Horizontal reputation

Matthieu Bouvard[*], Raphaël Levy[†]

November 2015

Abstract

We consider a career-concerned agent whose reputational reward is higher when perceived as closer to an interior bliss reputation. Career concerns give rise to multiple equilibria characterized by repositioning towards the ideal reputation. A better equilibrium in which repositioning is moderate and reputation-building increases welfare coexists with a less efficient equilibrium where repositioning is extreme, and welfare may be even lower than in the absence of reputation concerns. In the presence of multiple receivers, the inefficiency of the worse equilibrium is exacerbated by the (endogenous) selection of inefficiently narrow and congruent audiences.

[*]McGill University, Desautels Faculty of Management. E-mail: matthieu.bouvard@mcgill.ca.
[†]Mannheim University. E-mail: raphael.levy@uni-mannheim.de.

1 Introduction

The literature has carefully discussed how reputation or career concerns provide implicit incentives in the absence of formal commitment, and how these incentives may either improve or worsen welfare.[1] However, it has almost exclusively focused on reputation in a vertical sense, in that the market has a clear (i.e., monotonic) preference over the actions taken by the career-concerned party, or his type.[2] In many situations, though, quality is horizontal rather than (or on top of being) vertical. A sizeable literature in Industrial Organization and Marketing on horizontal differentiation has underlined the many dimensions along which products or services cannot be objectively ranked (e.g., design, taste, image). Along one such dimension, the "ideal quality" does not necessarily coincide with more quality, and, much the same way a monopolist optimally locates in the middle of the Hotelling segment, a career-concerned agent aims at an interior bliss reputation.[3] The attempt to reach an intermediate reputation may also reflect the desire to compromise between various clienteles with different preferences. For instance, a politician willing to raise money from two lobbies, one in favour of a given reform, and one against, raises more money when competition between lobbies is tough, i.e., when his reputation is such that both lobbies could reasonably expect him to take actions close to their interests.[4]

In line with this idea, a recent literature has begun to explore reputation in horizontal or multi-audience contexts, but in two period environments only.[5] One first important contribution of this paper is to introduce a tractable infinite horizon framework. Several insights emerge from our analysis. First, we derive the existence of two distinct

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

[2]For instance, in Holmström (1999), a manager perceived as more productive commands higher future wages.

[3]For instance, a garment firm selling clothing to eco-friendly consumers should target the right mix of natural fibers to cater to the environmental motivation of customers, and synthetic fabrics, which typically allow for a better performance in terms of strength, warmth or waterproofness.

[4]In the same vein, politicians derive an electoral payoff decreasing in the distance between the policy they are expected to implement and the median voter's preferred policy.

[5]For instance, 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 (2013) establish that a certifier who needs to attract sellers and buyers reaches his maximum profit when his reputation for accuracy is interior. In Shapiro and Skeie (2015), a bank regulator faces ambiguous reputational incentives: a stronger tendency to bail out distressed institutions reassures depositors but induces banks to take excessive risk.

equilibria. Second, we show that welfare may be lower in one equilibrium than in the infinitely-repeated static game, where reputational concerns are absent. This provides a new rationale for "bad reputation" relative to expert models (Morris, 2001; Ely and Välimäki, 2003), or partial observation of actions (Bar-Isaac and Deb, 2014a). Finally, we evidence that the standard intuition that more salient career concerns generate a higher investment in reputation is valid in the two-period case, but is not always robust in the stationary game.

Specifically, we consider a stylized model with non-monotonic career concerns, where the non-monotonicity is driven by the shape of the market demand. We build on Holmström (1999)'s signal jamming model, in which a decision maker tries to influence the market's perception of his type by exerting costly unobservable effort. However, instead of assuming that the market's willingness to pay is increasing in the perceived quality which the decision maker provides (quality is vertical), we assume it to be (quadratic) single-peaked: the decision maker's revenue increases when he supplies a quality perceived as closer to the (interior) market's preferred quality (quality is horizontal). In a stationary environment, we establish the existence of two linear equilibria in which the decision-maker tries to reposition closer to his bliss reputation. In these equilibria, the impact of reputation on incentives is simply captured by a multiplier measuring how responsive the decision maker is to his reputational deficit, that is, the distance between his actual and his bliss reputation. Repositioning is moderate in one equilibrium, but extreme in the other. In the high-responsiveness equilibrium, the decision maker "overshoots", i.e., reacts so much to his reputational deficit that he ends up supplying a quality on average on the other side of the market's bliss point as compared to what he would intrinsically do.[6]

These two equilibria have markedly distinct welfare properties. In the moderate equilibrium, the decision maker becomes more aligned with the market's preferences when he has reputational concerns than when he has none. Although the decision maker fails to manipulate the market's beliefs in equilibrium, reputational concerns help him commit to a course of action closer to efficient, i.e., closer to the course of action he

[6] In politics, such reversals occur when a given policy is more likely to be undertaken by unlikely parties. For instance, market-oriented reforms are often achieved by left-wing parties. The metaphor "Nixon goes to China" has become proverbial to describe such reversals.

would pick under full commitment. Accordingly, reputation in this equilibrium provides a welfare-enhancing (though only imperfect) substitute to commitment. However, in the high-responsiveness equilibrium, the reactivity is excessive, which makes welfare strictly lower than in the moderate equilibrium, and possibly even lower than in the infinitely repeated static game. Actually, the extreme reactivity may get the decision maker to go as far as to supply a quality on average farther from the market's preferred quality than when he behaves myopically. This contrast between the welfare properties of each equilibrium allows to grasp the intuition behind equilibrium multiplicity. Multiplicity arises because non-monotonic career concerns generate intertemporal complementarities between the *current responsiveness* to one's reputational deficit and the *efficiency of future responses*. In the high-responsiveness equilibrium, the strong reactivity to future reputational deficits is inefficient, which makes those deficits more costly to withstand. This in turn raises the current marginal benefit from reaching a better reputation, hence a high responsiveness today. By the same logic, a moderate future responsiveness makes future adjustments more efficient, which justifies current moderation.

Our model naturally extends to the case where the market explicitly consists of multiple audiences with different preferences. We raise two complementary questions regarding the optimal way for the decision-maker to segment the audience. First, within a given market, we study the decision maker's choice to offer the same good/service to each segment or to differentiate his offer. Second, we let the decision maker select the subset of the market he wants to trade with (and therefore the subset he excludes), when constrained to deliver a single quality. In both cases, the optimal strategy depends on the welfare impact of reputation-building. When segmenting the market, the decision maker faces a tradeoff between commitment and flexibility: fragmenting the market allows more personalized service, which boosts the total revenue, but also individualizes pandering, raising the overall costs incurred. The option which is ultimately chosen precisely depends on the relative benefits and costs of reputation: when reputation is welfare-enhancing, the decision maker holds as many reputations as possible, i.e., individualizes his relation with each receiver; when reputation is welfare-decreasing, on the contrary, he builds a global reputation, which allows to commit not to cater to every single receiver in the market. This may take the form of centralized decision-making in politics or in organizations, worldwide brands or undifferentiated advertising. As for the optimal audience selection, our results point to a complementary source of inefficiency: when reputation is less efficient, it is costlier to the DM to trade with more receivers in general, and to trade with less congruent receivers in particular. This results in narrower, more homogenous and more congruent audiences. Overall, our results on heterogenous audiences suggest that the inefficiency of the high-responsiveness equilibrium is reinforced by inefficiencies in the way the decision maker endogenously composes and serves his clientele.

Our paper is most related to Holmström (1999), who models an agent who jams the inference of the market about his productivity by exerting costly unobservable effort. In his model, the reputational payoff is linear, and the equilibrium strategy is accordingly independent of the reputation. By contrast, in our setting with non-monotonic reputational concerns, the equilibrium strategies always depend on the reputation, what is more in a way which generates both equilibrium multiplicity and the possibility of inefficient reputation-building. The paper also relates to Dewatripont, Jewitt, and Tirole (1999a,b) and Bar-Isaac and Deb (2014b), who extend Holmström's setup in different directions: Dewatripont, Jewitt, and Tirole generalize the technology, while Bar-Isaac and Deb generalize the reputational reward function in a multi-audience context. As we do, Dewatripont, Jewitt, and Tirole derive possible equilibrium multiplicity, and Bar-Isaac and Deb obtain repositioning towards the bliss reputation. However, these papers consider two-period environments only. Instead, we model the reputational reward in a way which both captures the horizontal and multi-sided feature of reputation and allows tractability of the infinite horizon analysis. We derive equilibrium multiplicity, but it does not arise from the technology (indeed, the equilibrium is always unique in finite horizon), but from the intertemporal complementarities between future and current incentives, which precisely result from the shape of the reputational reward function. Our focus on whether reputation improves or worsens welfare also relates us to models of "bad reputation" (Morris, 2001; Ely and Välimäki, 2003; Ely, Fudenberg, and Levine, 2008), in which bad reputation springs from the attempt by an "honest" type to separate from types with biased preferences, thereby taking actions detrimental to the market. By contrast, such separating strategies are impossible in our model, as information is and remains symmetric on the equilibrium path. Accordingly, bad reputation does not stem from the decision-maker's reputational incentives possibly going in the wrong direction, as in those papers. Instead, reputational incentives always go in the right direction, but sometimes lead the decision-maker to go too far in that direction, at the expense of welfare.

The remainder of the paper is structured as follows. In Section 2, we present the model and analyze reputation-building in the two-period case. In Section 3 we generalize the analysis to a stationary environment. We derive the existence of multiple equilibria, and examine their welfare and comparative statics properties. In Section 4, we introduce multiple receivers. Section 5 concludes.

2 The model

2.1 Setup

A. Preferences and technology

A long-lived decision maker (later "DM") supplies in each period t a good or service characterized by a positioning (horizontal quality) $x_t \in \mathbb{R}$. In each period, the market consists of a (representative) agent who lives only one period and attaches a value $\alpha - \frac{(x_t-\tilde{x})^2}{2}$ to a product with quality x_t, where $\tilde{x}$ is the agent's (time-invariant) preferred quality, and α is a constant.[7] The quality supplied by the DM is given by $x_t = \theta_t + a_t$, where θ_t is the DM's type in period t, and a_t is an action he chooses. There is moral hazard in that the action a_t is unobservable to the market and costly to the DM, with cost $c(a_t) = \frac{1}{2}\gamma a_t^2$ for all $a_t \in \mathbb{R}$.

B. Information structure

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 $\{\eta_t\}_{t\in\mathbb{N}}$ are i.i.d. normal variables with zero mean and precision h_η. Following the literature on career concerns, we assume that θ_t and η_t are unknown both to the DM and the agent. In addition, x_t is not observable either. However, a noisy signal s_t is observed by all players at the end of

[7] As explained in greater detail later, the preferences of the representative agent can be interpreted as a reduced form for the aggregate preference of heterogenous agents. See Section 4 for a formal analysis.

each period. This signal s_t is such that $s_t = x_t + \varepsilon_t = \theta_t + a_t + \varepsilon_t$, where $\{\varepsilon_t\}_{t\in\mathbb{N}}$ are i.i.d. normal variables with zero mean and precision h_ε. The variables θ_1, $\{\varepsilon_t\}_t$ and $\{\eta_t\}_t$ are mutually independent.

C. The stage game: timing and profit

The timing of the stage game is as follows.

1. The DM posts a price for his good/service.

2. The agent decides whether to buy or not.

3. If the agent does not buy, the game ends. Otherwise, the DM chooses an action a_t.

4. The signal $s_t = \theta_t + a_t + \epsilon_t$ is realized, and becomes public history.

Notice that the specification of the timing implies that the DM is paid in advance, capturing the idea that he cannot commit to charge a price contingent on the realization of s_t. In addition, the DM has all the bargaining power and can extract in each period the (short-lived) agent's expected surplus, which reads:[8]

$$\mathbb{E}[\alpha - \frac{(x_t - \tilde{x})^2}{2}] = \alpha - \frac{(\mathbb{E}(x_t) - \tilde{x})^2}{2} - \frac{\mathbb{V}(x_t - \tilde{x})}{2}$$

We assume that the DM gets infinitely negative utility if the agent does not buy, and is accordingly always willing to charge a price equal to this expected surplus, no matter how negative it may get.[9]

This stage game is infinitely repeated. In each period t, the DM and the new-born agent observe the past history of signals $\{s_\tau\}_{\tau<t}$. Given the normality and independence assumptions, standard Bayesian updating allows to derive that the conditional distribution of the DM's type at any date t is normal with mean m_t and precision h_t. When the

[8]Notice that we assume here that the agent's utility depends on the realization of x_t rather than $x_t + \varepsilon_t$, which is consistent with our interpretation of ε_t as observational noise. It would be equivalently possible to assume that the utility depends on $x_t + \varepsilon_t$, where ε_t would then capture randomness in the "production function". In this case, how much the DM is able to capture would decrease by the variance of ε_t, but since this is a constant, the analysis would be qualitatively unchanged.

[9]This assumption allows to keep the problem analytically tractable by ensuring that the DM's profit function is smooth everywhere. Alternatively, one may assume that the DM is bound to serve the market forever once he has entered it. In this spirit, we allow in section 4.2 the DM to *ex ante* select the pool of consumers he serves and exclude those which he finds unprofitable.

DM plays a_t, the motions of m_t and h_t are given by:[10]

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

and

$$h_{t+1} = \frac{(h_t + h_\varepsilon) h_\eta}{h_t + h_\varepsilon + h_\eta}. \tag{2}$$

We derive the static profit which the DM derives in period t :

$$\pi_t[\mathbb{E}(x_t)] \equiv \alpha - \frac{1}{2h_t} - \frac{(\mathbb{E}(x_t) - \tilde{x})^2}{2} = \alpha - \frac{1}{2h_t} - \frac{(m_t + a_t - \tilde{x})^2}{2} \tag{3}$$

This profit is maximized when the DM provides an expected quality $\mathbb{E}(x_t)$ equal to the preferred quality in the market $\tilde{x}$.[11]

C. Illustration

Our model captures situations in which market demand is maximized for some interior quality. Such single-peaked preferences may stem for instance from the desire to compromise between objectively desirable but possibly antagonistic dimensions (e.g., strength and eco-friendliness of a product, efficiency and equity concerns etc.). Alternatively, they could account for horizontal differentiation. Let us illustrate how the model for instance applies to campaign financing by a lobby. A lobby is willing to finance a politician's campaign all the more as it expects the quality of policies to be closer to its own preferred quality $\tilde{x}$.[12] The financing decision is made upfront, and the lobby cannot condition its contribution on future policies.[13] The ability of the politician to tailor quality to the preference of the lobby depends on his own ability and his action. Ability is

[10]Note that, although it is unobservable, the market knows a_t in equilibrium, and uses it to update beliefs.

[11]This non-monotonicity contrasts with Holmström, where the DM would like $\mathbb{E}(x_t) = m_t + a_t$ to be as high as possible. Except for this crucial difference, our specification is similar to his.

[12]The quality of policies has intertwined vertical and horizontal dimensions. A policy is generically better crafted if there are no or few loopholes, and if it is less likely to be subsequently undone or modified by amendments, courts, supranational authorities, or strikes. On the other hand, lobbies care about how policies actually apply to them, which depends on the existence of specific loopholes or amendments, which distinctly affect their own welfare.

[13]Because it is often impossible for politicians to sign contracts with their various stakeholders, and their reputation is accordingly critical to them, the application of Holmström's career concern setup to politics has been quite popular in political economy and political science (Persson and Tabellini, 2002; Alesina and Tabellini, 2007; Ashworth, 2005; Ashworth, de Mesquita, and Friedenberg, 2013). However, all these papers consider reputation in a vertical sense.

both individual-specific (personal experience, popularity, bargaining position, charisma) and party-specific (experience, ties with the unions or corporations), while the costly action captures the resources spent to reach out to other decision makers and draft a convincing case, but also the possible cost of shaping the reform in a way favorable to a specific constituency.[14] The ability to provide quality is subject to repeated changes (e.g., in political or economic conditions), but also exhibits persistence (party's experience or historical ties, quality of the technocratic support).[15] In this uncertain environment, the need to finance future campaigns leads the politician to distort his action so as to convince the lobby that their interests are congruent. The same logic would apply in the presence of several lobbies with diverging interests, in which case the politician would strive to be as close as possible to the average lobby.

More generally, the single-peakedness of the DM's profit π_t may capture in a stylized manner the need for the DM to strike a balance between several parties with heterogenous preferences. Politics is a byword for the art of managing multiple audiences, but multi-sided reputational concerns are pervasive in many other markets. For instance, platforms operating on two-sided markets or intermediaries need to carefully manage the expectations of the various clienteles they are serving, and reputation is typically instrumental in achieving this. Our model is for instance meant to help us understand to what extent the investigation efforts or the disclosure of sensitive information by the media comes as a response to their reputation. Note that, in this section, we capture the DM's desire for compromise when faced with an heterogenous audience in a reduced form. In Section 4, we show that the shape of reputational payoffs similarly remains single-peaked when we explicitly introduce multiple agents with different bliss points.

2.2 The two-period case

To provide a first intuition on our results, let us start with the analysis of the two-period game. In period 2, the DM has no reputational concerns and selects $a_2^* = 0$ no matter his reputation m_2, hence derives a period 2 profit $\pi_2(m_2)$. Denoting δ the discount factor of

[14]For instance, it may be costly to design a loophole which favors a given clientele without jeopardizing the judicial validity or the public acceptability of the law.

[15]Notice that the fact that θ_t captures the talent of the politician and the environment he is facing rather than his preferences/ideology is consistent with the politician being uncertain about his type.

the DM, and using (1), the equilibrium action in period 1, a_1^*, should satisfy

$$a_1^* \in \underset{a_1}{\operatorname{argmax}}\, \delta \mathbb{E}\pi_2 \left\{ \frac{h_1}{h_1+h_\varepsilon} m_1 + \frac{h_\varepsilon}{h_1+h_\varepsilon} [\theta_1 + \varepsilon_1 + a_1 - a_1^*)] \right\} - c(a_1)$$

Since $\pi_2(.)$ is concave and $c(.)$ is convex, a_1^* is the unique solution of

$$\delta \frac{h_\varepsilon}{h_1+h_\varepsilon} \mathbb{E}\pi_2' \left\{ \frac{h_1}{h_1+h_\varepsilon} m_1 + \frac{h_\varepsilon}{h_1+h_\varepsilon} [\theta_1 + \varepsilon_1] \right\} - c'(a_1^*) = 0$$

$$\Leftrightarrow a_1^* = \frac{\delta h_\varepsilon}{\gamma(h_1+h_\varepsilon)} (\tilde{x} - m_1)$$

Proposition 1 *The two-period game admits a unique equilibrium:* $a_1^* = \kappa_1(\tilde{x} - m_1)$, *where* $\kappa_1 \equiv \frac{\delta h_\varepsilon}{\gamma(h_1+h_\varepsilon)}$.

The equilibrium action of the DM takes a simple and intuitive form: it is the product of the distance between the DM's current reputation m_1 and the agent's bliss point $\tilde{x}$ with a multiplier κ_1, which captures the strength of reputational concern. In other words, the DM tries to "catch up" with what would be his ideal reputation, $\tilde{x}$, the one that perfectly aligns him with the agent. In politics, one may interpret this as pandering to the median voter. In the two-audience interpretation, this suggests that the DM caters to the one of his audiences he is perceived as more remote to.[16] For instance, a newspaper suspected of having its hands tied by the advertising business puts in extra investigation effort, which may ultimately result in disclosing evidence against these corporate interests. On the contrary, when perceived as "too" impartial, reputational concerns provide incentives to cut on investigation efforts, which may eventually result in covering up sensitive information.

Since the DM has no superior information on his type, he is unable to jam the signal in equilibrium because the market has rational expectations and is able to "undo" the impact of a_t on s_t. However, the DM has an actual incentive to match the market's expectations, because the market would make an inference biased at his expense if he does not (as in a "rat race"). Therefore, the DM should take a costly action in the direction of what the market expects him to do, that is, try to reposition in the direction

[16] This result echoes analogous results in Bouvard and Levy (2013); Bar-Isaac and Deb (2014a,b); Shapiro and Skeie (2015). While this result is accordingly not novel, it provides a first intuition useful to understand the analysis in the stationary case, which provides original insights absent in the finite case.

of his bliss reputation; in addition, the magnitude of his reaction is determined by how much the market expects the DM to be willing to spend in order to build a reputation, which is precisely what κ_1 captures. In particular, κ_1 increases with the discount factor δ, with the signal-to-noise ratio $\frac{h_\varepsilon}{h_1+h_\varepsilon}$, which measures how informative the signal s_t is on the DM's type, and decreases with the cost parameter γ.

Because the DM cares about his future rather than his current revenue, the impact of reputation on the DM's welfare is ambiguous. To see this, note that for any multiplier κ, the DM's period 1 profit can be written as follows:

$$\pi_1[m_1 + \kappa(\tilde{x} - m_1)] - \gamma \frac{\kappa^2(\tilde{x} - m_1)^2}{2} = \alpha - \frac{1}{2h_1} - \frac{K}{2}(\tilde{x} - m_1)^2,$$

where

$$K \equiv (1 - \kappa)^2 + \gamma \kappa^2. \tag{4}$$

Since K is U-shaped in κ, an increase in κ first increases welfare, but at some point becomes inefficient. Reputational concerns improve the DM's welfare in equilibrium relative to the myopic case ($\kappa = 0$) if and only if $K < 1 \Leftrightarrow \kappa_1 < \frac{2}{1+\gamma}$, which is equivalent to γ small enough. Reputational concerns can therefore be either beneficial or detrimental to the DM. In particular, an increase in δ or in the signal-to-noise ratio $\frac{h_\varepsilon}{h_1+h_\varepsilon}$, by making career concerns more salient, can make the equilibrium more inefficient.

We now turn to the analysis of the infinite-horizon version of the model, and show that, while the profile of equilibrium actions remains similar, the welfare implications change in a significant way.

3 The stationary case

3.1 Linear Markov equilibrium

In this section, we analyze the asymptotic state of the infinite horizon game, where the precision of agents' 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 s_t. Since there is persistence in the DM's type, information on θ_t increases the precision of the conditional distribution of θ_{t+1}, that is, h_{t+1}. On the other hand, because θ_t changes across periods according to unobservable shocks $\{\eta_t\}_t$, each period brings additional uncertainty, which lowers h_{t+1}. As time goes by, the precision always converges to a steady state value such that these two effects exactly offset each other:

$$h_t \underset{t\to+\infty}{\to} h \text{ with } h = \frac{(h+h_\varepsilon)h_\eta}{h+h_\varepsilon+h_\eta} \Leftrightarrow h = \frac{\sqrt{h_\varepsilon^2+4h_\eta h_\varepsilon}-h_\varepsilon}{2} \tag{5}$$

Focusing on this steady state simplifies the analysis, as the variance of the conditional distributions becomes time-independent: the beliefs about θ_t given any history of the game are fully characterized by the mean of the posterior distribution. We will henceforth refer to the mean of this conditional distribution, m_t, as the (public) reputation of the DM. Since the DM's profit function does not depend on calendar time, one may rewrite the period t profit as

$$\pi[\mathbb{E}(x_t)] \equiv \alpha - \frac{1}{2h} - \frac{(\mathbb{E}(x_t)-\tilde{x})^2}{2} \tag{6}$$

Let a_t^e denote the market's expectation about the DM's action a_t. In equilibrium, this expectation has to be correct but, generically, the motion of the reputation is determined as follows:

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

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

While, in the two-period game, the beliefs of the DM over his own type are irrelevant at $t=2$, these beliefs matter with longer horizons. Indeed, the conditional distributions of the DM's type such as perceived by the DM and the market must coincide on the equilibrium path, but they do differ off path, as a deviation is observed by the DM but not by the market. Therefore, we introduce an additional variable which keeps track of the DM's private beliefs, which we denote m_t^{DM} and call the DM's private reputation. We have:

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

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 market must believe that DM plays $a_t^e = a(m_t, m_t)$.

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 market beliefs about his action a_t^e, the DM chooses the period-t action optimally:

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

ii) the market has rational expectations

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

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

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

Note that strategies describe the DM's behaviour both on and off-path. In particular, condition (9) states that the DM's action is optimal, even following an undetected deviation (i.e., if $m_t^{DM} \neq m_t$).[17]

Let us first remark that, if $\delta = 0$ (reputation has no value) or $\lambda = 1$ (the reputation remains constant), there is a unique equilibrium in which the DM repeatedly plays the static action $a^{static} = 0$. In what follows, we assume that $\delta > 0$ and $\lambda < 1$, and establish that multiple equilibria can coexist, which differ both in their on- and off-path strategies. Since we are interested in the equilibrium predictions, we do not distinguish here between different equilibria with different out-of-equilibrium behavior as long as they do not differ in the profile of actions played on path, and are therefore payoff-equivalent.[18] On the equilibrium path, $m_t^{DM} = m_t$, and, to simplify notation, we denote an equilibrium strategy by $a^*(m_t)$.

[17]Notice in this respect that, as long as the market expects 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.

[18]We refer the reader to the Appendix for the derivation of the full-blown equilibrium.

Proposition 2 *There exist two distinct linear (Markovian) equilibrium strategies:*

$$a^*(m_t) = \kappa(\tilde{x} - m_t), \qquad \textit{with } \kappa \in \{\underline{\kappa}, \overline{\kappa}\} \textit{ and } 0 \leq \underline{\kappa} \leq \overline{\kappa}.$$

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

$$V^{\kappa}(m_t) = \frac{1}{1-\delta}\left(\alpha - \frac{1}{2h} - \frac{(1-\kappa)^2 + \gamma\kappa^2}{2}(\tilde{x} - m_t)^2\right) - \frac{(1-\kappa)^2 + \gamma\kappa^2}{2}\sum_{s=1}^{+\infty}\delta^s \mathbb{V}(m_{t+s})$$

In the Appendix, we establish that $\sum_{s=1}^{+\infty}\delta^s\mathbb{V}(m_{t+s}) = \frac{1}{h_\eta}\frac{\delta}{(1-\delta)^2}$. Therefore, recalling that $K = (1-\kappa)^2 + \gamma\kappa^2$, we re-write V^κ as:

$$V^{\kappa}(m_t) = \frac{1}{1-\delta}\left(\alpha - \frac{1}{2h} - \frac{K}{2}(\tilde{x} - m_t)^2 - \frac{K}{2h_\eta}\frac{\delta}{1-\delta}\right) \tag{12}$$

The full proof is in the Appendix. The intuition for the existence of a linear equilibrium is as follows. Suppose that the value function is quadratic, so that the marginal value function is linear in the future reputation. Since posterior reputations linearly depend on the current action, using (1), the marginal expected benefit of a_t is also linear in the current reputation, from the martingale property of beliefs. Since the marginal cost $c'(a_t)$ is also linear, the optimal action is linear in the current reputation. Finally, both the per-period profit and cost functions $\pi(.)$ and $c(.)$ are quadratic, so the value function is actually quadratic. The end of the proof then consists in identifying the solutions analytically.

Formally, viewed from period t, the DM's expected profit in period $t+i$, given a strategy $a(m_t) = \kappa(\tilde{x} - m_t)$ is

$$\alpha - \frac{1}{2h} - \frac{K}{2}[\tilde{x} - \mathbb{E}_t(m_{t+i})]^2 - \frac{K}{2}\mathbb{V}_t(m_{t+i}). \tag{13}$$

Note that the DM cares not only about the expected value of his future reputation, but also about its dispersion: the variance term in (13) captures the risk that future reputations m_{t+i} end up far away from $\tilde{x}$, which, given the curvature of the profit function π, is costly to the DM. Since a_t affects m_{t+i} in a deterministic way, leaving the variance term $\mathbb{V}_t(m_{t+i})$ unchanged, the DM only cares about the marginal impact of a_t on the expectation term

$\mathbb{E}_t(m_{t+i})$. The impact of a_t on m_{t+i} is linear and corresponds to the weight the agent puts on the period-t signal when updating his beliefs, $1-\lambda$. In turn, m_{t+1} has a persistent effect of magnitude λ^{i-1} on m_{t+i} (see (7)). Overall, the marginal effect of the action a_t on the profit in period $t+i$ is

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

Summing up across periods and using the martingale property of beliefs, the marginal benefit of a_t, that is, the impact of the DM's action on the discounted sum of profits reads

$$(1-\lambda)\sum_{i=1}^{+\infty}\delta^i\lambda^{i-1}K\left[\tilde{x}-\mathbb{E}_t(m_{t+i})\right]=\frac{\delta(1-\lambda)}{1-\delta\lambda}K(\tilde{x}-m_t), \tag{14}$$

while the marginal cost of a_t reads

$$\gamma\kappa(\tilde{x}-m_t). \tag{15}$$

In a stationary equilibrium, the multiplier κ must be the same in every period, meaning that the current κ and the future K must be mutually consistent. Accordingly, κ satisfies a fixed point condition given by the equality of (14) and (15). As illustrated in Figure 1, there are two fixed points, corresponding to two equilibria. In the low responsiveness equilibrium $\underline{\kappa}$, the DM moderately reacts to his reputation: $a^*(m_t)=\underline{\kappa}(\tilde{x}-m_t)$. As a result, the burden of adjustments to catch up future reputational deficits (measured by K) is small. This, in turn, justifies a moderate effort to adjust his reputation today. In the high responsiveness equilibrium $\overline{\kappa}$, the DM overreacts to his reputation: $a^*(m_t)=\overline{\kappa}(\tilde{x}-m_t)$, which makes future misalignments costly. This increases the marginal benefit of investing in his reputation today, hence the high responsiveness.

Notice that equilibrium multiplicity arises only when the horizon is infinite: indeed, multiplicity occurs because different expectations about future behavior generate different behaviors today, which would be impossible with a finite horizon, as there is a unique optimal last-period action. In this case, it is possible to pin down a unique equilibrium strategy by backward induction starting with the final date strategy $\kappa_T=0$.[19]

[19] As one can see from Figure 1, the equilibrium would then converge to $\underline{\kappa}$ as $T\to\infty$. Notice in this respect that both equilibria are stable, in the sense that a small shock to a parameter can never cause a

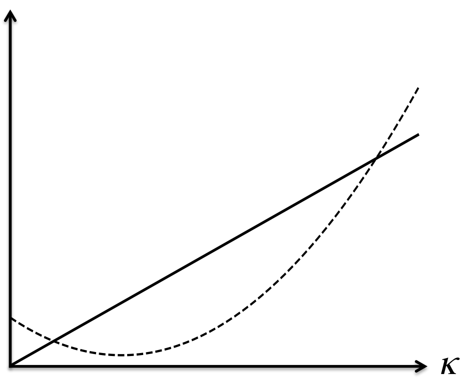

Figure 1: Marginal Cost (solid) and Marginal Benefit (dashed) of reactivity κ.

The intertemporal complementarity between current incentives and future responsiveness is driven by the curvature of the reputational reward π, which makes the present marginal value of reputation depend on the DM's future actions. In this respect, it is instructive to contrast our result with Holmström's, who obtains a unique equilibrium in the stationary case. In both our model and his, increasing a_t today changes the path followed by the future reputations, hence all future revenues. But, in our model, equilibrium actions depend on the reputation, so a change in future reputations affects not only future revenues, but also future costs, unlike in Holmström.

Finally, the source of equilibrium multiplicity in our model differs from Dewatripont, Jewitt, and Tirole (1999b). There, multiplicity arises when type and effort are complements *in the signal* which agents observe (for instance, that signal takes a multiplicative form $s_t = \theta_t a_t + \varepsilon_t$), even in the two-period game. In that case, the signal is more informative about the type θ_t when effort is high, as effort magnifies the weight of the type in the signal relative to the noise. As a result, when effort is expected to be high, agents put more weight on the signal when updating their beliefs, which in turn justifies high effort to manipulate the signal. A similar logic holds for low effort. By contrast, we assume an additive form for the signal $s_t = \theta_t + a_t + \varepsilon_t$, which shuts down this type of complementarity.

switch from an equilibrium with high responsiveness to one with low responsiveness, or vice-versa.

3.2 Welfare: Good and bad reputation

Before deriving the welfare properties of both equilibria, let us first consider the benchmark case where the DM can *ex ante* commit to a course of action. In this case, he maximizes the per-period profit $\pi(m_t + a_t) - \gamma \frac{a_t^2}{2}$, and chooses

$$a^{FB}(m_t) = \frac{1}{1+\gamma}\left(\tilde{x} - m_t\right). \tag{16}$$

Note that the first-best level of reactivity $\kappa^{FB} = \frac{1}{1+\gamma}$ minimizes $K = (\kappa - 1)^2 + \gamma\kappa^2$. Accordingly, K measures the *(in)efficiency* of the equilibrium profile of actions. This is apparent in the expression of the value function V^κ (see Eq. (12)), which is decreasing in K. Since the DM only internalizes the benefit of his actions through their impacts on future reputations and profits, and not their current value, the equilibrium is generically inefficient. However, reputational concerns provide a substitute to commitment, though an imperfect one, and may still improve welfare relative to the case where the DM behaves myopically and repeatedly plays $a^{static} = 0$. We show in the next Proposition that this is not always the case, and that the impact of reputational concerns on welfare is ambiguous.

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

$$V^0(m_t) = \frac{1}{1-\delta}\left(\alpha - \frac{1}{2h} - \frac{1}{2}(\tilde{x} - m_t)^2 - \frac{1}{2h_\eta}\frac{\delta}{1-\delta}\right) \tag{17}$$

Proposition 3 *$V^{\underline{\kappa}}, V^{\overline{\kappa}}$ and V^0 are such that, for any m_t :*

- $V^{\underline{\kappa}}(m_t) > V^0(m_t)$

- $V^{\underline{\kappa}}(m_t) > V^{\overline{\kappa}}(m_t)$

- $V^{\overline{\kappa}}(m_t) > V^0(m_t) \Leftrightarrow \overline{\kappa} < \frac{2}{1+\gamma} \Leftrightarrow \gamma < \frac{\delta(1-\lambda)}{(1-\delta\lambda)+(1-\delta)}$

The first result derives from the fact that $0 \leq \underline{\kappa} \leq \frac{1}{1+\gamma}$. The low-responsiveness equilibrium exhibits the familiar pattern that career concerns alleviate the moral hazard problem in helping the DM commit to take actions closer to the efficient action than in the no-reputation case, but are generically insufficient to reach efficiency.[20] On the contrary,

[20]See Holmström (1999).

the equilibrium $\overline{\kappa}$ features excessive responsiveness: $\overline{\kappa} \geq 1 \geq \frac{1}{1+\gamma}$, and welfare comparisons are a priori unclear. Proposition 3 states that the low-responsiveness equilibrium is always better for the DM than the high-responsiveness one ($V^{\underline{\kappa}} > V^{\overline{\kappa}}$). This ranking of welfare across equilibria is intrinsically related to the complementarity leading to equilibrium multiplicity: in the high-responsiveness equilibrium, the inefficiency of future responses is precisely what induces the DM to be more (inefficiently) reactive today. In addition, when $\overline{\kappa} > \frac{2}{1+\gamma} \Leftrightarrow \gamma > \frac{\delta(1-\lambda)}{(1-\delta\lambda)+(1-\delta)}$, welfare in the high responsiveness equilibrium is lower than in the infinitely repeated static game ($V^{\overline{\kappa}} < V^0$). In this case, the DM not only overreacts to his reputational deficit compared to the first-best action, but the over-reaction is so large than he ends up being worse off than in the no-reputation case. Notice that the welfare loss has two components: first, the fact that the DM is excessively responsive implies that the cost of his actions is excessively high; second, by overshooting, the DM increases the distance between the quality he provides x_t and the bliss point $\tilde{x}$. Actually, the result that reputation decreases welfare may hold even if one abstracts from the costs borne by the DM. Indeed, the expected distance between the quality provided by the DM and the preferred quality in the market reads:[21]

$$\mathbb{E}[(x_t - \tilde{x})^2] = (\kappa - 1)^2(m_t - \tilde{x})^2 + \frac{1}{h}. \tag{18}$$

Therefore, the quality provided is on average farther from the market's preferred quality if the DM has reputational concerns than if he has none if and only if

$$(\kappa - 1)^2 > 1 \Leftrightarrow \kappa > 2$$

Overall, career concerns lead the DM to care about the market, but while they induce moderation in the low-responsiveness equilibrium, they may result in extreme reactions in the high-responsiveness equilibrium. Going back to the political example, our paper therefore establishes that policies are closer to the median voter's preferred policy than under no reputation in the good equilibrium, but may be farther in the bad equilibrium. In addition, since $\overline{\kappa}$ is always larger than 1 and can take arbitrarily large values, (18) suggests that even slight changes in the market's preferred quality $\tilde{x}$ (e.g., a change in the

[21]We consider the average squared distance for simplicity, but it would be equivalent to consider the expected absolute value of the distance between x_t and $\tilde{x}$.

preferences of the median voter) might entail swift adjustments. Finally, notice that the high-responsiveness equilibrium features reversals, in the sense that the expected quality $\mathbb{E}(x_t) = (1 - \overline{\kappa})\, m_t + \overline{\kappa}\tilde{x}$ is decreasing in m_t, i.e., quality becomes negatively correlated with the reputation of the DM. Accordingly, 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. This is reminiscent of the "It takes a Nixon to go to China" effect (Cukierman and Tommasi, 1998). However, reversals have in our model nothing to do with the fact that the unlikely party has more credibility, but instead derive from his desire to build a reputation with respect to constituencies he is perceived as too far from.

3.3 Comparative statics

In both equilibria, the DM tries to reposition in the direction of the bliss point $\tilde{x}$. In this section, we examine how the magnitude of this repositioning depends on the key parameters of the model. We show that, while the direction in which parameter changes affect the DM's reactivity depends on the equilibrium one considers, the (qualitative) impact of these changes on (static) efficiency does not.

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

Proof In the Appendix.

An increase in δ, h_ε or $-h_\eta$ raises the salience of futures payoffs. While a higher δ mechanically makes future payoffs more valuable, a higher precision of the signal s_t (a higher h_ϵ), or a higher variability of θ_t (a lower h_η) induce the agent to put more weight on the latest observation s_t when updating beliefs on θ_{t+1}, which makes future profits more sensitive to the current action. In the 2-period case, this induces the DM to increase his investment in reputation in period 1 (κ_1 increases). This logic does not always hold in a stationary equilibrium. Actually, in the high-responsiveness equilibrium, the DM becomes *less* reactive when his future payoffs become more sensitive to his current action.

To understand this comparative statics result, note that when future actions become more efficient, the future value function becomes flatter (K decreases), which decreases the marginal value of reputation (14): accordingly, the benefit from closing his reputational deficit today is smaller to a DM who will react in a more efficient way to future deficits. Consider first the simpler case of the low-responsiveness equilibrium. Starting from the equilibrium $\underline{\kappa}$, a positive shock to, say, δ increases the marginal benefit of effort and pushes it above the marginal cost (in Figure 2, this corresponds to an upward shift of the dashed curve). In the low-responsiveness equilibrium, the only way to restore equality between marginal cost and benefit is to increase κ. This raises the marginal cost, while lowering the marginal benefit, as a higher κ improves future efficiency (lower K). The effects of an increase in h_ε or a decrease in h_η are analogous. The intuition for the comparative statics in the high-responsiveness equilibrium is slightly more complex because, around $\overline{\kappa}$, both the marginal cost and the marginal benefit are increasing in κ. However, the marginal benefit increases at a faster rate than the marginal cost. Here again, the key factor is that the marginal benefit of effort increases with the inefficiency of future actions. Because of the curvature of the value function, that inefficiency increases at an increasing rate in the region above $\overline{\kappa}$. By contrast, the marginal cost increases at a constant rate. Therefore, following a positive shock to δ, $\overline{\kappa}$ has to decrease for the equilibrium condition to hold again.

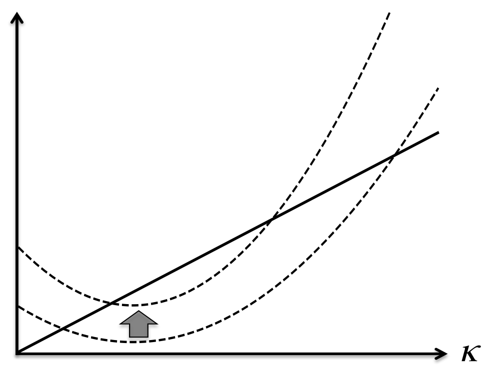

Figure 2: Comparative Statics

A natural question if one wants to derive policy implications from our results has to do with what a benevolent social planner would do if he could design or modify the DM's environment by an appropriate parameter choice. Of course, the answer to these questions depends on the social welfare function one has in mind. A utilitarian social planner has the same value function as the DM, as the latter captures all the surplus, but there are several reasons why the social planner's objective may differ from that of the DM.[22] In what follows, we consider the criterion of static efficiency, that is, how much welfare is lost in each period as compared to the full-commitment solution (the first best).

Let $\Delta W(m_t) \equiv \pi(m_t + a^{FB}(m_t)) - c(a^{FB}(m_t)) - (\pi(m_t + a^*(m_t)) - c(a^*(m_t)))$ denote the difference between the maximum surplus attainable and the equilibrium surplus in a given period. ΔW therefore measures the magnitude of the static inefficiency.

It is easy to show that $\Delta W(m_t) = (K - K^{FB})(\tilde{x} - m_t)^2$, where $K^{FB} = \frac{\gamma}{1+\gamma}$ is the minimum of the function $\gamma\kappa^2 + (1-\kappa)^2$.

Corollary 1 *In both equilibria, for any m_t, the inefficiency $\Delta W(m_t)$ is decreasing in δ and h_ε, and increasing in h_η.*

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 in the two-period case, where the impact of a change in, say, δ is an increase in the magnitude (in absolute value) of effort, regardless of whether the action is below or above the first best. In the stationary equilibrium, an increase in δ always increases welfare.[23] Importantly, while the equilibrium multiplicity limits the predictive power of the model, the fact that the (qualitative) welfare impact of a parameter change does not depend on the equilibrium one considers implies that the normative implications of the model are non-ambiguous.

Corollary 1 implies that the equilibrium is more efficient when δ increases. In particular, one easily shows that $\underline{\kappa}$ tends to κ^{FB} as δ goes to 1, a result reminiscent of folk theorems in repeated games. However, the fact that the inefficient equilibrium also becomes less inefficient when δ increases notably contrasts with the results derived in the literature on bad reputation (Morris, 2001; Ely and Välimäki, 2003; Ely, Fudenberg,

[22]For instance, she may not fully internalize the cost borne by the DM, or have a different discount factor.

[23]This shows that one should be careful in deriving policy implications from the two-period analysis. Indeed, one might be tempted to prescribe a decrease in δ if $\kappa_1 > \frac{1}{1+\gamma}$, but such a decrease worsens the static inefficiency in the stationary equilibrium. Another illustration of this difference is the impact of the cost parameter γ: when γ tends to 0 the action becomes infinitely inefficient in the 2-period equilibrium ($\kappa_1 \to \infty$), while κ tends to $\kappa^{FB} = 1$ in both equilibria of the stationary game.

and Levine, 2008). There, the very desire of the DM to build a reputation results in strategic behavior which ultimately generates welfare losses. The DM is then led to take less efficient actions when he cares more about the future, as reputation is then more salient.[24] On the contrary, the adverse welfare impact of reputation is not driven here by heightened reputational concerns: when the DM cares more about his reputation, the inefficiency actually diminishes.[25] Accordingly, the reason why reputation depresses welfare is essentially different, and actually stems from the distinctive feature of our model, that is, the non-monotonicity of the reputational payoff. Corollary 1 also states that a higher h_ε and a lower h_η improve static efficiency. That is, signals should be very informative in order to make the DM accountable, but should have little relevance to future incarnations of the DM in order not to jeopardize future incentives. The information which is learnt on the DM's type should then be immediately garbled by additional noise.

The combination of the result on δ and h_η provide an interesting insight as to the source of competition in this market. Suppose that the DM is an organization (firm, political party, news outlet), and let us interpret h_η as the degree of turnover of its personnel (managers, political leaders, journalists). The theory predicts that welfare is higher when external competition is soft (high δ : the organization is more likely to operate in the future), but when internal competition is tough (low h_η : turnover within the organization is important). In politics, this suggests that open primaries to select candidates are more efficient than appointment by executive party members, or even the grass roots, as there is more potential for renewal of ideas or emergence of new leaders when candidates are chosen by a larger and more diverse spectrum of voters.

Finally, notice that an efficiency-concerned social planner might also care about dynamic efficiency. In particular, a change in h_η directly affects the DM's welfare through its impact on the variability of future reputations, beside its impact on static efficiency (on K): a decrease in the precision h_η makes future types less predictable, hence increases the volatility of future reputations. Given the curvature of the value function, this is costly to the DM. This effect is reminiscent of Holmström and Ricart I Costa (1986), who show

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

[25] One may find surprising that a higher δ increases welfare after we have stressed that the DM could be better off in the game where he behaves myopically than in the high-responsiveness equilibrium. This is due to the fact that the equilibrium payoff of the DM in the high-responsiveness equilibrium is not continuous at $\delta = 0$: $\lim_{\delta \to 0} \overline{\kappa} = \infty$, while the unique equilibrium involves $a^{static} = 0$ when $\delta = 0$.

that a risk-averse agent with career concerns has an incentive to suppress any source of information that would introduce variability in the market's perception of his skills. Interestingly, in our setup, the welfare impact of such "risk-aversion" depends (through K) on the equilibrium one considers, a feature we will exploit in the next section on endogenous audiences. Overall, this suggests a trade-off between static and dynamic efficiency. On the one hand, a lower h_η improves incentives by preventing the DM from resting on his laurels. On the other hand, such a lower h_η exposes the DM to additional risk, which he is averse to.

4 Multiple receivers: Segmentation and Exclusion

We have assumed so far that the DM interacts with a unique receiver with single-peaked preferences, although we have underlined that these preferences may be a reduced form for the aggregate preference of heterogenous receivers trading with the DM. In this section, we examine the implications of our model in environments where the market explicitly consists of multiple agents (receivers) with similarly shaped preferences but different bliss points, and raise the following two questions:

1. *Optimal Segmentation*: Suppose the DM can arbitrarily partition the audience into independent segments and delegate the provision of the good/service in each specific segment to independent (but otherwise identical) local DMs. What is the optimal way for the DM to segment the total market?

2. *Optimal audience composition / Exclusion*: If the DM can *ex ante* select the receivers with whom he interacts and exclude the others, what are the characteristics of the audience he optimally selects in terms of size, diversity, and congruence?

In order to answer these questions, we extend our model and assume that the market consists of a mass 1 of receivers with period-t utility $\alpha - \frac{1}{2}(x_t - X)^2$ when quality x_t is provided, where X is randomly distributed with c.d.f. F on some support $\mathbb{S} \subset \mathbb{R}$. We also suppose that $h_1 = h$, so that the game is at the steady-state from the very beginning, and that the DM makes the segmentation or the exclusion decision at $t = 1$. The decision, once taken, cannot be adjusted when future reputations are realized, consistent with the idea that the penalty for deserting an ongoing relationship with a segment of the market is prohibitive. Finally, we assume that the cost function $c(.)$ is scaled up or down as a function of the size of the audience: if the audience has a mass μ, the cost function becomes $\mu\gamma\frac{a_t^2}{2}$. This allows to rule out technological effects driven by economies of scale and to focus on the role of reputation only.

Before writing down the optimization program in each of the two problems, let us specify the per-period profit functions and the value functions in the modified version of the game where the DM only faces receivers belonging to a subset $I \subset \mathbb{S}$. Assuming that the DM is still able to perfectly price discriminate, i.e., to extract the expected surplus of each receiver within I, his per-period profit reads

$$\int_{X\in I}\left(\alpha - \frac{1}{2}\mathbb{E}_{x_t}(x_t - X)^2 - \gamma\frac{a_t^2}{2}\right) dF(X)$$

$$= \int_{X\in I}\left(\alpha - \frac{1}{2h} - \frac{1}{2}(m_t + a_t - X)^2 - \gamma\frac{a_t^2}{2}\right) dF(X)$$

After some computations, this equals

$$p(I)\left(\alpha - \frac{1}{2h} - \frac{1}{2}\mathbb{V}(X|X\in I) - \frac{1}{2}\left(m_t + a_t - \mathbb{E}(X|X\in I)\right)^2 - \gamma\frac{a_t^2}{2}\right), \tag{19}$$

where $p(I) \equiv \int_{X\in I} dF(X)$ denotes the mass of receivers in subset I.

Up to an affine transformation, the DM's per-period payoff is the same as in the original model with a single receiver. Therefore, if the DM faces an audience I, he behaves in equilibrium as if he was facing a single receiver with a bliss point $\tilde{x} = \mathbb{E}(X|X \in I)$ and selects $a^*(m_t) = \kappa\left(\mathbb{E}(X|X \in I) - m_t\right)$.[26] We derive the expected discounted value at date 1 :

$$p(I)\left\{\frac{1}{1-\delta}\left(\alpha - \frac{1}{2h} - \frac{1}{2}\mathbb{V}(X|X\in I)\right) - \frac{(\kappa-1)^2 + \gamma\kappa^2}{2}\sum_{i=0}^{+\infty}\delta^i\mathbb{E}_1(\mathbb{E}(X|X\in I) - m_{i+1})^2\right\}$$

[26] In deriving this, we implicitly assume that either I has positive mass or I is a singleton. To avoid technical complications, we disregard here the case of zero-mass subsets with more than two elements, but one could easy show that it can never be strictly optimal for the DM to have a segment of this kind in his clientele.

Recalling that $K = (\kappa - 1)^2 + \gamma\kappa^2$ and $\sum_{i=0}^{+\infty} \delta^i \mathbb{V}_1(m_{i+1}) = \frac{\delta}{(1-\delta)^2} \frac{1}{h_\eta}$, we derive the value for a DM with reputation m_1 :

$$\frac{1}{2(1-\delta)} p(I) \left\{ 2\alpha - \frac{1}{h} - \mathbb{V}(X|X \in I) - K[\mathbb{E}(X|X \in I) - m_1]^2 - \frac{K}{h_\eta} \frac{\delta}{(1-\delta)} \right\} \tag{20}$$

Notice that, if $\mathbb{S} = \{\tilde{x}\}$ as in Section 2, (20) becomes exactly (12).

4.1 Optimal Segmentation

Until now, we have assumed that the DM supplies the same quality to all receivers. In this section, we allow the DM to segment the market. We interpret segmentation as the problem of a "central" organization which may allocate a different "local" decision-maker to each segment of the market.[27] Formally, the DM can arbitrarily partition the support of types $\mathbb{S}$ and supply a different quality to each segment. Let $\mathbb{P}$ denote a partition and I an element of the partition $\mathbb{P}$. There is a continuum of ex ante similar local decision-makers indexed by i, with prior type θ_1^i, where θ_1^i are i.i.d. Normally distributed with mean m_1 and variance $\frac{1}{h}$. The central DM assigns to each element I of the partition a local DM i, who selects an action a_t^i in each period t. All segments are informationally independent in the sense that the shocks η_t^i to types θ_t^i, as well as the noises ε_t^i of the public signals about the qualities x_t^i are independent across DMs. As a consequence, the reputation of one local DM does not provide any information on the types of other DMs (no learning across segments). Finally, we assume no agency friction within the organization, so that each local DM simply maximizes the profit he contributes to the central organization, i.e., the discounted sum of revenues extracted from receivers within I minus the costs he incurs.

Given the informational independence between segments, we derive that the expected value which the DM derives from a partition $\mathbb{P}$ equals

$$\frac{1}{2(1-\delta)} \int_{I \in \mathbb{P}} p(I) \left\{ 2\alpha - \frac{1}{h} - \mathbb{V}(X|X \in I) - K[\mathbb{E}(X|X \in I) - m_1]^2 - \frac{K}{h_\eta} \frac{\delta}{(1-\delta)} \right\} \tag{21}$$

[27] For instance, a firm can decide to sell different goods under the same or different brand names. Similarly, an organization can choose to grant more or less authority or autonomy to local decision-makers.

The DM should pick the partition which maximizes (21), subject to

$$\int_{I\in\mathbb{P}} p(I) = 1 \tag{22}$$

$$\int_{I\in\mathbb{P}} p(I)\mathbb{E}(X|X\in I) = \mathbb{E}(X) \tag{23}$$

(22) reflects the fact that the total audience has mass 1, while (23) is the Law of Iterated Expectations. After simplification, using (22) and (23), the DM maximizes

$$2\alpha - \frac{1}{h} - K\left(\mathbb{E}(X) - m_1\right)^2 - K\mathbb{V}(X) - \frac{K}{h_\eta}\frac{\delta}{(1-\delta)} - (1-K)\int_{I\in\mathbb{P}} p(I)\mathbb{V}(X|X\in I)$$

The solution to this problem is bang bang and only depends on whether K is larger or smaller than 1. If $K=1$, the composition of the audience is irrelevant. Intuitively, the payoff the DM gets is the same as when $\kappa=0$, i.e., when he plays $a_t=0$ in each period regardless of the average preference in the market, a profile of actions which makes the composition of the audience indifferent. If $K\neq 1$, the choice of the partition only affects the DM's profit through $\int_{I\in\mathbb{P}} p(I)\mathbb{V}(X|X\in I) = \mathbb{E}[\mathbb{V}(X|X\in I)]$. It is easy to see, using the law of total variance, that $0\leq \mathbb{E}[\mathbb{V}(X|X\in I)]\leq \mathbb{V}(X)$, with $\mathbb{E}[\mathbb{V}(X|X\in I)]=\mathbb{V}(X)$ when the partition consists of a single element $I=\mathbb{S}$, and $E[\mathbb{V}(X|X\in I)]=0$ when each element I of the partition is a singleton. If $K>1$, the DM should maximize $\mathbb{E}[\mathbb{V}(X|X\in I)]$, and then offers a single quality to all receivers. On the contrary, if $K<1$, he should minimize $\mathbb{E}[\mathbb{V}(X|X\in I)]$ and then builds individualized reputations with each of the receivers. Notice that the result that the optimal segmentation is bang bang simplifies the analysis, as one may then ignore the question of which receivers the DM would like to pool together: there is either full pooling of receivers, or no pooling at all.[28] Since how K compares to 1 is precisely what determines whether reputation improves welfare, we derive the following result:

Proposition 5 *The DM customizes the quality he offers to each receiver in the market whenever reputation increases welfare, that is, when $V^\kappa > V^0$. Otherwise, he supplies the*

[28]It is clear that introducing some cost (or benefit) of having more fragmented audiences, for instance (dis)economies of scale, would possibly lead to an interior solution, and hence a smoother relationship between efficiency of reputation and optimal segmentation. In that case, the result of Lemma 1 in the next subsection suggests that that receivers should be optimally pooled in intervals.

same quality to all receivers in the market.

The intuition for this result is as follows: when segmenting the audience, the DM faces a tradeoff between commitment and flexibility. A more fragmented audience allows the DM to provide tailored quality to each segment, thus raising the revenue of the DM. Indeed, quadratic preferences are here akin to risk aversion, and imply that the DM has to forgo a larger (risk) premium when the variance of the receivers' preferences in the audience increases. However, with a more fragmented audience, the DM cares about his reputation vis-à-vis each of the specific segments, and there must be some segments in the market which the DM has to cater to more than if he only cared about the average receiver in the whole market. Put differently, dealing with a large unique audience provides some commitment for the DM not to pander to every single receiver. Since κ does not depend on the composition of the audience, fragmentation duplicates the benefits of reputation by the same order as it duplicates costs. Therefore, it is precisely when reputation is welfare-decreasing, that is, when the cost of reputation exceeds the benefits, that the value of commitment outweighs the value of flexibility, and that the DM should content himself of a single reputation.[29] Notice also that the benefit from choosing the optimal partition increases with the heterogeneity of preferences in the market measured by $\mathbb{V}(X)$: the gain from providing individualized services is (linearly) increasing in $\mathbb{V}(X)$, while the cost of catering to each receiver is (linearly) decreasing in $\mathbb{V}(X)$, and is weighted by the parameter measuring the inefficiency of reputation K.

A direct consequence of Proposition 5 is that, under full commitment, the DM would always provide fully personalized quality ($K^{FB} = \frac{\gamma}{1+\gamma} < 1$). This suggests an additional source of inefficiency. When $V^{\overline{\kappa}} < V^0$, not only does the DM excessively reacts, but he chooses to maintain a unique reputation, while efficiency would prescribe personalized service and reputations. In other words, the suboptimal fragmentation decision exacerbates the adverse welfare impact of reputation due to the overreaction.

Finally, let us remark that one may also interpret the choice of segmentation as a choice of communication mode: the DM can decide to communicate directly with the whole audience, or to delegate communication to independent identical representatives.

[29] Notice that we require here all local DMs to be ex ante identical. It is obvious that the existence of DMs with different prior reputations would give an extra edge to individualized reputations, as it would then be possible to allocate DMs to a segment they are a priori close to.

Under delegation, each segment of the audience is allocated a specific interlocutor, hence feels more special, which boosts their willingness to listen. However, delegation also increases the total cost of catering, which justifies a preference for global communication when reputation entails excessive demagogy.

Proposition 5 suggests a relationship between marketing strategies of multi-product or multi-national firms and their reputation costs: when reputation costs are low enough, firms should offer more differentiated products. This could take the form of wider brand portfolios, with for instance country-specific brands or advertising campaigns etc; conversely, firms facing high costs of reputation-building should try and manage a unique global reputation, by selling for instance different goods under the same brand, or by having uniform advertising campaigns. In politics, centralization provides commitment not to pander to each of the local constituencies when pandering costs are high, while delegation to local decision-makers (decentralization) dominates when pork-barrels are not too much of a concern. This suggests a relationship between centralization and efficiency where centralization is not a cause, but a byproduct of inefficiencies in policy-making.[30]

4.2 Optimal audience composition

In this section, we address the complementary question of the optimal audience the DM would select if he could decide to "exclude" some receivers from the market (that is, to commit not to extract surplus from them). Formally, we let the DM choose one single subset $I \subset \mathbb{S}$ of receivers he wants to interact with. To make things simple, we assume here that X is uniformly distributed on $[-1, 1]$. We look for the subset I^* which maximizes the expected discounted value of the DM. Formally, $I^* \in \underset{I}{\operatorname{argmax}}\, \Pi(I)$, where

$$\Pi(I) \equiv p(I) \left\{ 2\alpha - \frac{1}{h} - \mathbb{V}(X|X \in I) - K[\mathbb{E}(X|X \in I) - m_1]^2 - \frac{K}{h_\eta} \frac{\delta}{(1-\delta)} \right\} \tag{24}$$

The choice of I impacts the expected payoff Π along three dimensions. First, expanding I increases the mass of receivers $p(I)$ from which the DM collects revenues. Second, expanding I increases the dispersion of preferences in the audience $\mathbb{V}(X|X \in I)$, which

[30] In relation with this issue, an important literature in public finance has focused on the efficiency impact of fiscal federalism. See for instance Oates (1999).

decreases the total revenue of the DM, as service is less personalized. These two effects go in opposite directions, and the DM has an incentive to increase the size of the audience as long as the profit derived from trading with the marginal receiver is positive. This tradeoff arises even holding the distance between the DM's reputation and the average bliss point in the audience fixed. When deciding on I, the DM should in addition take into account how aligned he is going to end up with the average receiver in the audience, because the degree of congruence with the audience determines the equilibrium profile of actions and, hence, his equilibrium payoff.

Before characterizing I^*, let us make the assumption that $2\alpha > \frac{1}{h}$. If this assumption does not hold, $\Pi(I) < 0$ for all I. This implies that $I^* = \varnothing$ (or I^* has zero mass) for any m_1 and any K. We first derive the following lemma:

Lemma 1 *If I^* has positive mass, then I^* must be an interval.*

Proof Suppose that I has positive mass but is not convex. Consider the alternative interval

$$I' \equiv (\mathbb{E}(X|X \in I) - \frac{p(I)}{2}, \mathbb{E}(X|X \in I) + \frac{p(I)}{2}).$$

By construction, $\mathbb{E}(X|X \in I') = \mathbb{E}(X|X \in I)$ and $p(I') = p(I)$, but $\mathbb{V}(X|X \in I') < \mathbb{V}(X|X \in I)$. This implies $\Pi(I') - \Pi(I) = p(I)\left[\mathbb{V}(X|X \in I) - \mathbb{V}(X|X \in I')\right] > 0$. Therefore, the DM is strictly better off choosing I' rather than I. □

Intuitively, if I is not an interval, one can always reshuffle some mass from the extremes to the center, without changing neither the total mass of I nor the conditional expectation in I. Such a change decreases the mass of receivers with extreme bliss points, hence the "risk-premium" which the DM has to leave on the table. We now derive the following Proposition.

Proposition 6 *The DM selects a wider and less congruent audience when the equilibrium is more efficient: $p(I^*)$ and $|\mathbb{E}(X|X \in I^*) - m_1|$ decrease in K.*

Proof In the Appendix.

Inspecting (24), one sees that the impact of K on the optimal audience is two-fold. First, as captured by the last term in (24), each additional receiver generates lower profits to the DM in a less efficient equilibrium, i.e., when K increases. As a result, when expanding the audience, the DM reaches more quickly the point where the marginal receiver becomes unprofitable. Second, in a less efficient equilibrium, the expected loss stemming from the incongruence between the DM and the audience has a stronger (negative) impact on the DM's payoff. This effect is captured by the second term in (24). As a consequence, the audience should be wider (hence more dispersed) and less congruent when reputation is more efficient.[31] Combining the results of Propositions 3 and 6, we derive the following corollary:

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

In terms of welfare, the result of Proposition 6 is contrasted. On the one hand, there are positive effects from exclusion: first, and quite trivially, it may be inefficient to serve the whole audience, even regardless of the initial reputation of the DM, simply because of the DM's implicit "risk aversion", so that some exclusion is warranted from a welfare perspective. Second, selection curbs the "over-reaction" problem: the fact that the DM chooses a more congruent audience in a less efficient equilibrium partly corrects the inefficiency, which does increase with incongruence. On the other hand, the result points to a possible other source of inefficiency, namely that less efficient equilibria may be more prone to unwarranted exclusion when the DM does not internalize social welfare. This may happen notably when excluded receivers still derive utility from the good (for instance, in politics, non-targeted constituencies and minorities still benefit or suffer from the policies implemented). In this case, the endogenous audience choice creates an alignment between the DM's preference and the average preference in the selected audience, but at the expense of the rest of the population.

Overall, Proposition 6 suggests that the potential cost of being trapped in an inefficient equilibrium takes not only the form of costly over-reactions, but also of narrower and less

[31] As should be clear from the intuitions, nothing here particularly relies on the fact that the distribution of receivers' tastes is uniform. This assumption only simplifies the analysis by ensuring that the second-order conditions of the problem are satisfied. More generally, one would need to make assumptions on the distribution to ensure the convexity of the problem, which we do not derive here to keep the extension simple.

diverse audiences. In markets where reputation is more costly to maintain (low discount factor, low renewal of incumbents), we should expect to see niche markets rather than mass markets. For instance, suppose that competitive threats increase, say because of a decrease in entry costs, which, in the model, we may interpret as a decrease in δ. Our results predict that such an increase in competitive threats should result in narrower clienteles. This is consistent with the fact that the development of the Internet has seen the entry in the market for news of many small "niche" information providers which target a very specific audience (see, e.g., Stroud (2011)).[32] In politics, one should expect politicians in countries or regimes where policy-making is less efficient to rely more on smaller constituencies. In line with this prediction, it is interesting to notice that dictatorial or authoritarian regimes, which the model would predict as being less efficient (low renewal of politicians and ideas, little accountability), are indeed characterized by the presence of a narrow and highly patronized base of support (Acemoglu, Robinson, and Verdier (2004); Acemoglu and Robinson (2008)).

5 Conclusion

This paper explores the question of the impact of career concerns in environments where reputation is horizontal and the reward from reputation non-monotonic. As compared to the vertical reputation case *à la* Holmström (1999), we establish new results: first, the non-monotonicity in the reward function creates an intertemporal complementarity between future and current incentives. This complementarity results in equilibrium multiplicity. In a first equilibrium, reputation provides a good but imperfect substitute to palliate the lack of commitment, that is, reputation-building is insufficient to reach the first-best but improves welfare over the benchmark of the infinitely repeated static game. But this equilibrium coexists with another equilibrium characterized by "overshooting," where the decision maker takes extreme actions in order to maintain a reputation. This equilibrium is less efficient than the first one, and may even be less efficient than when the decision

[32]Notice that, in the presence of competing DMs, the endogenous choice of audience would be akin to an optimal location choice problem *à la* Hotelling, with different firms serving different segments to soften competition. If potential entrants have sufficiently distinct reputations, and the cost of reputation-building is large enough, one may even conjecture that the market could consist of several monopolists, each of which operates on a distinct segment of the total demand.

maker has no reputational concerns. We also show that the inefficiency of the overreaction equilibrium is dampened when the decision maker can strategically shape the audience he his facing. The additional inefficiency takes the form of excessively narrow and congruent audiences, and the provision of insufficiently personalized service.

References

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

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

Alesina, A., and G. Tabellini, 2007, "Bureaucrats or politicians? Part I: a single policy task," *The American Economic Review*, 97(1), 169–179.

Ashworth, S., 2005, "Reputational dynamics and political careers," *Journal of Law, Economics, and Organization*, 21(2), 441–466.

Ashworth, S., E. B. de Mesquita, and A. Friedenberg, 2013, "Accountability traps," working paper.

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.

Benabou, R., and G. Laroque, 1992, "Using privileged information to manipulate markets: Insiders, gurus, and credibility," *Quarterly Journal of Economics*, 107(3), 921–958.

Bouvard, M., and R. Levy, 2013, "Two-sided reputation in certification markets," working paper.

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, 1999a, "The Economics of Career Concerns: Part 1," *Review of Economic Studies*, 66, 183–98.

——— , 1999b, "The Economics of Career Concerns: Part 2," *Review of Economic Studies*, 66, 199–217.

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.

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

Holmström, B., and J. Ricart I Costa, 1986, "Managerial Incentives and Capital Management," *The Quarterly Journal of Economics*, 101(4), 835–860.

Kreps, D., and R. Wilson, 1982, "Reputation and imperfect information," *Journal of economic theory*, 27(2), 253–279.

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

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

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

Persson, T., and G. E. Tabellini, 2002, *Political economics: explaining economic policy*. MIT press.

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

Stroud, N. J., 2011, *Niche News: The politics of news choice*. Oxford University Press.

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

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

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

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

In order to satisfy the equilibrium conditions (9) and (10), the following condition must hold for any pair (m_t^{DM}, m_t):

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

Given $V(.,.)$, there exists a unique linear strategy, $a_t(m_t^{DM}, m_t) = b_1 m_t^{DM} + b_2 m_t + b_3$, which satisfies (26). (b_1, b_2, b_3) must satisfy

$$\begin{aligned} & 2\alpha_2 \left\{ [\lambda - (1-\lambda)(b_1+b_2)] m_t + (1-\lambda) m_t^{DM} - (1-\lambda) b_3 \right\} + \alpha_3 m_t^{DM} + \alpha_5 \\ & = \left[\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda) \right] (b_1 m_t^{DM} + b_2 m_t + b_3) \end{aligned}$$

for all (m_t^{DM}, m_t).

This gives

$$b_1 = \frac{\alpha_3 + 2\alpha_2(1-\lambda)}{\frac{\gamma}{\delta(1-\lambda)} - 2\alpha_2(1-\lambda)}. \tag{27}$$

$$b_2 = \frac{\delta(1-\lambda)}{\gamma}(2\alpha_2 + \alpha_3) - b_1 = \frac{\delta(1-\lambda)}{\gamma} 2\alpha_2[\lambda - (1-\lambda)b_1] \tag{28}$$

$$b_3 = \frac{\delta(1-\lambda)}{\gamma}\alpha_5 \tag{29}$$

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

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

We will also make use of the following expectations, derived using (7) and (8), where (7) is rewritten as $m_{t+1} = \lambda m_t + (1-\lambda)\left[\theta_t+\varepsilon_t+b_1(m_t^{DM}-m_t)\right]$.

$$\begin{aligned}
\mathbb{E}(m_{t+1}^{DM}) &= m_t^{DM}\\
\mathbb{E}(m_{t+1}) &= [\lambda-(1-\lambda)b_1]m_t + (1-\lambda)(1+b_1)m_t^{DM}\\
\mathbb{E}[(m_{t+1}^{DM})^2] &= m_t^2 + \frac{h_\varepsilon}{h(h+h_\varepsilon)}\\
\mathbb{E}(m_{t+1}^2) &= [\lambda-(1-\lambda)b_1]^2 m_t^2 + (1-\lambda)^2(1+b_1)^2(m_t^{DM})^2\\
&\quad +2[\lambda-(1-\lambda)b_1](1-\lambda)(1+b_1)m_t^{DM}m_t + \frac{1}{h_\eta}\\
\mathbb{E}(m_{t+1}^{DM}m_{t+1}) &= (1-\lambda)(1+b_1)(m_t^{DM})^2 + [\lambda-(1-\lambda)b_1]m_t^{DM}m_t + \frac{1}{h_\eta}
\end{aligned}$$

Since all the previous terms are quadratic in (m_t^{DM}, m_t), $\pi(.)$ and $c(.)$ are 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:

$$\begin{aligned}
\pi[m_t + a_t(m_t,m_t)] &= \alpha - \frac{1}{2h} - \frac{1}{2}[(1+b_1+b_2)m_t + b_3 - \tilde{x}]^2\\
&= \alpha - \frac{1}{2h} - \frac{1}{2}(1+b_1+b_2)^2 m_t^2 + (\tilde{x}-b_3)(1+b_1+b_2)m_t - \frac{1}{2}(\tilde{x}-b_3)^2
\end{aligned}$$

$$\begin{aligned}
\gamma\frac{a_t^2(m_t^{DM},m_t)}{2} &= \gamma\frac{(b_1 m_t^{DM} + b_2 m_t + b_3)^2}{2}\\
&= \frac{\gamma}{2}[b_1^2(m_t^{DM})^2 + b_2^2 m_t^2 + 2b_1 b_2 m_t^{DM} m_t + 2b_1 b_3 m_t^{DM} + 2b_2 b_3 m_t + b_3^2]
\end{aligned}$$

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

$$\alpha_1 = \delta[\alpha_1 + \alpha_2(1-\lambda)^2(1+b_1)^2 + \alpha_3(1-\lambda)(1+b_1)] - \frac{\gamma}{2}b_1^2 \tag{30}$$

$$\alpha_2 = -\frac{1}{2}(1+b_1+b_2)^2 + \delta\alpha_2[\lambda-(1-\lambda)b_1]^2 - \frac{\gamma}{2}b_2^2 \tag{31}$$

$$\alpha_3 = \delta\{2\alpha_2[\lambda-(1-\lambda)b_1](1-\lambda)(1+b_1) + \alpha_3[\lambda-(1-\lambda)b_1]\} - \gamma b_1 b_2 \tag{32}$$

$$\alpha_4 = \delta\alpha_4 + \delta\alpha_5(1-\lambda)(1+b_1) - \gamma b_1 b_3 \tag{33}$$

$$\alpha_5 = (\tilde{x}-b_3)(1+b_1+b_2) + \delta\alpha_5[\lambda-(1-\lambda)b_1] - \gamma b_2 b_3 \tag{34}$$

$$\alpha_6 = \alpha - \frac{1}{2h} - \frac{1}{2}(\tilde{x}-b_3)^2 + \delta\left[(\alpha_1+\alpha_2+\alpha_3)\frac{1}{h_\eta} + \alpha_6\right] - \frac{\gamma}{2}b_3^2 \tag{35}$$

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

$$\alpha_2 = \frac{\frac{\gamma}{\delta(1-\lambda)}b_2}{2[\lambda - b_1(1-\lambda)]} \qquad \text{and} \qquad \alpha_3 = \frac{\gamma}{\delta(1-\lambda)}\left[b_1 - b_2\frac{(1-\lambda)(1+b_1)}{\lambda - b_1(1-\lambda)}\right] \tag{36}$$

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

$$\frac{1}{\gamma}\alpha_2 = -\frac{1}{2\gamma}(1+b_1+b_2)^2 + \frac{1}{2}b_2\left[\frac{\lambda}{1-\lambda} - b_1\right] - \frac{1}{2}b_2^2 \tag{37}$$

$$\frac{1}{\gamma}\alpha_3 = \frac{\lambda}{(1-\lambda)}b_1 - b_1^2 - b_2 b_1 \tag{38}$$

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

$$\begin{aligned}\frac{1}{\delta(1-\lambda)}(b_1+b_2) &= -\frac{1}{\gamma}(1+b_1+b_2)^2 + b_2\left[\frac{\lambda}{1-\lambda} - b_1\right] - b_2^2 + \frac{\lambda}{(1-\lambda)}b_1 - b_1^2 - b_2 b_1 \\ &= -\frac{1}{\gamma}(1+b_1+b_2)^2 + \frac{\lambda}{1-\lambda}(b_1+b_2) - (b_1+b_2)^2\end{aligned}$$

Let $\kappa \equiv -(b_1+b_2)$.

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

$$\Leftrightarrow \quad F(\kappa) \equiv (1+\gamma)\,\kappa^2 - \left[\frac{\gamma(1-\delta\lambda)}{\delta(1-\lambda)} + 2\right]\kappa + 1 = 0. \tag{39}$$

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

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

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

Let us now check that there exist b_1 and b_2 solutions to (37) and (38). Rearranging (37),

$$\begin{aligned}
\frac{b_2}{\delta(1-\lambda)[\lambda - b_1(1-\lambda)]} &= -\frac{1}{\gamma}(1+b_1+b_2)^2 + b_2\left[\frac{\lambda}{1-\lambda} - b_1\right] - b_2^2 \\
\Leftrightarrow \quad \frac{b_2}{\delta(1-\lambda)^2} &= \left[-\frac{1}{\gamma}(1-\kappa)^2 + b_2\left(\frac{\lambda}{1-\lambda}+\kappa\right)\right]\left(\frac{\lambda}{1-\lambda}+\kappa+b_2\right) \\
\Leftrightarrow \quad & \left(\frac{\lambda}{1-\lambda}+\kappa\right)b_2^2 + \left[-\frac{1}{\gamma}(1-\kappa)^2 + \left(\frac{\lambda}{1-\lambda}+\kappa\right)^2 - \frac{1}{\delta(1-\lambda)^2}\right]b_2 \\
& -\frac{1}{\gamma}(1-\kappa)^2\left(\frac{\lambda}{1-\lambda}+\kappa\right) = 0
\end{aligned}$$

Letting $G(b_2)$ denote the polynomial in the last line and remembering that $\kappa > 0$, we derive that $G(.)$ has two roots of opposite signs. Consider the positive root first. $\kappa > 0$ and $b_2 > 0$ implies $b_1 < 0$. Using this and (36), (25) is equivalent to $-(1-\lambda)\kappa < \lambda$ which is always true. Turn now to the negative root of $G(.)$. $G[-\kappa - \lambda/(1-\lambda)] > 0$ implies $-\kappa - \lambda/(1-\lambda) < b_2$ and therefore $\lambda - (1-\lambda)b_1 > 0$. This implies in turn, from (36), that $\alpha_2 < 0$ so that (25) holds. In conclusion, for any κ solution to (39) there exist two pairs (b_1, b_2), such that (25), (36), (37) and (38) hold.

In order to fully characterize equilibrium strategies, it only remains to derive b_3. From (34),

$$\begin{aligned}
\alpha_5\{1-\delta[\lambda-(1-\lambda)b_1]\} &= (\tilde{x}-b_3)(1+b_1+b_2) - \gamma b_2 b_3 \\
\Leftrightarrow \frac{\gamma - \delta\gamma[\lambda-(1-\lambda)b_1]}{\delta(1-\lambda)}b_3 &= (\tilde{x}-b_3)(1+b_1+b_2) - \gamma b_2 b_3 \\
\Leftrightarrow \left[\gamma\frac{1-\delta\lambda}{\delta(1-\lambda)} + 1 - (1+\gamma)\kappa\right]b_3 &= \tilde{x}(1-\kappa) \\
\Leftrightarrow b_3 &= \kappa\tilde{x}
\end{aligned}$$

where the last equality makes uses of (39).

We therefore conclude that the strategy $a_t(m_t, m_t) = a_t^*(m_t) = \kappa(\tilde{x} - m_t)$, where

$\kappa \in \{\underline{\kappa}, \overline{\kappa}\}$, is an equilibrium strategy.

On the equilibrium path, the DM's profit in each period t reads

$$
\begin{aligned}
& \pi[m_t + a_t^*(m_t)] - \gamma \frac{a_t^*(m_t)^2}{2} \\
= \; & \alpha - \frac{1}{2}[m_t + \kappa(\tilde{x} - m_t) - \tilde{x}]^2 - \gamma \frac{\kappa^2(\tilde{x} - m_t)^2}{2} - \frac{1}{2h} \\
= \; & \alpha - \frac{1}{2}(\tilde{x} - m_t)^2 \left[(\kappa - 1)^2 + \gamma\kappa^2\right] - \frac{1}{2h}
\end{aligned}
$$

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

$$V(m_t) = -\frac{1}{2}\left[(\kappa - 1)^2 + \gamma\kappa^2\right] \sum_{i=t}^{+\infty} \delta^{i-t} \mathbb{E}_t(\tilde{x} - m_i)^2 + \frac{1}{1-\delta}(\alpha - \frac{1}{2h}).$$

One derives

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

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

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

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

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

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

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

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

This implies that

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

Finally, we conclude

$$V(m_t) = \frac{1}{1-\delta}(\alpha - \frac{1}{2h}) - \frac{1}{2(1-\delta)}K(\tilde{x} - m_t)^2 - \frac{1}{2h_\eta}\frac{\delta}{(1-\delta)^2}K$$

□

Proof of Proposition 3 Since the path of m_t does not depend on κ, it is easy to compare the equilibrium payoff of the DM in both equilibria and in the infinitely repeated stage game:

$V^{\underline{\kappa}} - V^{\overline{\kappa}}$ has the same sign as $(\overline{\kappa}-1)^2 + \gamma\overline{\kappa}^2 - (\underline{\kappa}-1)^2 - \gamma\underline{\kappa}^2 = (\overline{\kappa}-\underline{\kappa})\left[(\overline{\kappa}+\underline{\kappa})(1+\gamma) - 2\right] = \gamma\frac{1-\delta\lambda}{\delta(1-\lambda)}(\overline{\kappa}-\underline{\kappa}) > 0$, using (39).

$V^{\underline{\kappa}} - V^0$ has the sign of $1 - (\underline{\kappa}-1)^2 - \gamma\underline{\kappa}^2 = -\underline{\kappa}[(1+\gamma)\underline{\kappa} - 2] > 0$, as $0 < \underline{\kappa} < \frac{1}{1+\gamma}$.

$V^{\overline{\kappa}} - V^0$ has the sign of $1 - (\overline{\kappa}-1)^2 - \gamma\overline{\kappa}^2 = -\overline{\kappa}[(1+\gamma)\overline{\kappa} - 2]$. We derive that $V^{\overline{\kappa}} - V^0 > 0 \Leftrightarrow \overline{\kappa} < \frac{2}{1+\gamma}$.

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

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

$$\tilde{F}(\kappa, z) = (1+\gamma)\kappa^2 - (2+\gamma z)\kappa + 1 = 0 \qquad (39')$$

It is easy to see that $\tilde{F}_z \leq 0$. In addition, $\tilde{F}_\kappa(\underline{\kappa}, ., .) < 0$ and $\tilde{F}_\kappa(\overline{\kappa}, ., .) > 0$.

Using (5), 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:

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

Proof of Corollary 1 Using the fact that $K = (1-\kappa)^2 + \gamma\kappa^2$ is decreasing in κ at $\underline{\kappa}$ and increasing at $\overline{\kappa}$, we immediately derive from Proposition 4 that

$\frac{\partial K}{\partial\delta} \leq 0$, $\frac{\partial K}{\partial h_\varepsilon} \geq 0$, and $\frac{\partial K}{\partial h_\eta} \leq 0$.

Since K^{FB} is independent of δ and λ, this proves the result.

Proof of Proposition 6 From Lemma 1, if I^* has positive mass, then it must be an interval, so we can write $I^* = [\underline{a}, \overline{a}]$.[33] Given the assumption of uniform distribution that

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

For convenience of notation, let us write $P \equiv \frac{\overline{a}-\underline{a}}{2}$ and $\tilde{x} \equiv \frac{\overline{a}+\underline{a}}{2}$. One remarks that $\mathbb{V}(X|X \in I^*) = \frac{(\overline{a}-\underline{a})^2}{12} = \frac{P^2}{3}$. Let us also denote $\beta(K) \equiv 2\alpha - \frac{1}{h} - \frac{\delta}{(1-\delta)} \frac{K}{h\eta}$. Instead of maximizing over $\overline{a}$ and $\underline{a}$, one may equivalently maximize (24) over P and $\tilde{x}$:

$$\max_{P \in [0,1],\, \tilde{x} \in [-(1-P), 1-P]} P \left(\beta(K) - \frac{P^2}{3} - K(\tilde{x} - m_1)^2 \right) \tag{44}$$

Let us first fix $P \in [0, 1]$ and maximize (44) w.r.t. $\tilde{x} \in [-(1-P), 1-P]$.

This gives:

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

$\tilde{x} = 1-P$ if $1-P < m_1$

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

There are four cases:

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

We now maximize over P. Let us focus on the first two cases (the other two are symmetric), and start with the case $m_1 > 1$.

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

[33] If I^* has zero mass, then it gives a payoff of 0 to the DM, which is what he derives by choosing $\overline{a} = \underline{a}$, so there is no loss of generality from restricting attention to $I^* = [\underline{a}, \overline{a}]$.

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

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

Let us now consider the case $m_1 \in [0,1]$. Let $f(P) \equiv \beta(K)P - \frac{P^3}{3}$.

One maximizes a function equal to $f(P)$ on $[0, 1-m_1]$ and $g(P)$ on $[1-m_1, 1]$.

f is concave on $[0,1]$, and nonincreasing in K. g is concave on $[1-m_1, 1]$ when $m_1 \in [0,1]$. It is also easy to see that $g'(P) \leq f'(P)$ on $[1-m_1, 1]$, and $g'(1-m_1) \leq f'(1-m_1)$.

We conclude that the solution of the problem may either be

- $P^* = 0$ if $f'(0) < 0 \Leftrightarrow \beta(K) < 0$
- $P^* \in [0, 1-m_1]$ if $f'(1-m_1) \leq 0 \leq f'(0) \Leftrightarrow 0 < \beta(K) < (1-m_1)^2$
- $P^* \in [1-m_1, 1]$ if $g'(1) \leq 0 \leq g'(1-m_1) \Leftrightarrow (1-m_1)^2 < \beta(K) < 1 + Km_1^2 + 2Km_1$
- $P^* = 1$ if $0 < g'(1) \Leftrightarrow 0 < \beta(K) - 1 - Km_1^2 - 2Km_1$

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