)}\left(1 - \mathbb{V}(b) - K(\bar{b} - m_t)^2 - K\Sigma\right), \tag{13}$$

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

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

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

In the DM's value function (13), $K = (1-k)^2 + \gamma k^2$ 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

$$\pi[m_t + k(\bar{b} - m_t)] - \gamma\frac{k^2(\bar{b} - m_t)^2}{2} = \frac{1 - \mathbb{V}(b)}{2} - \frac{K}{2}(\bar{b} - m_t)^2. \tag{14}$$

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 (though vain in affecting receiver beliefs in equilibrium), the first term $(1-k)^2$ measures the impact of the DM's intervention on the efficiency of the 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

[9]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.

commit to. K accordingly measures how good a substitute for commitment reputation is (the lower K the more efficient reputation).

To understand where the form 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 his private belief m_t^{DM} was adjusted to the public belief m_t.[10]

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

$$\frac{1 - \mathbb{V}(b)}{2} - \frac{K}{2}[\bar{b} - \mathbb{E}_t(m_{t+i})]^2 - \frac{K}{2}\mathbb{V}_t(m_{t+i}). \tag{15}$$

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 (7)).[11] Overall, the effect of marginally increasing a_t on the payoff in 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, that is, its impact on the discounted sum of future payoffs reads

$$(1 - \lambda)\sum_{i=1}^{+\infty} \delta^i \lambda^{i-1} K \left[\bar{b} - \mathbb{E}_t(m_{t+i})\right] = \frac{\delta(1 - \lambda)}{1 - \delta\lambda} K(\bar{b} - m_t), \tag{16}$$

while its marginal cost reads

$$\gamma k(\bar{b} - m_t).^{12} \tag{17}$$

[10]I.e., the DM plays $a^*(m_t, m_t)$ even though $m_t^{DM} \neq m_t$. 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. Note that Cisternas (2017) uses an equivalent approach to derive a necessary equilibrium condition (see in particular Remark 5 in his paper).

[11]Notice that the impact of a_t on m_{t+i} is deterministic, meaning that the variance of future reputations does not depend on the DM's actions. From (15), 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.

[12]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

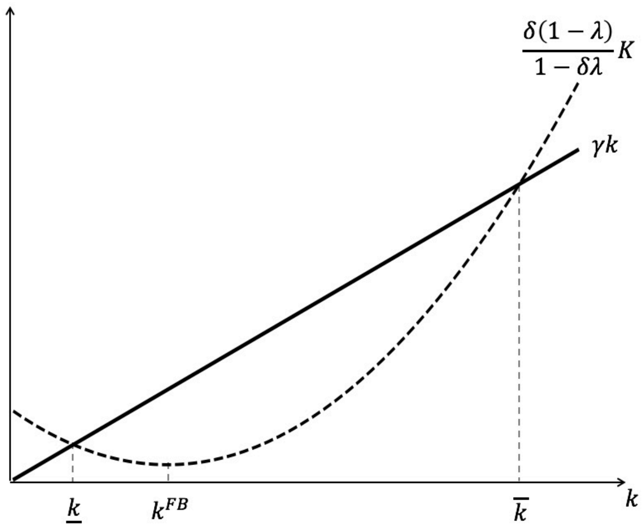

Figure 1: Marginal Cost (solid) and Marginal Benefit (dashed) of reactivity k.

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 (16) and (17). 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 $\overline{k}$ where the DM overreacts ($\overline{k} \geq k^{FB}$). In the latter equilibrium, 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 $\overline{b}$, and raises the marginal benefit from reacting today, 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. Remember that the two-period game features a unique equilibrium, as would any finite-horizon version of the model.[13] 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.[14]

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.

[13]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.

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

Second, the complementarity between current and future responsiveness is driven by the concavity of the payoff function π, which itself drives the concavity of the value function V.[15] 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 value function is very steep. This raises the marginal benefit of the DM's action, which sustains 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 value 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, makes reputational concerns less salient, hence lowers incentives to improve reputations. 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 about his type but not about his action, the existence of a linear equilibrium would require an additional restriction that the DM's action is costly enough relative to the discount factor.[16]

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

[15] In the linear version of Holmström (1999), the equilibrium actions are independent of the reputation, and the equilibrium is accordingly unique, even in the stationary game.

[16] Notice that Cisternas (2017) derives a similar condition in his quadratic specification. Indeed, if γ is too small, incentives to fool the market are too strong, precluding equilibrium existence. However, when the action bears an efficiency impact, as in our model, incentives to fool the market can be made smaller even when γ is low, as the DM is less sensitive to improvements in his reputation when actions are close to efficient.

therefore potentially ambiguous.[17] Before investigating the welfare properties of each equilibrium, let us introduce the following definition:

Definition 1. *We say that reputation is good (resp. bad) when the DM obtains an equilibrium payoff larger (smaller) than in the infinitely repeated static game where $a_t = 0$ for all t.*

We then refer to "bad reputation" to describe situations where the DM would like to commit not to build a reputation.[18]

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 more efficient actions than in the no-reputation case, but are generically insufficient to reach efficiency (Holmström, 1999). 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 this overreaction is so large that he ends up being worse off than in the no-reputation case. Notice that overreactions are directly costly because the action involves a convex cost $\frac{1}{2}a_t^2$, but they may also reduce welfare even ignoring the costs borne by the DM to build a reputation:

Remark 1. $\overline{k} > 2 \Leftrightarrow \pi(m_t + a_t^*) < \pi(m_t)$ *for all* m_t.

To see this, it is enough to write the gross surplus in period t,

$$\pi(m_t + a_t^*) = \frac{1}{2}\left(1 - \mathbb{V}(b) - (k-1)^2(m_t - \overline{b})^2\right). \tag{18}$$

[17] By welfare, we mean here the expected discounted payoff of the DM. In the political application we consider, we implicitly have in mind a benevolent planner, but the DM's payoff need not coincide with social welfare.

[18] 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).

It 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. Notice that, while equilibrium multiplicity is driven by the concavity of π, the result of Remark 1 stems from the non-monotony of π: if π were to be increasing, we would always have $a_t^* > 0$, hence $\pi(m_t + a_t^*) > \pi(m_t)$. Here, the over-reaction is so large that it leads to an excessive adjustment of receivers' incentives leading to over-investment (in the case $m_t < \overline{b}$) even from some receivers who otherwise tend to under-invest.

Finally, and relatedly, since $\overline{k} > 1$, the average level of investment $1 - \overline{b} + \mathbb{E}(x_t) = 1 - (1-\overline{k})(\overline{b} - m_t)$ 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.[19]

3.5 Comparative statics

In both equilibria, the DM tries to reposition in the direction of the bliss reputation $\overline{b}$. 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, where we have plotted the marginal benefit and marginal cost of increasing the responsiveness k (given by (16) and (17)). As one sees, the slope of the marginal benefit curve is smaller than the slope of the marginal cost line at $\underline{k}$

[19] 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 (forthcoming) show that an increase in a politician's reputation may hurt voters. This happens over some range of beliefs where a higher reputation implies more uncertainty on the politician's type, i.e., stronger incentives for the politician to build a reputation through inefficient pandering, which may ultimately outweigh the positive impact of a better reputation.

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 total surplus in period t is larger.[20] 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 reliability of past signals 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.[21]

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. 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.[22] 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.[23] Accordingly, the reason why reputation depresses welfare is essen-

[20]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 (13).

[21]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.

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

[23]One may find surprising that a higher δ increases welfare after we have stressed that the DM could be better off in the game where he behaves myopically than in the high-responsiveness equilibrium. This

tially different, and actually stems from the DM over-reacting to his reputation in the bad equilibrium.

4 Strategic audience management

In the previous section, we have analyzed how the DM builds a reputation given his audience, and shown that his intertemporal welfare at date 1 is (proportional to)

$$1 - \mathbb{V}(b) - K(\bar{b} - m_1)^2 - K\Sigma. \tag{19}$$

This expression evidences how the DM's value function relates to the initial distribution of receivers' biases b. First, since the impact of his intervention is uniform across receivers, the loss stemming from the variance of investments equals the variance of biases $\mathbb{V}(b)$. Second, he incurs a loss from the mismatch between his initial reputation m_1 and the average bias $\bar{b}$ (his bliss reputation), which is amplified by the inefficiency of reputational concerns measured by K. This suggests that altering the composition of the audience relevant to the decision maker could improve efficiency and that such audience design should be more desirable when K is larger. In this section, we take advantage of the flexibility of our framework to explore several audience management strategies. The upshot is that the optimal design ultimately depends on whether reputation is good or bad. Therefore, our results on equilibrium multiplicity (Proposition 2) and bad reputation (Proposition 3) have far-reaching implications: when a bad reputation equilibrium exists, the two equilibria of the reputation game result in distinct organizational forms as optimal responses.

4.1 The audience management model

We start by describing how our baseline model can accommodate strategic audience management. Consider a slightly modified version of the model where incentives are driven by reputational concerns with respect to a subset $\mathcal{I} \subset \mathcal{B}$ of the whole audience, and let $P(\mathcal{I}) \equiv \int_{b\in\mathcal{I}} dF(b)$ denote the mass of receivers in $\mathcal{I}$. The total welfare created within

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 $\delta = 0$.

audience $\mathcal{I}$ in period t reads:[24]

$$\begin{aligned}
& \int_{b\in\mathcal{I}} [1-b+\mathbb{E}_t(x_t) - \frac{1}{2}\left(1-b+\mathbb{E}_t(x_t)\right)^2 - \gamma\frac{a_t^2}{2}]\, dF(b) \\
= \quad & P(\mathcal{I})\left(\frac{1}{2} - \frac{1}{2}\mathbb{V}(b|b\in\mathcal{I}) - \frac{1}{2}\left(m_t + a_t - \mathbb{E}(b|b\in\mathcal{I})\right)^2 - \gamma\frac{a_t^2}{2}\right). \qquad (20)
\end{aligned}$$

Remarking the similarity between (20) and (6), one derives that the equilibrium strategy given an audience $\mathcal{I}$ is

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

where $k \in \left\{\underline{k}, \overline{k}\right\}$.

Note that the two levels of responsiveness $\underline{k}$ and $\overline{k}$ are the same as in Section 3, that is, they are independent of the audience $\mathcal{I}$. Therefore, the audience affects the equilibrium action only through the reputational deficit with respect to the average bias in the audience, $\mathbb{E}(b|b\in\mathcal{I}) - m_t$. This means that there is no difference between facing an heterogenous audience $\mathcal{I}$ and a single receiver with the average bias $\mathbb{E}(b|b\in\mathcal{I})$ in terms of the profile of actions. However, the welfare created within a given audience also depends on the measure of receivers in this audience as well as the dispersion of their actions (this is apparent, for instance, in (20)). Importantly, we allow actions driven by reputation with respect to an audience $\mathcal{I}$ to impact a wider audience $\mathcal{J}$. To make this distinction clear, we let $\tilde{V}^k(\mathcal{I}, \mathcal{J}, m_1)$ denote the surplus created within a subset $\mathcal{J}$ of the audience when actions are driven by reputational incentives with respect to an audience $\mathcal{I}$ that need not coincide with $\mathcal{J}$. Formally, one has

$$\begin{aligned}
& \tilde{V}^k(\mathcal{I}, \mathcal{J}, m_1) \equiv \\
& P(\mathcal{J})\sum_{t=1}^{\infty}\delta^{t-1}\mathbb{E}\left(\frac{1}{2} - \frac{1}{2}\mathbb{V}(b|b\in\mathcal{J}) - \frac{1}{2}\left(m_t + a_t^*(m_t,\mathcal{I}) - \mathbb{E}(b|b\in\mathcal{J})\right)^2 - \gamma\frac{a_t^*(m_t,\mathcal{I})^2}{2}\right)
\end{aligned}$$

This formulation nests the original model in that $\tilde{V}^k(\mathcal{B}, \mathcal{B}, m_1) = V^k(m_1)$, and provides a consistent framework to study different audience management strategies. In what follows, we examine the incentives of the DM to decentralize decision-making, to delegate control, and to exempt some receivers from his intervention. These design decisions are assumed to

[24] In order to rule out economies of scale and to focus the analysis on the role of reputation only, we assume that the cost of the action a_t is proportional to the size of the audience $P(\mathcal{I})$.

be enduring organizational choices made at $t = 1$, which, once taken, cannot be adjusted to future reputations.

4.2 Decentralization

We start by comparing the relative merits of centralization and decentralization in the light of our reputation model. Decentralization here consists in moving away from uniform treatment and tailoring interventions to the biases of receivers. Formally, we allow the DM to partition the audience into subsets and appoint a surrogate to each of them. We make two important assumptions. First, we assume that surrogates are identical to the DM in every respect, in particular, they have the same initial reputation m_1.[25] Second, we assume away learning across segments, that is, neither the surrogate nor receivers in a given segment can learn from the outcomes in other segments.[26] This may be because θ_t is a common shock capturing some aggregate state of the economy but players do not observe the payoffs in other segments, or because the shocks underlying these payoffs are orthogonal (e.g., θ_t could reflect some idiosyncratic feature of each surrogate).

Each surrogate is assumed to be benevolent, i.e., behaves in such a way to maximize welfare on his own segment $\mathcal{I}$, hence chooses $a_t^*(m_t, \mathcal{I}) = k\left(\mathbb{E}(b|b \in \mathcal{I}) - m_t\right)$. Since the DM cares about aggregate welfare, the problem he faces is then to choose a partition $\mathcal{P}$ of the set of types $\mathcal{B}$ that maximizes

$$\int_{\mathcal{I} \in \mathcal{P}} \tilde{V}^k(\mathcal{I}, \mathcal{I}, m_1) \tag{21}$$

At the two extremes, full centralization is equivalent to $\mathcal{P} = \{\mathcal{B}\}$, i.e., the original setup, and full decentralization to $\mathcal{P} = \{\{b\}, b \in \mathcal{B}\}$: each receiver is assigned to a different surrogate. In between are intermediate levels of decentralization where different surrogates each exert influence on a group of heterogenous receivers. While we allow for such partial decentralization, the next result shows that the optimal design is bang bang.

[25]Considering surrogates with different reputations would stack the choice in favor of more decentralization, as it would then be possible to allocate surrogates to a group of receivers they are intrinsically closer to, hence more able to influence.

[26]Notice that Bar-Isaac and Deb (2014a) contrast the value of reputation when there is cross learning and when each receiver only observes the action targeted at him in a context where receivers are treated differentially. Instead, we compare uniform and differential treatment in the absence of cross learning. Therefore, our approach is complementary to theirs.

Proposition 5. *The DM optimally chooses full decentralization when reputation is good ($K \leq 1$) and centralization when it is bad ($K \geq 1$).*

To understand the tradeoff underlying 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, the institutional choice only affects the variance terms.[27] First, the variance of the actions $\mathbb{V}(a_t)$ is nil under centralization (there is a single action taken by a single decision maker), but equals $k^2\mathbb{V}(b)$ under decentralization: since actions are adapted to receivers, the dispersion of actions is larger under decentralization, the more so the stronger the responsiveness k. However, since actions are better tailored to the specific bias of receivers, decentralization may improve 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. When the responsiveness is too large (when $k > 2$), $\mathbb{V}(y_t)$ also becomes larger under decentralization.[28] Centralization then unambiguously dominates. Actually, centralization allows to commit not to cater to every single receiver, which is more generally valuable when reputation costs are too large.[29] As k decreases, however, decentralization lowers the variance of investments, at some point sufficiently to offset the increase in $V(a_t)$, so that decentralization ultimately becomes dominant.[30] Whether the benefit from adaptation dominates the cost of duplicating interventions depends on whether reputation is good or bad. This is a counterpart result of Proposition 3 with

[27]Because y_t and a_t are linear in b, 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.

[28]This is the counterpart effect of the one stated in Remark 1: when $k > 2$, not only does the average investment gets further apart from the efficient investment, but the variance of investments also increases.

[29]This is reminiscent of the merits of public versus private communication in cheap talk games with multiple audiences (Farrell and Gibbons, 1989). In Farrell and Gibbons, whether public or private communication dominates depends on the extent of misalignment in preferences, with a strong positive bias in one state of the world being possibly outweighed by a strong negative bias in the other state. In our setup, however, there is no privately observed state of the world, and whether centralization dominates is independent of both the misalignment of preferences within the audience and between the audience and the DM.

[30]Notice that because the optimal partition 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 a finite number N of potential surrogates. In this case, the identity of receivers belonging to each segment matters when $K < 1$. When b is uniformly distributed, it is actually simple to derive that the DM maximizes his payoff by having N intervals of equal size.

variance terms: Proposition 3 establishes that reputation is good when the gain from a more efficient average investment outweighs the cost of reactions, while Proposition 5 shows that decentralization dominates when the gain from less dispersed investments outweighs the total cost from more dispersed actions.

Proposition 5 shows how the optimal institutional form arises as a response to the inefficiency of reputational concerns.[31] While decentralization would be the ideal organization form in a first best world where policy-makers could commit to a course of action, over-reactions to reputational incentives actually make decentralization unprofitable. Put differently, centralization arises in the worse equilibrium as a by-product of the inefficiency of reputation. In the debate on decentralization versus centralization (see, e.g., Oates (1999) for an account of this debate for fiscal policy), this suggests that the merits of each option should not be established per se, but in relation to the efficiency of decision-making in other dimensions.

Relatedly, Proposition 5 suggests that different institutional forms may arise in otherwise similar economic environments. For instance, banking supervision was, until recently, conducted at the country level in the European Union. Local supervision allows to adjust the amount of risk that banks can take to the specificities of regional markets. The recent development of the European Banking Union reversed this architecture by transferring prudential supervision and interventions in the banking system at the Union level. Interestingly, this change happened in the wake of the 2012 Eurozone crisis, and amid growing concerns about the credibility of government intervention and the intricate connexions between national governments and banks.[32] This is consistent with our idea that centralized intervention may be a remedy for the regulator in environments where expectations lock him into an inferior equilibrium. In line with our multiple equilibria result, such a regime switch could also be interpreted as the consequence of a change in expectations regarding the responsiveness of policy responses to economic conditions that need not be related to any change in the fundamentals.

[31] In the no reputation case ($k = 0$, i.e., $K = 1$), the centralization decision is irrelevant.

[32] See, e.g., the 7th report of the 2012-13 session of the British House of Lords' European Committee, "European Banking Union: Key issues and challenges."

4.3 Delegation

In the previous section, we have allowed receivers to be treated differentially, but imposed that each receiver matter, i.e., belong to one surrogate's audience. We now consider the possibility for the DM to insulate himself from reputational concerns vis-à-vis some selected receivers. Since the DM cannot change his own preferences (audience), we consider how he can indirectly achieve this goal by mandating a delegate who targets a subset $\mathcal{I}$ of the whole audience $\mathcal{B}$, even though his actions affect all receivers in $\mathcal{B}$. To underline this audience-based rationale for delegation, we assume that the delegate differs from the DM through the composition of his audience only.[33]

We first consider a situation where the DM can mandate the delegate to target an audience $\hat{\mathcal{I}}$. The problem of the DM now reads

$$\max_{\hat{\mathcal{I}}} \tilde{V}^k(\hat{\mathcal{I}}, \mathcal{B}, m_1). \tag{22}$$

Since the profile of actions depends on the audience through the average bias within the audience only, imposing a mandate $\hat{\mathcal{I}}$ is akin to imposing an average receiver $\hat{b} \equiv \mathbb{E}(b|b \in \hat{\mathcal{I}})$, hence a profile of actions $a_t^*(m_t, \hat{\mathcal{I}}) = k(\hat{b} - m_t)$. To simplify matters, we assume from now on that receivers' biases b are uniformly distributed on $\mathcal{B} = [-A, A]$, where $A > 1$.[34] This implies $\overline{b} = \mathbb{E}(b) = 0$.

Proposition 6. *The DM assigns to the delegate a mandate $\hat{b}^* = \left(1 - \frac{1}{(1+\gamma)k}\right) m_1$. The amount of dissent $|\hat{b}^*|$ increases with K, i.e., when reputation becomes less efficient. The value of 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, so delegation brings no value: $\hat{b}^* = 0$, and the delegate 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

[33]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 $\overline{b}$). 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).

[34]The assumption of a uniform distribution is made for analytical convenience, but is inconsequential.

suboptimal. Indeed, assigning a target $\hat{b}^*$ induces the optimal action in period 1:[35]

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

When $k \leq \frac{1}{1+\gamma}$, the delegate under-reacts to his reputational deficit, which can be compensated by artificially inflating this deficit. Conversely, when $k > \frac{1}{1+\gamma}$, reputation leads to over-reactions, and the DM lowers the reputational deficit by requesting the delegate to target a more congruent audience. In either case, the level of dissent necessary to correct for this inefficiency increases with K, as the equilibrium actions then become less efficient. Relatedly, the DM is less efficient himself, which increases the value of delegation.

Such dissent between the objective of the DM and the mission assigned to the delegate can be implemented through "narrow mandates." For instance, public agencies are often assigned narrow and well-identified missions by public authorities with broader objectives (Wilson, 1989). Such delegation strategies can play the role of insulating decision-making from the influence of certain constituencies (e.g., from the public opinion). Here also, independent central banks provide a useful illustration. For instance, the stated objective of the European Central bank is price stability, and some central banks, like the Bank of England, are even assigned an explicit target for the inflation level (see, e.g., Blinder, Ehrmann, Fratzscher, Jakob, and David-Jan (2008)), although pursuing such a restricted objective might have adverse effects on a range of macroeconomic variables such as investment, employment and growth.

There are instances, however, where the DM is not able to perfectly shape the contours of the delegate's audience. For instance, regulators may engage in side-contracting with selected firms allowing them to appropriate a fraction of the surplus generated by these firms (e.g., through direct funding or future employment opportunities).[36] There, the fact the delegate has a private motive to target an audience that differs from the one the DM would optimally assign to him introduces an agency problem.[37]

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

[36] See for instance Laffont and Tirole (1993) for a detailed account on regulatory capture.

[37] Alternatively, 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.

To endogenize this friction, we allow the delegate (when given control) to freely determine the composition of his own audience. Specifically, the delegate selects an audience $\tilde{\mathcal{I}}^*$ to maximize

$$\max_{\tilde{\mathcal{I}}} \tilde{V}^k(\tilde{\mathcal{I}}, \tilde{\mathcal{I}}, m_1), \tag{23}$$

As (23) shows, the delegate actually internalizes the investments of receivers who belong to $\tilde{\mathcal{I}}$ only. While the DM cares about the whole audience, the delegate can strategically forgo the payoff generated by a subset of constituents if this makes him overall better off. Notice that

$$\tilde{V}^k(\tilde{\mathcal{I}}, \tilde{\mathcal{I}}, m_1) = \frac{1}{2(1-\delta)} P(\tilde{\mathcal{I}}) \left(1 - \mathbb{V}(b|b \in \tilde{\mathcal{I}}) - K[\mathbb{E}(b|b \in \tilde{\mathcal{I}}) - m_1]^2 - K\Sigma\right). \tag{24}$$

In choosing the composition of his audience, the delegate faces a trade-off between the value of increasing the mass of receivers in his audience $P(\tilde{\mathcal{I}})$, the associated cost of increasing the variance of their investments $\mathbb{V}(b|b \in \tilde{\mathcal{I}})$, and the potential cost that more receivers may imply to include receivers with biases b that are further away from his own reputation m_1. Letting $\tilde{b}^* = \mathbb{E}(b|b \in \tilde{\mathcal{I}}^*)$ denote the average bias in $\tilde{\mathcal{I}}^*$, we obtain the following lemma.

Lemma 1. *The audience $\tilde{\mathcal{I}}^*$ is more congruent (i.e., $(\tilde{b}^* - m_1)^2$ is smaller) in the less efficient equilibrium than in the more efficient equilibrium.*

The impact of the efficiency parameter K on the optimal audience $\tilde{\mathcal{I}}^*$ is transparent from Equation (24). Since the mismatch with the average receiver becomes more costly when the equilibrium is less efficient (for higher K), more inefficient equilibria are associated with more congruent audiences. Coming back to the delegation decision, the benefits of dissent now have to be weighted against the cost of (endogenously) misaligned incentives between the DM and the delegate. Formally, the DM now delegates whenever

$$\tilde{V}^k(\tilde{\mathcal{I}}^*, \mathcal{B}, m_1) \geq \tilde{V}^k(\mathcal{B}, \mathcal{B}, m_1) \tag{25}$$

Proposition 7. *When the delegate can choose the composition of 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 $|m_1| < A - 1$.*

Proof In the Appendix.

Proposition 7 shows that delegation may become undesirable when the delegate can endogenously set his own target. Delegation is then optimal only when the reputation equilibrium is inefficient enough. There are two effects at play: first, even when the DM fully controls the mandate of the delegate, the value of delegation increases with K (Proposition 6); second, when K is large, the negative externality (that the delegate does not internalize the welfare of receivers he selects out) looms small for the DM as compared to the costs that would be required to influence these (non-congruent) receivers. In particular, in the extreme case where $K = +\infty$, the agency problem between the DM and the delegate has no bite: the DM would assign a target $\hat{b}^* = m_1$ to the delegate, but the delegate also chooses a perfectly congruent audience himself ($\tilde{b}^* = m_1$), in both cases to save on the costs of responsiveness. Therefore, the value of mandates and the alignment of preferences both increase with K, and delegation is optimal only when reputation is bad.

Here also, the multiplicity of equilibria derived in Proposition 2 translates into the coexistence of one (efficient) equilibrium with no delegation, and one (bad reputation) equilibrium where delegation acts as a remedy to the inefficiency of the DM's interventions. In line with our audience-based approach of delegation, it is worth noting that delegation often gives rise to *blame-shifting*, notably when delegation mandates are broad:[38] politicians that delegate 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 delegation.[39]

[38]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.

[39]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. As an example, 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).

4.4 Exemption

The previous analysis suggests that it is beneficial for the DM to leave a subset of receivers outside the scope of his intervention by empowering a delegate that targets a different audience. In this section, we study how the DM can alternatively narrow down his audience by introducing some differential treatment. Specifically, we allow the decision maker to exempt some receivers, i.e., to insulate them 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).[40]

Let $\mathcal{I}_0$ and $\mathcal{E}_0$ denote the sets of receivers that the DM chooses to influence and to exempt. Since receivers in $\mathcal{E}_0$ choose $y_t(b) = 1 - b$ regardless of the DM's reputation, the objective of the DM in period 1 is to maximize

$$\tilde{V}^k(\mathcal{I}_0, \mathcal{I}_0, m_1) + \frac{1}{1-\delta}\int_{b\in\mathcal{E}_0} [1 - b - \frac{1}{2}(1-b)^2]\, dF(b). \tag{26}$$

$$\text{s.t. } \{\mathcal{I}_0, \mathcal{E}_0\} \text{ is a partition of } \mathcal{B}.$$

Note that we maintain the simplifying assumption that receivers' biases b are uniformly distributed on $[-A, A]$.

Proposition 8. *In equilibrium, both $\mathcal{I}_0^*$ and $\mathcal{E}_0^*$ are intervals. In addition, $\mathcal{E}_0^* \neq \varnothing$: the DM always exempts some receivers. If $\mathcal{I}_0^* \neq \varnothing$, one has*

- $m_1 > 0 \Rightarrow A \in \mathcal{I}_0^*$ *and* $-A \in \mathcal{E}_0^*$
- $m_1 < 0 \Rightarrow -A \in \mathcal{I}_0^*$ *and* $A \in \mathcal{E}_0^*$

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 intrinsically less able to influence. The intuition is as follows: Extreme types generate the strongest loss for the DM (their incentives are very misaligned) and are accordingly those the DM needs to direct his intervention at. However, it is impossible to simultaneously

[40]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 (2018), who analyze the effects of net neutrality regulation on innovation incentives of major content providers.

impact receivers at both extremes of the spectrum for if the average bias in the DM's audience were too moderate extreme receivers would not expect an intervention that compensates their own biases. In order to tilt his target audience towards 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 it is generically difficult to characterize how $\mathcal{I}_0^*$ varies with K, this is possible in the symmetric case where $m_1 = 0$.

Corollary 2. *Suppose* $m_1 = 0$*. If* $K > 1 - \frac{\Sigma}{A^2}$*, then* $\mathcal{I}_0^* = \varnothing$*. If* $K \leq 1 - \frac{\Sigma}{A^2}$*, then* $\mathcal{I}_0^* \neq \varnothing$*, and one then has* $m_1 \in \mathcal{E}_0^*$*. In addition, the audience* $\mathcal{I}_0^*$ *is narrower and less congruent in the less efficient equilibrium than in the more efficient one.*

When K is above some threshold, the DM is better off abstaining from any intervention.[41] Even when K is sufficiently small for $\mathcal{I}_0^*$ to be non-empty, the audience is narrower in the less efficient equilibrium. Intuitively, in the equilibrium where K is higher and reputational incentives lead to more inefficient actions, the benefit of exerting influence on any receiver goes down. Since, as explained above, the DM benefits most from influencing extreme receivers, the first receivers to be exempted are the more moderate ones. This, in turn, implies that the DM's reputational deficit is higher in the less efficient equilibrium. This result suggests that different reputational equilibria are associated with different scopes for government intervention: while moderate expectations warrant (targeted) policy interventions, laissez-faire policies are both caused and justified by anticipations that politicians overreact to their reputational deficits.

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

[41] In particular, no intervention is optimal when reputation is bad. Note that this result is different from the notion of bad reputation, though obviously related: exempted receivers are insulated not only from the DM's action but also from the impact of his type θ_t. By contrast, the myopic DM in the benchmark case of Definition 1 always picks $a_t = 0$ but still influences receivers through θ_t.

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. We also show that multiple institutional forms can arise in equilibrium as a response to the costs of reputation-building. Delegation and centralization improve welfare when reputation is bad, but never when it is good, and exemptions are more beneficial in inferior equilibria. More generally, our approach suggests that strategic audience management can be a key lever in settings where reputation matters.

References

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.

Blinder, A., M. Ehrmann, M. Fratzscher, D. H. Jakob, and J. David-Jan, 2008, "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence," *Journal of Economic Literature*, 46(4), 910–945.

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

Choi, J. P., D.-S. Jeon, and B.-C. Kim, 2018, "Net neutrality, network capacity, and innovation at the edges," *The Journal of Industrial Economics*, 66(1), 172–204.

Cisternas, G., 2017, "Two-sided learning and the ratchet principle," *The Review of Economic Studies*, 85(1), 307–351.

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

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

Eccles, R. G., N. Nohria, and J. D. Berkley, 1992, *Beyond the hype: Rediscovering the essence of management*. Beard Books.

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.

Frenkel, S., 2015, “Repeated interaction and rating inflation: A model of double reputation,” *American Economic Journal: Microeconomics*, 7(1), 250–80.

Gertner, R., R. Gibbons, and D. Scharfstein, 1988, “Simultaneous signalling to the capital and product markets,” *The RAND Journal of Economics*, 19(2), 173–190.

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

Kartik, N., and R. Van Weelden, forthcoming, “Informative Cheap Talk in Elections,” *The Review of Economic Studies*.

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.

Laffont, J.-J., and J. Tirole, 1993, *A theory of incentives in procurement and regulation*. MIT press.

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.

Oates, W. E., 1999, "An Essay on Fiscal Federalism," *Journal of Economic Literature*, 37(3), 1120–1149.

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 (9) to be convex, we need to make sure that

$$2\alpha_2(1-\lambda) - \frac{\gamma}{\delta(1-\lambda)} < 0, \tag{27}$$

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 (9) and (10), the following condition must hold for any pair (m_t^{DM}, m_t):

$$2\alpha_2 \left\{ \lambda m_t + (1-\lambda)[m_t^{DM} - a(m_t, m_t)] \right\} + \alpha_3 m_t^{DM} + \alpha_5 = \left[\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda) \right] a_t(m_t^{DM}, m_t). \tag{28}$$

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 (28). For all (m_t^{DM}, m_t), $(\beta_1, \beta_2, \beta_3)$ must satisfy

$$\begin{aligned} & 2\alpha_2 \left\{ [\lambda - (1-\lambda)(\beta_1+\beta_2)] m_t + (1-\lambda) m_t^{DM} - (1-\lambda)\beta_3 \right\} + \alpha_3 m_t^{DM} + \alpha_5 \\ & = \left[\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda) \right] (\beta_1 m_t^{DM} + \beta_2 m_t + \beta_3). \end{aligned}$$

This gives

$$\beta_1 = \frac{\alpha_3 + 2\alpha_2(1-\lambda)}{\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda)}. \tag{29}$$

$$\beta_2 = \frac{\delta(1-\lambda)}{\gamma}(2\alpha_2 + \alpha_3) - \beta_1 = \frac{\delta(1-\lambda)}{\gamma} 2\alpha_2[\lambda - (1-\lambda)\beta_1] \tag{30}$$

$$\beta_3 = \frac{\delta(1-\lambda)}{\gamma}\alpha_5 \tag{31}$$

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 (7) and (8), where (7) 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:

$$\alpha_1 = \delta[\alpha_1 + \alpha_2(1-\lambda)^2(1+\beta_1)^2 + \alpha_3(1-\lambda)(1+\beta_1)] - \frac{\gamma}{2}\beta_1^2 \tag{32}$$

$$\alpha_2 = -\frac{1}{2}(1+\beta_1+\beta_2)^2 + \delta\alpha_2[\lambda-(1-\lambda)\beta_1]^2 - \frac{\gamma}{2}\beta_2^2 \tag{33}$$

$$\alpha_3 = \delta\{2\alpha_2[\lambda-(1-\lambda)\beta_1](1-\lambda)(1+\beta_1) + \alpha_3[\lambda-(1-\lambda)\beta_1]\} - \gamma\beta_1\beta_2 \tag{34}$$

$$\alpha_4 = \delta\alpha_4 + \delta\alpha_5(1-\lambda)(1+\beta_1) - \gamma\beta_1\beta_3 \tag{35}$$

$$\alpha_5 = (\bar{b}-\beta_3)(1+\beta_1+\beta_2) + \delta\alpha_5[\lambda-(1-\lambda)\beta_1] - \gamma\beta_2\beta_3 \tag{36}$$

$$\alpha_6 = \frac{1-\mathbb{V}(b)}{2} - \frac{1}{2}(\bar{b}-\beta_3)^2 + \delta\left[(\alpha_1+\alpha_2+\alpha_3)\frac{1}{h_\eta} + \alpha_6\right] - \frac{\gamma}{2}\beta_3^2 \tag{37}$$

Notice also that the relations in (29) and (30) can be rewritten as

$$\alpha_2 = \frac{\frac{\gamma}{\delta(1-\lambda)}\beta_2}{2[\lambda-\beta_1(1-\lambda)]} \qquad \text{and} \qquad \alpha_3 = \frac{\gamma}{\delta(1-\lambda)}\left[\beta_1 - \beta_2\frac{(1-\lambda)(1+\beta_1)}{\lambda-\beta_1(1-\lambda)}\right] \tag{38}$$

Using (38) to substitute α_2 and α_3 in the RHS of (33) and (34),

$$\frac{1}{\gamma}\alpha_2 = -\frac{1}{2\gamma}(1+\beta_1+\beta_2)^2 + \frac{1}{2}\beta_2\left[\frac{\lambda}{1-\lambda} - \beta_1\right] - \frac{1}{2}\beta_2^2 \tag{39}$$

$$\frac{1}{\gamma}\alpha_3 = \frac{\lambda}{(1-\lambda)}\beta_1 - \beta_1^2 - \beta_2\beta_1 \tag{40}$$

$2\times(39)+(40)$ yields, using (30),

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

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

$$-\frac{1}{\delta(1-\lambda)}k = -\frac{1}{\gamma}(1-k)^2 - \frac{\lambda}{(1-\lambda)}k - k^2$$

$$\Leftrightarrow \quad \varphi(k) \equiv (1+\gamma)\,k^2 - \left[\frac{\gamma(1-\delta\lambda)}{\delta(1-\lambda)} + 2\right]k + 1 = 0. \tag{41}$$

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 (41) 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 (39) and (40). Rearranging (39),

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

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 (38), (27) 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 (38), that $\alpha_2 < 0$ so that (27) holds. In conclusion, for any k solution to (41) there exist two pairs (β_1, β_2), such that (27), (38), (39) and (40) hold.

In order to fully characterize equilibrium strategies, it only remains to derive β_3. From (36),

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

where the last equality makes uses of (41).

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 $s \geq 1$,

$$m_{t+s} = \lambda^s m_t + (1-\lambda) \sum_{i=0}^{s-1} \lambda^{s-1-i}(\theta_{t+i} + \epsilon_{t+i}) \tag{42}$$

in equilibrium.

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

$$\mathbb{V}_t(m_{t+s}) = (1-\lambda)^2 \left\{ \frac{1}{h} \left(\sum_{i=0}^{s-1} \lambda^i \right)^2 + \frac{1}{h_\varepsilon} \sum_{i=0}^{s-1} \lambda^{2i} + \frac{1}{h_\eta} \sum_{i=1}^{s-1} \left(\sum_{j=0}^{i-1} \lambda^j \right)^2 \right\} \tag{43}$$

Recalling $h = (1-\lambda)h_\eta$ and $h_\varepsilon = \frac{1-\lambda}{\lambda} h_\eta$, and using simple algebra, one can simplify (43) as

$$\mathbb{V}_t(m_{t+s}) = \frac{s}{h_\eta} \text{ for all } s \geq 0. \tag{44}$$

This implies that

$$\sum_{s=0}^{+\infty} \delta^s \mathbb{V}_t(\bar{b} - m_{t+s}) = \frac{\delta}{(1-\delta)^2} \frac{1}{h_\eta} \tag{45}$$

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 conclude

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

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 (41).

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}

$$\lim_{t \to +\infty} \delta^t \mathbb{E}_1[V^k(m_t^{DM}, m_t) | m_1, m_1^{DM}, a^*(.,.)] = 0. \tag{46}$$

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 (44)

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

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

$$m_{t+1} = m_{t+1}^{DM} + [\lambda - (1-\lambda)\beta_1]^t \left(m_1 - m_1^{DM}\right) \tag{47}$$

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 (46) to hold is

$$\delta[\lambda - (1-\lambda)\beta_1]^2 < 1 \tag{48}$$

We have shown that for $k \in \{\overline{k}, \underline{k}\}$, there exists a solution to equations (29) to (37) such that $\alpha_2 < 0$. Then (33) 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 (48) is true and the transversality condition is verified for the equilibrium strategy. This shows in particular that the DM's value function coincides 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 a strategy $a \equiv \{a_t\}_{t>0}$ instead.

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

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

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

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 \left\{ \frac{1-\mathbb{V}(b)}{2} - \frac{[m_t(a) - \bar{b}]^2}{2} - \gamma\frac{a_t^2}{2} \right\} = 0. \tag{49}$$

Since

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

and

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

(49) implies

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1[m_t(a) - \bar{b}]^2 = 0 \tag{50}$$

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

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

Hence,

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

and therefore,

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

Hence (50) that can be written as

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

implies

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 m_t^2(a) = 0. \tag{52}$$

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

$$\lim_{t\to+\infty} \delta^t \mathbb{E}_1 [m_t(a) m_t^{DM}] = 0 \tag{53}$$

Combining (51), (52) and (53)

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

□

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 where the DM plays $a_t = 0$ in each period. 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, $K = 1$:

$$V^0(m_t) = \frac{1}{2(1-\delta)} \left(1 - \mathbb{V}(b) - (\bar{b} - m_t)^2 - \Sigma\right) \tag{54}$$

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 (41) as

$$\tilde{\varphi}(k,z) = (1+\gamma)k^2 - (2+\gamma z)k + 1 = 0 \tag{41'}$$

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 (3.3), 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 $\mathcal{P}$ 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

$$\int_{\mathcal{I}\in\mathcal{P}} \tilde{V}^k(\mathcal{I}, \mathcal{I}, m_1)$$

Using $a_t^*(m_t, \mathcal{I}) = k\left(\mathbb{E}(b|b\in\mathcal{I}) - m_1\right)$, one can write

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

Overall, the DM therefore maximizes

$$\frac{1}{2(1-\delta)} \int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I}) \left\{1 - \mathbb{V}(b|b\in\mathcal{I}) - K[\mathbb{E}(b|b\in\mathcal{I}) - m_1]^2 - K\Sigma\right\} \tag{55}$$

subject to

$$\begin{cases} \displaystyle\int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I}) = 1 & \text{(56a)} \\ \displaystyle\int_{\mathcal{I}\in\mathcal{P}} P(\mathcal{I})\mathbb{E}(b|b\in\mathcal{I}) = \mathbb{E}(b) & \text{(56b)} \end{cases}$$

(56a) reflects the fact that the total audience has mass 1, while (56b) is the Law of Iterated Expectations. After simplification, using (56a) and (56b), 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 $K = 1$, one immediately sees that any partition gives the same payoff. 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 decentralizes decision making at the most granular (i.e., at the individual) level. □

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

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

The present value at date 1 is therefore

$$\tilde{V}^k(\hat{\mathcal{I}}, \mathcal{B}, m_1) = \frac{1}{2(1-\delta)}\left(1 - \mathbb{V}(b) - Km_1^2 - K\Sigma - (1+\gamma)k^2\hat{b}^2 - 2\left(k(1-k) - \gamma k^2\right)\hat{b}m_1\right)$$

This function reaches a maximum at

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

It is clear that, if $\hat{b}^* < -A$ (resp. $\hat{b}^* > A$), one should optimally choose $\hat{b}^* = -A$ (resp.

$\hat{b}^* = A$).

Therefore, the net benefit from delegation is

$$\tilde{V}^k(\hat{\mathcal{I}}^*, \mathcal{B}, m_1) - \tilde{V}^k(\mathcal{B}, \mathcal{B}, m_1) \propto -\frac{1}{2}(1+\gamma)k^2\hat{b}^{*2} - \left(k(1-k) - \gamma k^2\right)\hat{b}^* m_1$$

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

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

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. $\square$

Proof of Lemma 1 Let us denote by $\Pi(\tilde{\mathcal{I}}) \equiv \tilde{V}^k(\tilde{\mathcal{I}}, \tilde{\mathcal{I}}, m_1)$ the function which the delegates wants to maximize.

One has

$$\Pi(\tilde{\mathcal{I}}) \equiv \frac{1}{2(1-\delta)} P(\tilde{\mathcal{I}}) \left(1 - \mathbb{V}(b|b \in \tilde{\mathcal{I}}) - K[\mathbb{E}(b|b \in \tilde{\mathcal{I}}) - m_1]^2 - K\Sigma\right). \quad (57)$$

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

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

By construction, $\mathbb{E}(b|b \in \mathcal{I}') = \mathbb{E}(b|b \in \tilde{\mathcal{I}}^*)$ and $P(\mathcal{I}') = P(\tilde{\mathcal{I}}^*)$, but $\mathbb{V}(b|b \in \mathcal{I}') < \mathbb{V}(b|b \in \tilde{\mathcal{I}}^*)$. This implies $\Pi(\mathcal{I}') - \Pi(\tilde{\mathcal{I}}^*) = \frac{1}{2(1-\delta)} P(\tilde{\mathcal{I}}^*) \left[\mathbb{V}(b|b \in \tilde{\mathcal{I}}^*) - \mathbb{V}(b|b \in \mathcal{I}')\right] > 0$. Therefore, the delegate is strictly better off choosing $\mathcal{I}'$ rather than $\tilde{\mathcal{I}}^*$, and $\tilde{\mathcal{I}}^*$ cannot be the solution of the delegates's problem.

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

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

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 $\mathbb{V}(b|b \in \tilde{\mathcal{I}}^*) = \frac{(\overline{a}-\underline{a})^2}{12} = \frac{A^2P^2}{3}$. Let us also denote $\rho(K) \equiv 1 - K\Sigma$.

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

$$\max_{P \in [0,1],\, \tilde{b} \in [-A(1-P), A(1-P)]} P\left(\rho(K) - \frac{A^2P^2}{3} - K(\tilde{b} - m_1)^2\right) \tag{58}$$

Let us first fix $P \in [0, 1]$ and maximize (58) w.r.t. $\tilde{b} \in [-A(1 - P), A(1 - P)]$.

This gives:

$$\begin{cases} \tilde{b} = m_1 \text{ if } -A(1-P) \leq m_1 \leq A(1-P) \\ \tilde{b} = A(1-P) \text{ if } m_1 > A(1-P) \\ \tilde{b} = -A(1-P) \text{ if } m_1 < -A(1-P). \end{cases}$$

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 $m_1 \in \mathcal{B} = [-A, A]$, then $\tilde{\mathcal{I}}^*$ 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 $g(P) \equiv \rho(K)P - \frac{A^2P^3}{3} - K[A(1 - P) - m_1]^2P$ 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.[42] We conclude that the solution P^* is an nonincreasing function of K.

Notice that the delegate chooses an empty audience (i.e., refuses delegation) 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 delegate is too far away from even the closest receiver 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 lower in the less efficient equilibrium than in the more efficient one. In addition, one has $\tilde{b} = m_1$ as long as $P^* \leq 1 - \frac{m_1}{A}$, and $|\tilde{b} - m_1|$ decreases in K otherwise, for the same reason as in the case $m_1 > A$. We can conclude that the degree of congruence $-|\tilde{b} - m_1|$ is larger in the less efficient equilibrium than in the more efficient equilibrium. □

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

$$\tilde{V}^k(\tilde{\mathcal{I}}^*, \mathcal{B}, m_1) - \tilde{V}^k(\mathcal{B}, \mathcal{B}, m_1) \propto -(1+\gamma)k^2\tilde{b}^{*2} - 2\left(k(1-k) - \gamma k^2\right)\tilde{b}^* m_1 \tag{60}$$

[42] In the way we define the function g (and later h), K and Σ are treated as orthogonal variables. In reality, they are not, because δ and h_η jointly affect K and Σ. However, we are interested here in comparing the optimal audience across equilibria, which means that we compare two values of K for fixed parameter values (hence a fixed Σ), so that this shortcut does not invalidate the comparison we intend to make.

From the proof of Lemma 1, 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 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 (60) > 0.

Let us now consider the case $k \leq \frac{2}{1+\gamma}$. Provided $m_1 \in (-A+1, A-1)$, one always has $\rho(K) = 1 - K\Sigma < (A - m_1)^2$, which implies $\tilde{b}^* = m_1$ in the optimal audience (see the proof of Lemma 1).[43] Then the benefit from delegating (60) becomes (proportional to)

$$-\left((1+\gamma)k^2 + 2\left(k(1-k) - \gamma k^2\right)\right) m_1^2.$$

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

Proof of Proposition 8 The objective of the DM is to choose a partition $\{\mathcal{I}_0, \mathcal{E}_0\}$ of $\mathcal{B}$ that maximizes

$$\tilde{V}^k(\mathcal{I}_0, \mathcal{I}_0, m_1) + \frac{1}{1-\delta}\int_{b\in\mathcal{E}_0}[1 - b - \frac{1}{2}(1-b)^2]\, dF(b),$$

which can be rewritten (up to a constant multiplier)

$$1 - \mathbb{V}(b) + P(\mathcal{I}_0)\left(\mathbb{E}(b|b\in\mathcal{I}_0)^2 - K[\mathbb{E}(b|b\in\mathcal{I}_0) - m_1]^2 - K\Sigma\right) \tag{61}$$

Let $b_0 \equiv \mathbb{E}(b|b \in \mathcal{I}_0)$. A first remark is that it is always optimal to exempt some receivers. Indeed, choosing $\mathcal{I}_0 = \mathcal{B}$ entails $P(\mathcal{B}) = 1$ and $b_0 = \bar{b} = 0$, hence yields a payoff strictly lower than $1 - \mathbb{V}(b)$. Since the DM could secure $1 - \mathbb{V}(b)$ by choosing $\mathcal{I}_0 = \varnothing$, one infers that $\mathcal{E}_0^* \neq \varnothing$.

In order that the DM does not exempt everyone, one needs that there exists at least one value of b_0 such that $b_0^2 - K(b_0 - m_1)^2 - K\Sigma > 0$. If $K < 1$, the function $b_0^2 - K(b_0 - m_1)^2 - K\Sigma$ is convex in b_0, so is maximum either at $b_0 = -A$ or at $b_0 = A$. One also remarks that the maximal value is attained at $b_0 = A$ if $m_1 \geq 0$ and at $b_0 = -A$ if $m_1 \leq 0$.

Therefore, a necessary condition to have a positive mass of non-exempted receivers is

[43] Note that if $\rho(K) < 0$, then $\mathcal{I}^* = \varnothing$, in which case we can impose that $\tilde{b}^* = m_1$. This is innocuous, as we can anyway assume that the DM does not want to leave control to a delegate who has an empty audience.

$A^2 - K(A - |m_1|)^2 - K\Sigma > 0$. Let us assume this is the case from now on.

In the case where $K > 1$, the function $b_0^2 - K(b_0 - m_1)^2 - K\Sigma$ 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 b_0 such that $b_0^2 - K(b_0 - m_1)^2 - K\Sigma > 0$, the DM wants to maximize $P(\mathcal{I}_0)$. It is easy to see that the maximum $P(\mathcal{I}_0)$ compatible with a conditional expectation b_0 is $\frac{A-b_0}{A}$ if $b_0 \geq 0$ and $\frac{A+b_0}{A}$ if $b_0 \leq 0$.

Therefore we are interested in the maximum of

$$f(b_0) \equiv (A - b_0)\left(b_0^2 - K(b_0 - m_1)^2 - K\Sigma\right)$$

on $[0, A]$ and

$$g(b_0) \equiv (A + b_0)\left(b_0^2 - K(b_0 - m_1)^2 - K\Sigma\right)$$

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 $\tilde{b_0} \leq 0$ which maximizes g on $[-A, 0]$. One easily check that $f(-\tilde{b_0}) - g(\tilde{b_0}) = -4K(A + \tilde{b_0})m_1\tilde{b_0} > 0$. Therefore, the global maximum b_0 attainable to the DM must have the same sign as m_1. $\square$

Proof of Corollary 2 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 $(A - b_0)\left((1-K)b_0^2 - K\Sigma\right)$. First, it is easy to see that this function is nonpositive for any $K \geq 1$. In this case, the DM exempts everyone ($\mathcal{I}_0^* = \varnothing$). If $K < 1$, the solution to this problem is $b_0 = \frac{1}{3}\left(A + \sqrt{A^2 + \frac{3K}{1-K}\Sigma}\right)$. It is easy to see that b_0 is nondecreasing in K and that $b_0 > A/2$, which implies that $0 = m_1 \in \mathcal{E}_0^*$. $\square$