

Reputation for a Servant of Two Masters*

Heski Bar-Isaac[†]
New York University

Joyee Deb[‡]
New York University

May 2012

Preliminary Draft: Please do not circulate

Abstract

Classic models of reputation consider an agent taking costly actions to affect a single, homogeneous audience's beliefs about his ability, preferences or other characteristic. However, in many economic settings, agents must maintain a reputation with multiple parties with diverse interests. In this paper we study reputation incentives for an agent who faces two audiences with opposed preferences. We ask if the existence of multiple audiences *per se* changes reputation incentives. Further, should the agent deal with the different audiences commonly or separately? Our analysis yields some new qualitative insights. Specifically, the presences of heterogeneous audiences is more likely to lead the agent towards "pooling" equilibria in which he takes an intermediate compromise action. Instead, dealing with only one audience leads the agent to cater towards that audience's preferences, giving rise to a "separating" outcome or pooling on some extreme action. We analyze the welfare implications, and show that the agent most prefers that both audiences commonly observe all the actions that he takes.

In our setting, reputation acts as an informal contract that enforces desirable behavior through future continuation payoffs. Our analysis highlights that the presence of multiple heterogeneous audiences can, naturally, lead these rewards to be non-monotonic in an agent's reputation. We show different ways that this non-monotonicity arises. In an infinite horizon setting, it can emerge through endogenous interactions between the audiences, through equilibrium expectations of the agent's choice of action. It can also arise, perhaps more trivially, through direct payoff interactions.

***Acknowledgements.** We thank Jack Fanning, Bengt Holmstrom, Navin Kartik, Jin Li, David Miller, Volker Nocke, Martin Peitz, Debraj Ray, Larry Samuelson, Ron Siegel, Steve Tadelis, Michael Waldman, and audiences at the University of Mannheim, Econometric Society North American Winter Meeting, University of British Columbia, Olin School of Business, University of Maryland, Simon Graduate School of Business, University of Essex, New York University and NBER Working Group in Organizational Economics.

[†]Email: joyee.deb@nyu.edu

[‡]Email: joyee.deb@nyu.edu

1 Introduction

The problem of maintaining a reputation with multiple parties who have diverse interests arises in many economic settings. For example, a manager's promotion in an organization may depend on the evaluations of two superiors who have conflicting interests, a political candidate works to gain support of his local constituency as well as the central party leadership, a credit rating agency's payoff depends directly on payment from issuers but indirectly on its credibility among investors. The central question that we ask in this paper is how the existence of multiple audiences with heterogeneous preferences affects incentives to build a reputation.

Most of the literature on reputation studies reputation formation with a single audience with homogeneous preferences.¹ Perhaps, the leading example in the literature is that of a firm of a privately known type, making quality choices, and building a reputation with a consumer base that uniformly prefers higher quality to lower quality. In contrast, we want to think about horizontal quality differentiation, and reputation formation when the consumer base may have heterogeneous or opposed quality preferences.

The presence of multiple audiences raises several new questions. Fundamentally, does the existence of multiple audiences *per se* change the agent's incentive to build a reputation compared to when she faces a single audience? With multiple audiences, agents may be able to interact separately or commonly with each audience, and thus build a common, public reputation, or separate, private ones.² What difference does this make to outcomes? We answer these questions by studying an environment in which a single agent (of a private type) builds a reputation with two audiences with opposed preferences over an infinite-horizon.

We find that, when there are multiple audiences, equilibrium behavior depends crucially on whether the agent's actions are observed separately or commonly by the audiences. Reputational incentives are qualitatively different in these two cases. Under separate observations, compromise is never optimal in that agents do not take intermediate actions, but rather prefer extreme actions. On the other hand, under common observations, catering to an audience, by choosing that audience's favored (extreme) action, is much harder to sustain. Finally, an environment with a single audience is qualitatively similar to one with multiple audiences with separate observations. This last result should not be surprising: If the actions of

¹There are exceptions. Notably, Gertner et al. (1988) consider a firm that would like to signal to lenders that is low cost, while concerned that this will cause product market rivals to be more aggressive. Their analysis highlights that pooling equilibria might naturally arise. More recently, Bouvard and Levy (2010) and Frenkel (2011) consider credit rating agencies who would like issuers to believe they are lax, and investors to believe they are tough. Intermediate reputations can be optimal.

Somewhat further from our analysis, where an agent's type is one-dimensional and an agent takes two actions in each period, Austen-Smith and Fryer (2005) develop a model where an agent's type is two-dimensional. In their model, different audiences care about orthogonal dimensions, but the agent has only a single action with which to signal to both audiences.

²The question of dealing with different audiences separately or commonly has been explored in related literatures. Notably, Farrell and Gibbons (1989) and more recently Goltsman and Pavlov (2011) address this question in a cheap talk setting. In the classic reputation literature, we know of no paper addressing this issue, though Fingleton and Raith (2005) examine career concerns of bargainers seeking to develop reputations for the quality of their information on rivals' reserve prices and contrast open- and closed-door bargaining.

the agent are not observed by both audiences, then the agent can build independent, private reputations with each. In effect, he faces two independent strategic situations, each with a single audience.

The model delivers a clear normative result: The agent always prefers that the audiences commonly observe all actions. In our setting, catering to one audience harms the other audience, and so generates no additional payoffs but does involve additional costs. Ensuring that audiences commonly observe the agent's actions can, in effect, allow the agent to credibly commit not to engage in wasteful pandering behavior (and under separate observation, such a commitment would not be as credible since each agent only observes part of what the agent does in each period). Moreover, common observation can allow for pooling equilibria that feature compromise but we show that this can only arise when it is more efficient than the alternative, where, reputation effects do not discipline behavior and there is separation. Thus, in this case also, the agent prefers that audiences commonly observe his actions.

Reputation enforces desirable behavior much like a contract (see MacLeod (2007)), but where the terms of the contract arise endogenously as continuation values that can be sustained in equilibrium. The "rewards" in the implicit contract depend on audience preferences and expectations of the actions the agent will take in equilibrium. A subtlety, in the dynamic setting is that an endogenous interaction arises between the separate audiences, through the agent's choice of actions. This payoff interaction of the audiences can make an intermediate reputation more attractive to the agent than an extreme one. Consequently, pooling equilibria arise, in which agents take a compromise action that helps maintain uncertainty about her type (and an intermediate reputation).

A key theoretical insight here is that the presence of multiple audiences in a dynamic setting changes reputational incentives qualitatively compared to the single audience-case, because it affects the curvature of the agent's rewards as a function of her reputation. In particular, the endogenously determined value of a reputation to the agent can be non-monotonic in the current reputation level, making intermediate levels of reputation optimal.

Interestingly, we find this effect does not arise in the two-period version of our model. With only two periods, the return to reputation is always linear, pushing the agent to choose extreme actions in equilibrium in order to get an extreme reputation. However, "compromise" arises in an extension of the two-period model where we impose complementarity in the audience's payoffs exogenously.

To establish these results, we analyze some related models. In our baseline environment, an agent interacts with two audiences over an infinite horizon. The two audiences have opposed preferences over the agent's action choices. In every period, the agent produces a good or service that requires two tasks. The agent has three choices for each task—an action that is favored by one audience, an action favored by the other audience, and a compromise action. The agent can be one of two privately known types: Each type is inherently favored by one audience since it can take the audience's most preferred action at a lower cost. The agent's type is realized at the start of the game and is fixed forever. Audiences are uninformed about the agent's type, and share a common prior belief about it at the start of the game.

As is standard in the reputation literature, we assume that the audiences are myopic and risk neutral and therefore reward the agent based on their expectation of the agent's action. The agent's payoff is a function of the payments it receives from each audience. We want to allow for the possibility of the agent building common and separate reputations. To this end, we study two different environments: One in which the audiences commonly observe both of the agent's tasks, and another in which each audience sees only one of the tasks of the agent. We characterize the Markov perfect equilibria in this setting. A leading example of this environment might be an organizational setting, in which a manager reports to a Finance Director and a Marketing Director. The manager is inherently suited for either a quantitative Finance project, or a more qualitative Marketing project. A compromise project is one that involves a mix of these two skills. In each period, the manager must undertake two projects. The Finance Director and the Marketing Director prefer finance and marketing projects respectively. The manager's compensation is a function of the ratings that she receives from each Director. Internal review systems might be design choices that affect the extent to which the Finance and Marketing Directors can observe different aspects of the manager's performance.

In this baseline model, we establish that compromise arises in equilibrium with two audiences and common observations, but is impossible with a single audience, or under separate observations. We also show that catering to an audience, by pooling on that audience's favored action, is hard to sustain with two audiences and common observations. The intuition is that catering to one audience implies losing the support of the other. In particular, we show that catering to both audiences (choosing the favored task for each audience) is impossible with two audiences and common observations. Analogously, we find that non-reputational (or separating) equilibria are easier to sustain under common observations. The intuition is that with common observations, it is easier for the agent to convince one audience about how she will interact with the other audience as well. Full separation is less credible if the audiences cannot observe both tasks of the agent.

Since observability seems to crucially impact the nature of reputation incentives in the presence of multiple audiences, we are led naturally to ask questions regarding design or welfare. We compare the equilibrium payoffs across equilibrium regimes, and find that, under separate observations, the agent strictly prefers the non-reputational, fully separating equilibrium to any other reputational equilibrium. Indeed, the agent would strictly prefer his less favorite action to be very costly so that he is never expected to cater in equilibrium. The opposite is true for common observations: In this case, whenever feasible, the agent strictly prefers the reputational (pooling) equilibria with compromise. Further, we find that, among all equilibria under both separate and common observations, the agent most prefers the pooling equilibrium with full compromise (whenever it is feasible).

The pooling equilibria in the infinite horizon model have the feature that pooling sustains the uncertainty about the agent's type in the long-run, and pooling remains optimal.³ This implies that the learning

³In Section 4.1, we discuss, in detail, the relationship between our results and the literature on type-based reputation.

dynamics are not very interesting in this environment: Either there is no learning in equilibrium or types are learnt with certainty. This may make the reader wonder whether dynamics play any role at all in the reputational equilibria. To address this question, we present a two-period analog of the baseline model and show that compromise no longer arises in equilibrium.

We extend the two-period model to show that compromise can occur in equilibrium, if we impose (exogenously) some complementarity between the payoffs from the two audiences in the agent's utility function. Dynamics thus seem to play an important, but subtle role here: The payoff complementarity emerges endogenously in an infinite-horizon model, and makes the value of reputation non-monotonic in the reputation level.

We explore further this insight that the shape of the value of reputation affects reputational incentives. We do this directly by analyzing reputation formation by an agent with two audiences in a setting without any dynamics or asymmetric information. We present a two-period career concerns model, in which we depart from the standard setting by allowing the agent's payoff to be a general function of her reputation (rather than a linear function.) We discuss simple micro-foundations for such reward functions, and establish that some well-known career concerns results (Holmström, 1982/99) no longer hold. In particular, we find that reputation incentives are not independent of reputation level, but can vary non-monotonically with reputation and that incentives may be non-monotonic in the precision of prior beliefs. Martinez (2009), Casas-Arce (2010) and more recent work by Miklos-Thal and Ullrich (2012) examine some applications of such career concerns models.

The rest of the paper is organized as follows: In Section 2, we present the baseline model. In Section 3, we characterize the reputational and non-reputational equilibria, and present the welfare comparisons. Section 4 contains a two-period version of our model to explain the role of dynamics in a setting with multiple audiences. In Section 5, we further investigate the role of the shape of the returns-to-reputation in determining reputation incentives. Section 6 concludes.

2 A Model of Multiple Audiences

We present an infinite horizon model with two audiences who have opposed preferences for an agent's actions. As in the standard approach to reputation, knowing the agent's type is helpful for predicting the agent's action. The agent can be one of two (privately known) types $\theta \in \{\theta_L, \theta_R\}$. His type is realized at the start of the game, and is fixed forever. In each period, the agent works for two audiences, L and R respectively. The audiences are uninformed of the agent's type. At the start of the game, the audiences have a common belief λ_0 , where λ_0 is the probability of the agent being of type $\theta = \theta_L$. An agent of type θ_L is inherently more favored by L , since he can (and is more likely) to take actions preferred by L at lower cost, as described below. Similarly, an agent of type θ_R is more favored by R .

Time is discrete, and the horizon infinite. In every period, the agent produces a good or service that

requires two actions $(a^1, a^2) \in \{a_L, a_M, a_R\} \times \{a_L, a_M, a_R\}$. The cost of an action depends on the agent's type: For an agent of type θ_L (θ_R), the a_L (a_R) action is costless, the action a_R (a_L) is very costly, and a_M has intermediate cost. Formally, we assume the following: For $\theta \in \{\theta_L, \theta_R\}$,

$$\begin{aligned}
c((a_L, a_L), \theta_L) &= c((a_R, a_R), \theta_R) &= 0 \\
c((a_M, a_L), \theta_L) &= c(a_L, a_M, \theta_L) = c((a_M, a_R), \theta_R) = c((a_R, a_M), \theta_R) &= c \\
c((a_R, a_L), \theta_L) &= c((a_L, a_R), \theta_L) = c((a_L, a_R), \theta_R) = c((a_R, a_L), \theta_R) &= C \\
c((a_M, a_R), \theta_L) &= c((a_R, a_M), \theta_L) = c((a_M, a_L), \theta_R) = c((a_L, a_M), \theta_R) &= c + C \\
c((a_M, a_M), \theta_L) &= c((a_M, a_M), \theta_R) &= 2c \\
c((a_R, a_R), \theta_L) &= c((a_L, a_L), \theta_R) &= 2C
\end{aligned}$$

For each type, say θ_L , we interpret the costless action a_L as one that the agent is inherently better suited for and therefore finds easy to do. The opposite extreme action a_R is very costly. We refer to the action a_M as a “compromise” —an action that is of intermediate and symmetric cost for both types of agents. Formally $C \geq c > 0$. In the organizational application, we can think of a manager reporting to a Finance Director and a Marketing Director. Managers are required to complete two projects in every period, and a manager is inherently suited for either a quantitative Finance project, or a more qualitative Marketing project. A “compromise project” is one that involves a mix of these two skills.

Signal Structure: We compare two different environments: We consider an environment with “separate observations”, in which, the L -audience observes a^1 , the R -audience observes a^2 . In this situation, as far as reputation-building is concerned, the agent effectively faces two separate audiences. We also consider the polar case of “common observations” in which both audiences observe both action choices of the agent.

Payoffs: The two audiences L and R have opposed preferences. In keeping with the literature on reputation, we assume that the audiences are myopic, and risk neutral, and we characterize the payments by each audience, given its expectation of the agent's action.⁴ We denote the L - (and R -) audience's payments to the agent given that it expects the agent to take action a^* by $w^L(a^*)$ (and $w^R(a^*)$). The preferences for the audiences are such that the L -audience prefers a_L to a_M to a_R . The opposite is true for the R -audience. Formally, we assume:

$$\begin{aligned}
w^L(a_R, a_R) &= w^R(a_L, a_L) &= 0 \\
w^L(a_M, a_R) &= w^L(a_R, a_M) = w^R(a_M, a_L) = w^R(a_L, a_M) &= m \\
w^L(a_R, a_L) &= w^L(a_L, a_R) = w^R(a_L, a_R) = w^R(a_R, a_L) &= 1
\end{aligned}$$

⁴Typically, this is justified by supposing that there are many constituents in the audience who bid for a single unit of good/service and so pay their full valuation. Qualitatively similar results follow from assuming that rather than each audience paying its full valuation in each period, it pays a constant fraction of its valuation.

$$\begin{aligned}
 w^L(a_M, a_M) &= w^R(a_M, a_M) &= 2m \\
 w^L(a_M, a_L) &= w^L(a_L, a_M) = w^R(a_M, a_R) = w^R(a_R, a_M) &= 1 + m \\
 w^L(a_L, a_L) &= w^R(a_R, a_R) &= 2,
 \end{aligned}$$

where $m \in (0, 1)$. The agent's payoff in any period t is a function of the payments it receives from each audience. We suppose that the agent's per-period utility is given by

$$u_t = w_t^L + w_t^R,$$

where w_t^L and w_t^R are the payments received from the L and R -audiences respectively.⁵ We assume a discount factor $\delta \in (0, 1)$. The agent's total payoff is therefore $\sum_{\tau=1}^{\infty} \delta^{\tau-1} u_t$. In the organization application, we interpret u as the overall payoff of the manager, which is a function of the ratings that she receives from each of the two Directors.

Strategies: For the baseline model, we restrict attention to pure strategies. Note that with separate observations, the beliefs held by the two audiences can be different. Therefore, the relevant state for the agent is given by a pair of beliefs (λ^L, λ^R) . Let $a_\theta(\lambda^L, \lambda^R)$ denote a pure strategy of an agent of type θ : it specifies the pair of actions $a_\theta \in \{a_L, a_M, a_R\} \times \{a_L, a_M, a_R\}$ an agent of type θ will play, given prior beliefs $(\lambda^L, \lambda^R) \in [0, 1] \times [0, 1]$.

Solution Concept: Pure-strategy Markov perfect equilibrium (MPE), where the state is given by (λ^L, λ^R) . Note that the assumption of Markov perfection, common in the literature (for example, Mailath and Samuelson (2001)), ensures that the agent's incentives are "reputational" in the sense that the agent takes actions to affect audiences' beliefs about his type. Without the restriction to Markov perfection, repeated game constructions can allow standard folk-theorem effects to arise.

Off-Equilibrium Beliefs: In a pure strategy equilibrium, it is clear that characterizing equilibrium requires us to specify off-equilibrium beliefs in case of a deviation. Most deviations appeal to one type rather than another and so the standard forward induction refinement D1 suggests that almost all deviations would lead to some degenerate beliefs.⁶ We make the following standard assumption about off-path beliefs: Once the posterior belief of an audience becomes extreme (degenerate), the audience stops updating.⁷

⁵It is a straightforward extension to allow for the audiences to be asymmetric in their influence on the agent's payoffs. For a discussion, please see Section 3.4.

⁶An exception is that under common observation if the agent is anticipated to choose (a_L, a_R) it seems perverse to update beliefs if (a_R, a_L) is observed instead.

⁷This assumption has the immediate implication that in a Markov perfect equilibrium, agents will not take any costly actions once beliefs become extreme, further that future action choices will have no impact on the future beliefs of the audiences. This may seem unappealing in practice.

If we suppose, instead, that audiences become disabused if they observe inconsistent behavior at degenerate beliefs, so that off-equilibrium deviations must be maintained beyond a single period, we still get qualitatively similar results. Details are available from the authors.

3 Analysis: Reputational and Non-reputational Equilibria

At any particular state, there are nine different pure strategies (a^1, a^2) for each agent, leading to 81 possible strategy profiles. Fortunately, the problem simplifies considerably: The first simplification stems from the fact that with pure strategies, the learning process of the audiences is very straightforward. When agents choose pure strategies, two cases can arise. First, pooling can arise, in which case the state, or, rather, one dimension of the two-dimensional (λ^L, λ^R) -state, remains unchanged—no learning occurs in equilibrium. Otherwise, there is separation and the beliefs becomes degenerate so that λ^L, λ^R or both become either 0 or 1.⁸ This means that we can restrict attention to equilibrium play in the states $(1, 1)$, $(0, 0)$, and (λ_0, λ_0) (for both common and separate observation), and $\{(\lambda_0, 0), (\lambda_0, 1), (0, \lambda_0), (1, \lambda_0)\}$, for separate observations.

Further, note that when types separate in equilibrium, they will do so by playing their costless actions; i.e., θ_L plays a_L and θ_R plays a_R (or (a_L, a_L) and (a_R, a_R) respectively, under common observations). The reasoning is a little subtle inasmuch as it requires an assumption on off-equilibrium beliefs: However, at any separating equilibrium it is reasonable (and consistent with forward induction reasoning) to suppose that a deviation to a_L (a_R) reflects that the agent's type is θ_L (θ_R). Since separating leads to the same beliefs (or continuation values), regardless of the choice of separating action, it is immediate that equilibrium separation must arise by taking costless actions.

Finally, it is worthwhile to note that, at degenerate beliefs, the agent's action cannot affect audience beliefs and (by the Markov restriction), thereby cannot affect continuation payoffs. Thus, trivially, at any degenerate belief, the agent switches to playing his costless action forever. Therefore, we are left with six types of pure strategy equilibria that can arise at non-degenerate beliefs.

- **Full Separation/No reputation:** Agent types fully separate in equilibrium. Reputation plays no disciplining role in such an equilibrium in the sense that the agent always takes the costless action (as he would in a one-shot play of the game).
- **Full Compromise:** Both types of agents play only the compromise action (a_M, a_M) .
- **Catering and Compromise:** Both types cater to one audience and play the compromise action to the other. There are two types of such an equilibrium; one where both types of agent pool on (a_L, a_M) and another in which they pool on (a_M, a_R) .⁹
- **Catering to Both Audiences:** Both types play a_L for the L -audience and a_R for the R -audience; that is both types pool by playing (a_L, a_R) .¹⁰

⁸Off-equilibrium, of course, there is in principle considerably more flexibility in how beliefs can move; however, as we argue below, standard forward induction intuition suggests that off-equilibrium beliefs would be degenerate.

⁹It seems reasonable to treat (a_L, a_M) and (a_M, a_L) as identical as far as on- or off- equilibrium updating is concerned, in the case of common observation. As will be shown below although there is a distinction between these action profiles under separate observation, neither would arise in equilibrium.

¹⁰Again, under common observation we treat (a_R, a_L) as identical.

- **Catering and Separation:** Types pool on catering to one audience by playing the favorite action of that audience, and separate on the action to the other audience. There are two types of such equilibrium. In one the catering is to the R -audience, so that θ_L -type plays (a_L, a_R) and the θ_R -type plays (a_R, a_R) . In the other, catering is to the L -audience so that the θ_L -type plays (a_L, a_L) and the θ_R -type plays (a_L, a_R) .
- **Catering to Only One Audience:** Both types pool by playing either (a_R, a_R) , or (a_L, a_L) .

Note that the first of these is fully separating, the next four are pooling equilibria, and the last involves partial pooling. We refer the fully separating equilibrium as “non-reputational” since the agent would play in the same way as in a one-shot game where reputation cannot be effective. Instead, the other types of equilibria are “reputational” equilibria as they involve at least one type of the agent choosing costly actions. We proceed by characterizing parameters under which each type of equilibrium arises; highlighting how these differ depending on whether audiences observe the agent’s actions commonly or separately.

3.1 Compromise

A key result in this paper is that with separate observations, there can be no equilibrium in which the agent chooses the compromise action. On the other hand, with common observations, it is possible for agents to choose the compromise action in equilibrium. The intuition is that, under separate observations, an agent will always have an incentive to deviate to his costless action for the audience that prefers this action. Such a deviation would increase the agent’s payoff from one audience without adversely affecting the payoff from the other. It is worthwhile to point out that this intuition does not rely on the use of pure strategies. Indeed, incentives are strict, and therefore even if we allowed for mixed strategies, compromise would not arise in equilibrium with separate observations. The next two propositions state these results formally.

Proposition 1 (No Compromise with Separate Observations). *In a setting with separate observations, there is no equilibrium with compromise.*

Proof. Suppose that, there is an equilibrium in which, without loss of generality, an agent of type θ_L chooses action a_M with some non-zero probability. Under separate observations, regardless of the choice of a^2 , deviating to $a^1 = a_L$ gives the agent of θ_L a higher wage from L ; i.e., $w^L(a_L, a_L) \geq w^L(a_M, a_L)$, $w^L(a_L, a_R) \geq w^L(a_M, a_R)$ and $w^L(a_L, a_M) \geq w^L(a_M, a_M)$. This is immediate from the audience’s payoffs. Moreover, playing a_L is costless for him, and does not affect the future continuation payoff from the R -audience. Therefore, θ_L will have an incentive to deviate to $a^1 = a_L$. This in turn implies that the θ_R -agent never chooses $a^1 = a_M$ either: Playing $a^1 = a_M$ would reveal the θ_R -agent’s type to the L -audience with certainty. But, the θ_R agent can separate costlessly by playing $a^1 = a_R$ instead of a_M . An analogous argument shows that there is no equilibrium in which any agent plays $a^2 = a_M$. \square

However, under common observations, it is no longer possible to deviate with one audience without affecting the other audience's beliefs. Indeed, now compromise can be chosen in equilibrium. To see this, consider the strategy profile of full compromise, i.e., agents of both types pool on (a_M, a_M) . This can be optimal if getting positive intermediate payments from both audiences is more valuable than getting the highest payment from only one audience, relative to the cost of the compromise action.

Proposition 2 (Full Compromise with Common Observations). *In a setting with common observations, suppose that*

$$c \leq \delta(2m - 1).$$

Then, there exists an MPE in which, for all $\lambda \in (0, 1)$, both types of agents play (a_M, a_M) . At $\lambda \in \{0, 1\}$, each type of agent takes her costless action.

Proof. Suppose that, at all beliefs $\lambda \in (0, 1)$, agents of both types pool to play (a_M, a_M) , and at degenerate beliefs, agents choose their respective costless actions. Then, the equilibrium payments paid by the audiences are given by:

$$\forall \lambda \in (0, 1), \quad w^L(\lambda) = w^R(\lambda) = 2m.$$

$$\text{For } \lambda \in \{0, 1\}, \quad w^L(1) = w^R(0) = 2 \quad \text{and} \quad w^L(0) = w^R(1) = 0.$$

Given these strategies, we can derive the value functions of the two types of agents. For interior beliefs, we have

$$V_L(\lambda) = 4m - 2c + \delta V_L(\lambda) \implies V_R(\lambda) = V_L(\lambda) = \frac{4m - 2c}{1 - \delta}.$$

At extreme beliefs, $V_R(1) = V_L(1) = V_L(0) = V_R(0) = \frac{2}{1 - \delta}$. Any deviation from (a_M, a_M) will change posterior beliefs of both audiences to an extreme (either 0 or 1, depending on the choice of off-equilibrium beliefs). For the θ_L agent, the cheapest deviation involving play of a_L is playing a_L to both audiences. So, for optimality, we need

$$-2c + \delta V_L(\lambda) \geq \delta V_L(1) \iff c \leq \delta(2m - 1),$$

which is true by the hypothesis of the proposition. It is easy to check that this condition also implies that neither type of agent will deviate to (a_L, a_R) . For the θ_R agent, the cheapest deviation involving play of a_R is playing a_R to both audiences. For compromise to be optimal, we need $-2c + \delta V_R(\lambda) \geq \delta V_R(0)$, which is identical to the condition above. \square

The parameter condition that ensure the existence of an equilibrium with compromise is quite intuitive. We get the natural comparative static that this equilibrium (like any other reputational equilibrium) is easier to sustain with more patient agents. More importantly, the condition highlights that the equilibrium is more likely to exist, the lower the cost (c) of taking the compromise action and the more the compromise action is valued by the audiences (m). Note that, there is a sense in which concavity plays

a role: Given the symmetry between the audiences, with a horizontal interpretation of the model it is natural to think of a_M as “half way” between the a_L and the a_R actions; the compromise equilibrium can only arise if each audience values the compromise action at more than the average of its valuation for the a_L and a_R action. This is a first hint that the “shape” of reputational incentives (which in turn derive from the shapes of audience preferences and underlying cost technologies) plays an important role—we return to this theme at some length below.

As we describe in Proposition A.1 in the appendix, “catering and compromise” equilibria can also arise under common observation; however, whenever such an equilibrium exists, the full compromise equilibrium also exists.

3.2 Catering

A second key result is that, catering is “harder” to sustain in equilibrium under common observations, compared to separate observations. The intuition is that under common observation, catering to one audience comes at the cost of alienating the other, whereas under separate observations an audience would not observe whether or not the agent is catering to the other audience.

Formally, we show that it is not possible for agents to pool on the strategy of catering to both audiences (playing each audience’s favored action (a_L, a_R)) under common observations. However, catering to both audiences can be sustained under a wide range of parameters under separate observations.

Proposition 3 (Catering to Both Audiences Impossible under Common Observations). *With common observations, catering to both audiences is not sustainable in equilibrium. Under separate observations, this can arise in equilibrium if $C \leq \delta$.*

Proof. First, consider the environment with common observations. Suppose that there exists an equilibrium with catering to both audiences: At all $\lambda \in (0, 1)$, agents of both types play (a_L, a_R) (or (a_R, a_L)), and at degenerate beliefs, agents choose their respective costless actions. Then, the equilibrium payments made by the audiences would be $w^L(\lambda) = w^R(\lambda) = 1$, for any $\lambda \in (0, 1)$, and $w^L(1) = w^R(0) = 2$ and $w^L(0) = w^R(1) = 0$. Further, at any $\lambda \in (0, 1)$, the value functions of the agents would be

$$V_L(\lambda) = 2 - C + \delta V_L(\lambda) \implies V_R(\lambda) = V_L(\lambda) = \frac{2 - C}{1 - \delta}.$$

At extreme beliefs, $V_R(1) = V_L(1) = V_L(0) = V_R(0) = \frac{2}{1-\delta}$. Any deviation will change posterior beliefs of both audiences to an extreme. For the θ_L agent, the cheapest deviation that results in a continuation payoff of $V_L(1)$ is playing (a_L, a_L) . So, for catering to both audiences be optimal, we must have $-C + \delta V_L(\lambda) \geq \delta V_L(1)$, which cannot hold for any positive cost C . Therefore, such an equilibrium cannot exist.

Now consider the environment with separate observations. We show that it is possible to cater to both audiences in equilibrium. We impose the following off-equilibrium beliefs: If the L (R)-audience

observes a deviation, he assigns probability 0 (1) to the agent being of θ_L type. In an equilibrium with catering to both audiences, the payments made by the audiences are as follows. For any $(\lambda^L, \lambda^R) \in (0, 1) \times (0, 1)$, $w^L = 1 = w^R$. Also, $w^L(0, \lambda) = w^R(\lambda, 1) = 0$ and $w^L(\lambda, 1) = w^R(0, \lambda) = 1$. For interior beliefs, we have, for $(\lambda^L, \lambda^R) \in (0, 1) \times (0, 1)$,

$$V_L(\lambda^L, \lambda^R) = 2 - C + \delta V_L(\lambda^L, \lambda^R) \implies V_R(\lambda^L, \lambda^R) = V_L(\lambda^L, \lambda^R) = \frac{2}{1 - \delta}.$$

Similarly, we have:

$$V_L(\lambda, 1) = \frac{1}{1 - \delta} \quad \text{and} \quad V_L(0, \lambda) = \frac{1 - C}{1 - \delta} \quad \text{and} \quad V_L(0, 1) = 0.$$

Consider the incentives of the θ_L -type agent to deviate from the equilibrium strategy.

Her payoff from playing (a_L, a_R) is given by $-C + \delta V_L(\lambda^L, \lambda^R)$. Her most profitable deviation is potentially to deviate to (a_L, a_L) . (To see why, note deviating on $a^1 = a_L$ does not make sense, since this reduces the wage from the L -audience and is costly. The cheapest way to separate is to play $a^2 = a_L$.) The payoff from deviating to (a_L, a_L) is given by $\delta V_L(\lambda^L, 1) = \frac{1}{1 - \delta}$. It follows that for catering to both audiences to be optimal, we require $-C + \frac{\delta}{1 - \delta}(2 - C) \geq \frac{\delta}{1 - \delta}$, which reduces to

$$C \leq \delta.$$

Analogous arguments for θ_R -type agent lead to the same condition. \square

The intuition of the above result applies to “catering and separation” equilibria as well. Under common observations, there do not exist any equilibria in which agents cater to one audience and separate with the other. However, such equilibria can be sustained under separate observations as we show in Proposition A.2 in the appendix.

The last type of pooling equilibria that remains to be analyzed are those in which agents cater to only one audience, i.e. both types pool on either (a_L, a_L) or (a_R, a_R) . We show in the the appendix in Propositions A.3, that such equilibria do not exist. The intuition here is simple. By catering to a single audience, the agent earns per-period wages of 2 from the audience that he caters to, and nothing from the other audience. In the pooling equilibrium, this involves costs of $2C$ from one type of agent. However, this type of agent could earn the same but at no cost by separating and catering to his own “natural” audience at no cost.

3.3 Full separation

Just as it is harder to sustain catering under common observations, it is easier to sustain full separation. Under common observation when the agent separates on even one of the actions (a^1, a^2) he demonstrates to an audience that he is of the preferred type. For example, he can convince the R -audience that he is the

θ_R -type and will play the (a_R, a_R) action. Instead, under the separate observations case, the R -audience even if it assigns probability 0 to the agent being the θ_L -type may be unsure of L -audience beliefs and may think that the agent will cater to the L -audience on the a^1 task.

Proposition 4 (Full Separation Harder with Separate Observations). *A fully separating equilibrium always exists under common observations, Under separate observations, a fully separating equilibrium exists if $C \geq \frac{2\delta}{1-\delta}$.*

Proof. Consider the environment with common observations. The only fully separating equilibrium is for the θ_L -agent to choose (a_L, a_L) and the θ_R -agent to choose (a_R, a_R) . We impose the off-equilibrium beliefs, that after any off-equilibrium observation, the audiences assigns probability 0 to the agent being of type θ_L , if they observe (a_R, a_R) , and assigns probability 1 to the agent being of type θ_L otherwise. The equilibrium payments made by the audiences are then $w^L(\lambda) = 2\lambda$ and $w^R(\lambda) = 2(1 - \lambda)$. The equilibrium value functions are given by $V_L(\lambda) = 2 + \delta V_L(1)$ and $V_R(\lambda) = 2 + \delta V_R(0)$. In particular, the value functions are identical at degenerate beliefs, i.e, $V_L(1) = V_R(1) = V_R(0) = V_L(0) = \frac{2}{1-\delta}$. Any deviation would involve a costly action and would take audience beliefs to the opposite extreme without any change in continuation payoff. Clearly, costless separation is optimal.

Let us contrast this now with the setting with separate observations. In a fully separating equilibrium, the equilibrium wages would be $w^L(\lambda) = \frac{2\lambda}{1-\delta}$ and $w^R(\lambda) = \frac{2(1-\lambda)}{1-\delta}$. The best deviation would be for a θ_L agent deviating to (a_L, a_R) . This would give a payoff of $-C + \frac{2\delta}{1-\delta} + \frac{2\delta}{1-\delta}$. For full separation to be an equilibrium, we require $\frac{2\delta}{1-\delta} \geq -C + \frac{2\delta}{1-\delta} + \frac{2\delta}{1-\delta}$, which reduces to $C \geq \frac{2\delta}{1-\delta}$. \square

3.4 Comparing Single and Multiple Audiences

We may want to ask whether the reputational effects that arise in our setting with two audiences also arise in a setting with a single audience. A straightforward extension of our baseline model can be used to answer this question. Assume that the payoff of the agent now depends asymmetrically on the beliefs of the two audiences. In particular, suppose that $u_t = \alpha w_t^L + w_t^R$ for some $\alpha \in [0, 1]$. Here $\alpha < 1$ captures the idea that the L -audience is less important in determining the agent's payoffs compared to the R -audience. This alternate formulation now makes it easy to compare the setting of two audiences with that of one audience, by considering the extreme asymmetric case of $\alpha = 0$. An analysis of this extended model shows that the case of the single audience is similar to the case of separate observations insofar as no form of compromise can arise. With a single audience "catering" emerges in the form of catering to the single audience (if $C \leq \delta$). With two audiences and separate observations, though there are several kinds of catering that can arise, these are qualitatively similar.¹¹ Non-reputational equilibria also exist under some conditions (if $C \geq \frac{\delta}{1-\delta}$). In this sense, the case of an audience with homogeneous preferences is similar to the separate observations case.

¹¹Note, however, these are not identical to the symmetric $\alpha = 1$ case. For example, when $\alpha \neq 1$ then an equilibrium with catering to a single audience (that is, with both types of agent pooling on (a_R, a_R)) can arise.

3.5 Welfare Implications

The equilibrium characterization highlights that, in the presence of two audiences, the agent's behavior depends critically on whether the agent's actions are observed separately or commonly by the audiences. Reputational incentives are qualitatively different in these two environments, and this leads naturally to questions regarding welfare. Is there a particular environment that is preferred by the agent, and what equilibria would he prefer? In Table 1, we summarize the parameter restrictions for the existence of the different types of equilibria. We also compute the per-period ex-ante expected value to the agent, for each type of equilibrium, which allows us to then make welfare comparisons.

Table 1: Summary of Equilibria and Equilibrium Payoffs

Equilibrium Type	θ_L -agent plays	θ_R -agent plays	Per-period Expected Payoffs	Separate Observations	Common Observations
Full Separation (No reputation)	(a_L, a_L)	(a_R, a_R)	2	$C \geq \frac{2\delta}{1-\delta}$	Always Exists
Full Compromise	(a_M, a_M)	(a_M, a_M)	$4m-2c$	X	$\delta(2m-1) \geq c$
Catering and Compromise	(a_M, a_R)	(a_M, a_R)	$\theta_L : 2m + 1 - c - C$ $\theta_R : 2m + 1 - c$	X	$\delta(2m-1) \geq c + C$
	(a_L, a_M)	(a_L, a_M)	$\theta_L : 2m + 1 - c$ $\theta_R : 2m + 1 - c - C$	X	$\delta(2m-1) \geq c + C$
Catering to both audiences	(a_L, a_R)	(a_L, a_R)	$2 - C$	$\delta \geq C$	X
Catering and Separation	(a_L, a_R)	(a_R, a_R)	$\theta_L : 2 - C$ $\theta_R : 2$	$\delta(2 - \lambda) \geq C \geq \frac{\delta}{1-\delta}$	X
	(a_L, a_L)	(a_L, a_R)	$\theta_L : 2$ $\theta_R : 2 - C$	$\delta(1 + \lambda) \geq C \geq \frac{\delta}{1-\delta}$	X

We find that the observability has important welfare implications. First, the agent prefers reputational equilibria (whenever possible) under common observations. Second, he prefers to fully separate under separate observations. Finally, if we compare all equilibria under separate and common observations, we find that the agent most prefers the equilibrium with full compromise. Below, we establish these results formally.

Corollary 1 (Full Separation Better than Reputational Equilibrium with Separate Observations).

Consider the setting with separate observations. Whenever feasible, the agent prefers the fully separating equilibrium to any reputational equilibrium under separate observations.

The proof follows in a straightforward way by comparing the agent's per-period expected payoffs in each equilibrium. To see the economic intuition, note that, since the compromise action is never played in equilibrium, the maximal per-period payoff that the agent can possibly get is 2. This is exactly the payoff that he receives in a fully separating equilibrium. Therefore, whenever feasible, the agent prefers this equilibrium. If the cost of the undesirable action is not too high, then equilibria with catering can arise; but these all involve at least one type of the agent taking a costly action, without any increase in the payments from the audiences. This indicates, that under separate observations, the agent would actually prefer the cost of his undesirable action to be high, so that he is not expected to cater in equilibrium.

We have the opposite result in terms of agent welfare in the setting with common observations.

Corollary 2 (Reputational Equilibria always better than Separation with Common Observations).

Consider the setting with common observations.

- i) When “full compromise” is sustainable in equilibrium, the agent strictly prefers it to a non-reputational equilibrium (full separation).*
- ii) When “catering and compromise” is sustainable in equilibrium, the agent strictly prefers it to a non-reputational equilibrium (full separation).*
- iii) The agent prefers an equilibrium with full compromise to one with catering and compromise, when both are feasible.*

Proof. Let's compare an agent's payoff in a fully separating equilibrium with that in an equilibrium with full compromise. An agent's ex-ante per-period expected payoff is $2\lambda + 2(1 - \lambda) = 2$ in a full separating equilibrium, and $-2c + 4m$ in an equilibrium with full compromise. Now, full compromise is sustainable only if $C \leq \delta(2m - 1)$. In this parameter range, $-2c + 4m$ is strictly larger than 2, thus making full compromise preferable to full separation.

A comparison of the payoffs in an equilibrium with catering and compromise with those in an equilibrium with full separation yields a similar result. An agent's ex-ante per-period expected payoff is $2\lambda + 2(1 - \lambda) = 2$ with full separation, and the minimum per-period expected payoff for an agent in an equilibrium with catering and compromise is $-c - C + (2m + 1)$. Now, catering with compromise is sustainable only if $C + c \leq \delta(2m - 1)$. In this parameter range, $-c - C + (2m + 1)$ is strictly larger than 2, thus making catering and compromise in equilibrium preferable to full separation. A similar comparison also shows that the agent's payoff in a full compromise equilibrium is higher than that in an equilibrium with catering and compromise. \square

Corollaries 1 and 2 together yield the unambiguous welfare result that, among all equilibria under either separate or common observations, the agent's most preferred equilibrium is “full compromise.”

Corollary 3 (Full Compromise is the Best Equilibrium). *Whenever feasible, full compromise is the equilibrium that gives the highest payoffs to the agent, among all equilibria under either separate or common observations.*

Proof. We know that the compromise equilibrium exists, under common observations, whenever $c \leq \delta(2m - 1)$. In this parameter range, it is easy to check that the payoff from full compromise, $4m - 2c$, is higher than 2, which is the maximal payoff obtainable in any equilibrium under separate observations. \square

These results suggest that there are two reasons for the agent to prefer common observation to separate observation. First, Corollary 3 states that if full compromise is feasible as an equilibrium it gives the highest payoff, and we have already established that it is feasible only if there is common observation. Second, even if parameters are such that full compromise is not an equilibrium (or if a different equilibrium is selected), a similar argument shows that any feasible equilibrium under common observation delivers at least as much payoff to the agent as the fully-separating, non-reputational equilibrium and Corollary 1 argues that this is the equilibrium that delivers the highest payoff to the agent under separate observations.

4 Two Period Model and the Role of Dynamics

So far, we presented an infinite-horizon model with multiple audiences, where agents can build reputations commonly or separately. We have shown that qualitatively different equilibria emerge. Since we restrict attention to pure strategies, the learning process for the audiences is very stark on the equilibrium path: Either they learn nothing (initial beliefs are unaltered), or their beliefs become degenerate. The dynamics are not important in as far as the learning process is concerned. The reader may, rightfully, wonder whether dynamics play any role in these reputational equilibria. Put differently, would the same qualitative effects arise with two audiences in a two-period model?

Consider an environment in which the agent interacts with the two audiences for exactly two periods. The other features of the setting are unchanged. We find, somewhat surprisingly, that the results of the infinite horizon do not carry over.

Proposition 5. *Consider the two-period version of our baseline model. Compromise cannot arise in equilibrium either under separate or common observations.*

The intuition behind the above result is as follows. It is immediate that in the final period, the agent will take only costless actions. It follows that at the beginning of the last period $w^L(\lambda) = 2\lambda$ and $w^R(\lambda) = 2(1 - \lambda)$. Consequently, in the first period, compromise cannot emerge in equilibrium: The θ_R agent would prefer to costlessly separate and earn 2, rather than incur a cost to pool and still earn 2. Depending on parameters, however, catering equilibria under separate or common observation may

emerge. The interested reader may refer to Proposition A.4 in the appendix for a full characterization of the equilibria in the two-period model under separate and common observations.

Proposition 5 contrasts strikingly with the possibility of compromise under common observations in the infinite-horizon (Proposition 2). This is surprising, since here, unlike in standard reputation models, dynamics does not seem to play a role in the learning process. The resolution comes from observing that in the infinite-horizon model the pooling “compromise” equilibrium, in effect acts as a commitment on the part of the agent, to keep compromising. This commitment ensures that it is valuable for the both types to pool on compromise, in order to maintain further compromise. Instead in the two-period model, because of the terminal period, no such commitment arises.

At this point, it is useful to think about the agent’s continuation value as a function of its current reputation. We refer to this equilibrium object as the “returns-to-reputation” function. In the infinite-horizon model, with common observations, in the compromise equilibrium, the agent’s returns-to-reputation function is non-monotonic in his reputation. In other words, it is optimal for the agent to maintain his reputation at interior λ , rather than allow his reputation to become extreme. Instead, in the two-period model, because of the agent’s inability to commit in the last period, the returns-to-reputation function is monotonic (constant). This ensures that, even under common observations, the agent’s most preferred action is to separate costlessly, and induce a degenerate posterior. It follows that only a catering or separating equilibrium can arise. The feasibility of catering then depends on the cost of taking the less desired, extreme action.

Dynamics therefore plays an important but subtle role here. In a dynamic setting (with common observations), an endogenous interaction arises between the two audiences, through the agent’s choice of actions. This payoff interaction of the audiences makes an intermediate reputation more attractive to the agent than an extreme one. Put differently, the presence of multiple audiences in a dynamic setting changes reputational incentives qualitatively, because it affects the curvature of the agent’s rewards as a function of her reputation. This indicates that, if there were some exogenous complementarities in the agent’s utility from the payments between the two audiences (thus making the agent’s returns-to-reputation function non-monotonic), then we might see compromise arise in equilibrium even in the two-period model. Indeed, this is the case: In Section 4.1 below, we make alternative assumptions on how the agent’s payoff depends on the audiences’ payments. We see that with complementarities, pooling and compromise can arise in equilibrium.

4.1 Direct Payoff Interactions

Consider the two-period model with two audiences and common observations. Suppose that there were complementarities between the different audiences’ payments in the agent’s utility. In particular, we alter

the agent's payoffs from $w^L + w^R$ to being

$$w^L + w^R + \gamma\sqrt{w^L w^R}.$$

Essentially, we introduce a term that allows for direct payoff interactions between the two audiences.¹² Technically, this complementarity in payoffs implies that the second period return-to-reputation function turns out to be non-monotonic in the agent's reputation. Payoff interactions between different audiences can arise fairly naturally: For example, in the organizational example of promotion, if either Director (Finance or Marketing) can veto a promotion then the manager would need the support of both. Payoff interactions also arise in the context of credit rating agencies (Bouvard and Levy, 2011; Frenkel (2011) where if investors think that the agency is too lenient then issuers would not pay for ratings, and if instead rating agencies are too tough that may put off issuers, so that the optimal common reputation may be intermediate.

In Proposition 6 below, we show that higher complementarities in payoffs (high γ) makes pooling equilibria more appealing under common observation—and compromise can arise in equilibrium. Instead, when payments from the audiences are substitutes ($\gamma < 0$) then reputational equilibria do not exist. Analogous reasoning to Proposition 1 ensures that, under separate observations, there is no equilibrium with compromise.

Proposition 6 (Compromise and Catering with Payoff Interactions). *Consider the two period model with common observations, in which the agent's utility is given by $w^L + w^R + \gamma\sqrt{w^L w^R}$. Suppose that the initial reputation is $\lambda \in (0, 1)$.*

- *If $\gamma\sqrt{\lambda(1-\lambda)} \geq c$, there exists an equilibrium with full compromise, i.e., pooling on (a_M, a_M) .*
- *If $2\gamma\sqrt{\lambda(1-\lambda)} \geq C$, there exists an equilibrium with catering to both audiences, i.e., pooling on (a_L, a_R) .*
- *If $\min\{C, 2c\} \geq 2\gamma\sqrt{\lambda(1-\lambda)}$, there exists a fully separating equilibrium, i.e. θ_L -agent plays (a_L, a_L) and the θ_R -type plays (a_R, a_R) .*

Proof. TBC. For separating equilibrium, we assume that beliefs following off-equilibrium actions are degenerate, other than for (a_M, a_M) , (a_L, a_R) and (a_R, a_L) where we assume that no updating occurs. \square

Note that in contrast to Proposition 2, the feasibility of an equilibrium with compromise is totally independent of its value m to audiences. Indeed such an equilibrium could arise here even if m were negative. Here, compromise is merely one way to pool and may be a relatively cheap way to do so. The

¹²We impose symmetry for simplicity.

possibility of pooling (either through compromise or playing the (a_L, a_R) action) is easiest when the payoff interactions are greatest—this is where γ is high and where λ is closer to $\frac{1}{2}$.

At this point, it may be useful to clarify the relationship between our results and the literature on reputation. The reader uninterested in the relationship with the literature can skip directly to Section 5.

4.2 Relationship to the literature on type-based reputation

The earlier literature on type-based reputation, starting with the work of the “gang-of-four” (Kreps et al. (1982), Kreps and Wilson (1982) and Milgrom and Roberts (1982)) and subsequently developed in Fudenberg and Levine (1989), Fudenberg and Levine (1992) focused on the incentives of an agent who can take strategic choices over actions, and seeks to develop a reputation as a “commitment” type who can only undertake a single action.¹³ Along the equilibrium path, the strategic agent perfectly mimics the “Stackelberg” type whose action is most preferred (at least in perfect monitoring environments and with sufficiently patient agents). The nature of the equilibrium is similar to the catering equilibria in our setting, where different types of agent pool on the behavior preferred by the audience.

We also have compromise pooling equilibria arising in our setting, under common observations: In these equilibria, the agent’s returns-to-reputation are maximized at an interior value leading the agent to want to maintain an interior reputation. In this situation, the agent has an incentive to pool (on compromise) in every period forever, because pooling sustains uncertainty about the agent’s type in the long run by preventing any learning by the audience. Finally, we have separating equilibria: In an equilibrium in which the agent has an incentive to separate, if he successfully separates, he has no further incentive to take costly actions, and so reputation effects die out immediately.

This latter observation has nicely been made in Mailath and Samuelson (2001) in a model where a “competent” strategic agent seeks to avoid a reputation as an “inept” type. Mailath and Samuelson (2001) argue when the reputational concern is to avoid being seen as an inept type, a competent, strategic agent can take an action that reveals himself as competent. By separating in this way, the agent will convince the audience that he is competent. Having separated, there is no uncertainty about the agent’s type and, there is no reason left for the agent to take costly actions.

Mailath and Samuelson (2001) and related work of Tadelis (2002) show that the possibility of trading reputation can sustain the uncertainty that ensures that an agent keeps having to prove his competence. Similarly, exogenous probabilities of type-changes can sustain the type uncertainty required for long-lived reputation effects Holmström (1999), Phelan (2006). A subsequent literature has sought other means to replenish type-uncertainty, either through exogenous factors (notably bounded memory in Liu and Skrzypacz (2010) and Monte (2010) or endogenous mechanisms (team production and overlapping generations in Bar-Isaac (2007), limited memory as a design choice in Ekmekci (2009) and strategic

¹³For useful overviews of the economic literature on reputation see Bar-Isaac and Tadelis (2008), Cripps (2006), Macleod (2007) and Mailath and Samuelson (2006).

choice to acquire historical observation in Liu (2011).

We contribute to this strand of the reputation literature by presenting a setting, without a Stackelberg type, in which reputation effects would not arise with a finite horizon, but in which pooling incentives can arise in the infinite horizon. In our environment, there is no need to (exogenously or endogenously) replenish uncertainty about the agent’s type. Instead, the presence of audiences with heterogeneous preferences can lead to a returns-to-reputation function that is non-monotonic, which implies that both types of agent prefer to commit to an intermediate action. Pooling behavior over an infinite horizon allows them to effectively do so.

5 Shape of Returns-to-reputation

A key theoretical insight that comes out of our analysis is that the existence of multiple audiences alters reputation incentives because it affects the curvature of the agent’s rewards as a function of his reputation. The natural next step would be to investigate, more abstractly, how the shape of the “returns-to-reputation” function affects reputation in a general model. Discreteness in action choices and the presence of two strategic types in our setting makes a general analysis cumbersome. In this sub-section, we try to isolate the effect of the shape of the returns-to-reputation function on reputation incentives, by analyzing a stripped-down model with no asymmetric information or dynamics. Simplifying along these dimensions allows us to introduce more richness along other dimensions; in particular, we allow types and actions to be drawn from a continuum, and allow for general reduced form returns-to-reputation function.

5.1 Career Concerns and Shape of Returns-to-Reputation

Consider an agent whose type θ is normally distributed with mean μ and precision h . We suppose that the agent holds this prior commonly with the audience: Thus, this is a classical “career concerns” model à la Holmström (1999), rather than a signalling or reputation model. The agent can take an action $a \in \mathbb{R}$ at a cost $c(a)$, where $c'(|a|) > 0$ and $c''(|a|) > 0$.¹⁴ The agent’s action, together with his type together generate a (noisy) public signal $s = \theta + a + \varepsilon$ where $\varepsilon \sim \mathcal{N}(0, 1)$.

An audience observes the signal and uses it (together with its expectations about the agent’s equilibrium action) to form a posterior belief concerning the agent’s type. The agent then earns a payment that depends on the audience’s belief about his type, i.e., he earns payment $R(\mathbb{E}[\theta|s])$.¹⁵ We call $R(\nu)$ the returns-to-reputation function, as it captures the final agent’s period return to having a reputation ν . In

¹⁴Typically, modellers restrict attention to non-negative actions. However, we do not do this, since this would be a substantive restriction in our context: For example, when a denotes a “horizontal” action, there is no natural interpretation or rationale for such a non-negativity restriction.

¹⁵Note that since we fix the prior distribution and signal structure throughout, in equilibrium $E[\theta|z]$ is sufficient to characterize the posterior distribution.

general, the agent's problem is to choose an action to maximize his expected return, i.e., he chooses a in order to maximize $E_{\varepsilon, \theta} [R(\mathbb{E}[\theta|s])] - c(a)$.

By assuming that the agent's type is normally distributed, we can fully separate the effect of changes in the mean prior from changes in the precision. Further, by assuming that the signal is additive in ability, action and noise, we eliminate effects that arise technologically from the signal-to-noise ratio varying with underlying ability. These features allow us to focus squarely on how the shape of the returns-to-reputation function affects the strength of reputation incentives.

Proposition 7 (Optimal effort in the career concerns model). *The agent's equilibrium effort choice a^* satisfies*

$$c'(a^*) = \frac{1}{h+1} \int R' \left(\mu + \frac{x}{\sqrt{h(h+1)}} \right) \phi(x) dx. \quad (1)$$

The interested reader may refer to the Appendix for a proof. At this level of generality, it still requires a proof that a solution exists and is unique and in writing (1), we implicitly assume that $R(\cdot)$ is differentiable almost everywhere. From here on, let us suppose that this assumption holds, and that there exists a unique solution.

In the canonical case, the reward function is linear. Holmström (1999) presents such a model, where the agent's return is his expected productivity. With a linear reward $R(\cdot)$, $R'(\cdot)$ is a constant k . From (1), it follows that $c'(a^*) = \frac{k}{1+h}$. So, there is a unique equilibrium action that is independent of the reputation (μ). The equilibrium action also decreases in h , the precision of the public signal.

Both these results rely critically on the linearity of the return function, and below, we show that these results are overturned in case of a general concave or convex return function. It is no longer true that reputation incentives always work in the same direction: In particular, $R(\cdot)$ may be non-monotonic and the equilibrium effort can be positive or negative (based on the reputation).¹⁶

Proposition 8. *Suppose that the second order condition holds and there is a unique differentiable solution $a^*(\mu)$ then $a^{*'}(\mu) > 0$ if $R(\cdot)$ is convex, and $a^{*'}(\mu) < 0$ if $R(\cdot)$ is concave.*

We provide the proof in the appendix. It is easy to see that if agents know their own private type then such returns-to-reputation functions might naturally lead to the kind of "compromise" and pooling discussed in Sections 3 and 4.1.


¹⁶To see that effort need not be monotonic in precision, consider the case where $R(x) = x$ for $x \in (0, 1)$ but $R(x) = 0$ for $x \leq 0$ and $R(x) = 1$ for $x \geq 1$. In this example, effort makes a difference only if the posterior lies in $(0, 1)$. Suppose that the prior is $\frac{1}{2}$. Then, high precision means the posterior is more often in the region where effort makes a difference to rewards; instead, low precision, even though it means that actions will shift the posteriors to a greater extent, will often do so in regions where this makes no difference to rewards. For a more complete discussion of how shape of the returns-to-reputation can lead to non-monotonicity of incentives in precision see Miklos-Thal and Ullrich (2012). See Martinez (2009) for a specific example.

5.2 Micro-foundations for shape

Since the shape of the returns-to-reputation function is critical for determining the strength of reputation incentives, it is natural to ask what leads to a particular shape. Our work clearly shows that the existence of heterogeneous audiences in a dynamic setting generates non-monotonicities in the returns to reputation. More broadly, as discussed in the context of the credit-rating agency example, non-monotonicity of the returns-to-reputation function might be consequences of externalities that the audiences impose on each other through their own choices of actions, or even as consequences of heterogeneous beliefs held by the different audiences.

A comprehensive investigation of the micro-foundations for different returns-to-reputation functions is (somewhat orthogonal to) and beyond the scope of this paper. But in this section, we make a small digression: We present two examples of horizontal and vertical differentiation, that are well-understood in the static environment, and we show that the reputation incentives that arise in these examples are non-monotonic.

Example 1 (Horizontal Reputation). Suppose that a monopolist faces two consumers with different preferences over a horizontal characteristic θ of the product. Consumer 1's valuations of a product of quality θ is given by $v_1(\theta) = V - (1 - \theta)^2$, and consumer 2's valuation is given by $v_2(\theta) = V - (-1 - \theta)^2$. We can think of consumer 1 having her bliss point of quality at $\theta = 1$ and consumer 2 at $\theta = -1$. Each consumer can get a maximum value V from consuming the product, and suffers a quadratic loss from a product whose quality is not at her bliss point. The monopolist considers three possible optimal strategies (i) not to sell, (ii) sell to only one of the two consumers and (iii) sell to both consumers. It is easy to see that the monopolist's value of having a reputation ν is given by $R(\nu) = \max\{0, V - (-1 - \nu)^2 - h^{-1}, V - (1 - \nu)^2 - h^{-1}, (2V - (-1 - \nu)^2 - 2h^{-1}) 1_{\nu > 0} + (2V - (1 - \nu)^2 - 2h^{-1}) (1 - 1_{\nu > 0})\}$, where $1_{\nu > 0}$ is an indicator function that takes the value 1 if $\nu > 0$ and 0 otherwise. Given the returns-to-reputation function, and taking a specific functional form for the monopolist's cost function, it is easy to calculate optimal effort. The figure below plots the monopolist's return (on the left) and optimal action as a function of the current belief about quality (for $V = 3$, $h = \frac{1}{2}$ and quadratic costs). In a market with



horizontal differentiation, the monopolist's returns-to-reputation are non-monotonic. So, the optimal action can display sharp reversals.

A simple application of this model is in advertising in a market with horizontal differentiation. A firm can choose advertising to highlight some aspects of its product rather than others: By investing in marketing an image, it can develop a "horizontal reputation." A good example is the tobacco industry, where relatively homogeneous products have been marketed to develop particular reputations.¹⁷ It is noteworthy, that brands have entirely reversed their marketing strategies; for example, Marlboro, associated with the "Marlboro man" was originally marketed as a feminine brand to appeal to women.¹⁸ This is consistent with our simple model where stochastic realizations might lead a firm to reverse the direction of its branding.¹⁹

Example 2 (Vertical Reputation with a Rival of Known Quality). Suppose that a firm with uncertain quality has to compete against an incumbent of known quality. The reward to establishing a particular reputation can clearly be non-monotonic in the reputation, as there is a benefit to differentiating from the incumbent. Depending on the prior, it may be beneficial to differentiate as a worse quality competitor or a higher quality competitor. In addition, lower quality differentiation is limited since the firm prefers not to appear so low quality as to lose credibility as a competitor.

This intuition can be illustrated by adapting a standard model of vertical differentiation, such as Gabszewicz and Thisse (1979), Gabszewicz and Thisse (1980) and Shaked and Sutton (1982) and Shaked and Sutton (1983). Suppose that a consumer of type t anticipates obtaining utility $ts - p$ from consuming a good of expected quality s and price p . Risk-neutral consumers vary in their tastes for quality t , where t is uniformly distributed on $[0, 1]$. Suppose that there is incumbent whose quality is fixed at 1. By examining price-setting behaviour for the entrant and incumbent, and the associated profits for the entrant, at each possible posterior belief about the agent's type, it can be shown that the entrant's returns-to-reputation function is non-monotonic in the prior. Specifically, it is given by $R(E(\mu)) =$


$$\left\{ \begin{array}{ll} 0 & \text{if } E(\mu) < 0 \\ \frac{1-E(\mu)}{9} & \text{if } 0 \leq E(\mu) < 1 \\ \frac{4}{9}(E(\mu) - 1) & \text{if } 1 < E(\mu) \end{array} \right\}. \text{ The figure below plots the optimal effort, when the precision}$$

¹⁷As an example, see Vaknin (2007) who cites Alan Blum that "...the brand of cigarette of cigarette you smoked often marked you as a fan of a particular baseball team: New York Giants fans would probably smoke a Chesterfield, a Yankee fan Camels and Lucky Strike would be preferred by Dodgers supporters p.9."

¹⁸See Vaknin (2007). "Marlboro was originally produced by Philip Morris as a woman's cigarette. They were advertised as being 'Mild as May' for the female palate and had 'Ivory Tips' to 'protect the lip'...quite a different image from the masculine symbol it was to become...(p.45) Even in 1951, Philip Morris was using this particularly strange image of an adorable infant with a baby-pink background to sell cigarettes to mothers... The early 'new' Marlboro advertisements in 1954 pictures images of men who typified 'masculine confidence'... Later the campaign was refined by the Leo Burnett advertising agency to the image that was to endure all over the world for the next thirty years, the Marlboro cowboy and 'Marlboro Country'. (p.69-70)"

¹⁹There is a related literature on advertising in markets with horizontal differentiation. In particular, Grossman and Shapiro (1984) show that the market-determined level of informative advertising may be socially excessive, and that cheaper advertising technologies may lead to more severe price competition and reduced profits. Anand and Shachar (2011) provide empirical support for an informative rather than persuasive role for advertising, highlighting that, exposure to informative advertising on a horizontal characteristic leads some consumers to reduce their demand for the good.

$h = 1$ and cost of effort $c(a) = \frac{a^2}{2}$. Again, we see that not only can the optimal effort be non-monotonic,



but it can also switch direction depending on the prior.

6 Discussion and Conclusions

In this paper, we are interested in environments in which an agent builds a reputation with multiple audiences who have potentially opposed preferences. We ask how the presence of multiple audiences with diverse preferences affects the reputational incentives for an agent. Further, we ask whether the agent should deal with the audiences separately and build private reputations, or deal with them commonly and build a public reputation. We find that the presence of multiple audiences gives rise to qualitatively different reputational equilibria. Multiple audiences who commonly observe an agent's actions lead the agent towards pooling equilibria in which he takes intermediate, compromise actions rather than extreme actions. Instead, with a single audience or with audiences who observe separately, equilibria involve separation or catering to audiences by choosing extreme actions. Our analysis also highlights that the existence of multiple audiences affects the reputation incentives, since it affects the way in which the value of an agent's reputation varies with his current reputation.

There are some natural extensions and robustness checks for the baseline model of Section 2. We contrast common observation with separate observation, but we could also consider intermediate situations. Suppose that the L (R)-audience observed one action a^1 (a^2) with certainty, and the other a^2 (a^1) with probability $p \in (0, 1)$ (and there is perfect correlation in whether the two audiences observe one or two actions in a period). Analysis of such an environment does not yield any qualitatively different insights, and turns out to be a mixture of the two polar cases in this paper.

We have characterized pure strategy equilibria; however, note that our key result establishing the possibility of compromise with common observations and the impossibility of compromise under separate observations is true among all equilibria: pure and mixed. Characterizing all the mixed strategy is not a trivial extension. In most mixed strategy equilibria, there will be (non-degenerate) learning on the

equilibrium path, so that the equilibrium imposes conditions on all (or many) beliefs simultaneously, and behaviors at different beliefs will interact.

In our baseline model, we restrict attention to perfect monitoring of actions. We do this mainly for tractability: Perfect monitoring makes the learning process very simple, though extending the analysis to allow for imperfect monitoring may generate additional insight. For instance, we can consider a more general signal structure in which the agent's actions generate signals $y^1(a^1)$ and $y^2(a^2)$. If the agent took action a , with probability μ , the realized signal is correct ($y^i(a) = a^i$), and with probability $1 - \mu$ the realized signal is one of the other actions (with equal probability).

References

- Anand, B. and Shachar, R.: 2011, Advertising, the matchmaker, *RAND Journal of Economics* **42**(2), 205–245.
- Austen-Smith, D. and Fryer, R.: 2005, An economic analysis of acting white, *Quarterly Journal of Economics* **120**(2), 551–583.
- Bar-Isaac, H.: 2007, Something to prove: Reputation in teams, *RAND Journal of Economics* **38**(2), 495–511.
- Bar-Isaac, H. and Tadelis, S.: 2008, Seller reputation, *Foundations and Trends in Microeconomics* **4**(4), 273–351.
- Bouvard, M. and Levy, R.: 2010, Humouring both parties: a model of two-sided reputation. Working paper.
- Casas-Arce, P.: 2010, Career tournaments, *Journal of Economics and Management Strategy* **19**(3), 667–698.
- Cripps, M. W.: 2006, Reputation, *New Palgrave Dictionary of Economics, 2nd Edition* .
- Ekmecki, M.: 2009, Sustainable reputations with rating systems, *Journal of Economic Theory* **146**(2), 479–503.
- Farrell, J. and Gibbons, R.: 1989, Cheap talk with two audiences, *American Economic Review* **79**(5), 1214–1223.
- Fingleton, J. and Raith, M.: 2005, Career concerns of bargainers, *Journal of Law, Economics, and Organization* **21**, 179–204.
- Frenkel, S.: 2011, Repeated interaction and rating inflation: A model of double reputation.
- Fudenberg, D. and Levine, D.: 1989, Reputation and equilibrium selection in games with a patient player, *Econometrica* **57**, 759–778.
- Fudenberg, D. and Levine, D.: 1992, Maintaining a reputation when strategies are imperfectly observed, *Review of Economic Studies* **59**, 561–579.
- Gabszewicz, J. and Thisse, J.: 1979, Price competition, quality and income disparities, *Journal of Economic Theory* **20**, 340–359.
- Gabszewicz, J. and Thisse, J.: 1980, Entry (and exit) in a differentiated industry, *Journal of Economic Theory* **22**, 327–338.

- Gertner, R., Gibbons, R. and Scharfstein, D.: 1988, Simultaneous signalling to the capital and product markets, *RAND Journal of Economics* **19**(2), 173–190.
- Goltsman, M. and Pavlov, G.: 2011, How to talk to multiple audiences, *Games and Economic Behavior* **72**(1), 100–122.
- Grossman, G. and Shapiro, C.: 1984, Informative advertising with differentiated products, *Review of Economic Studies* **51**(1), 63–81.
- Holmström, B.: 1999, Managerial incentive problems—a dynamic perspective, *Review of Economic Studies* **66**(1), 169–182.
- Kreps, D., Milgrom, P., Roberts, J. and Wilson, R.: 1982, Rational cooperation in the finitely repeated prisoners' dilemma, *Journal of Economic Theory* **27**(2), 245–252.
- Kreps, D. and Wilson, R.: 1982, Reputation and imperfect information, *Journal of Economic Theory* **27**(2), 253–278.
- Liu, Q.: 2011, Information acquisition and reputation dynamics, *Review of Economic Studies* **78**(4), 1400–1425.
- Liu, Q. and Skrzypacz, A.: 2010, Limited records and reputation. Working paper.
- MacLeod, W. B.: 2007, Reputations, relationships and contract enforcement, *Journal of Economic Literature* **45**(3), 597–630.
- Mailath, G. J. and Samuelson, L.: 2001, Who wants a good reputation?, *Review of Economic Studies* **68**, 415–441.
- Mailath, G. J. and Samuelson, L.: 2006, *Repeated Games and Reputations*, Oxford University Press.
- Martinez, L.: 2009, Reputation, career concerns, and job assignments, *B.E. Journal of Theoretical Economics* **9**(1).
- Miklos-Thal, J. and Ullrich, H.: 2012, Belief precision and effort incentives in promotion contests.
- Milgrom, P. and Roberts, D.: 1982, Predation, reputation and entry deterrence, *Journal of Economic Theory* **27**, 280–312.
- Monte, D.: 2010, Bounded memory and permanent reputations. Working paper.
- Phelan, C.: 2006, Public trust and government betrayal, *Journal of Economic Theory* **130**, 27–43.
- Shaked, A. and Sutton, J.: 1982, Relaxing price competition through product differentiation, *Review of Economic Studies* **49**, 3–14.

Shaked, A. and Sutton, J.: 1983, Natural oligopolies, *Econometrica* pp. 1469–1483.

Tadelis, S.: 2002, The market for reputations as an incentive mechanism, *Journal of Political Economy* **92**(2), 854–882.

Vaknin, J.: 2007, *Smoke Signals: 100 Years of Tobacco Advertising*, Middlesex University Press.

A Additional Results and Proofs

Proposition A.1 (Cater and Compromise under Common Observations). *Suppose that*

$$C + c \leq \delta(2m - 1).$$

Then, in the setting with common observations, there exist MPE with catering to one audience and compromise with the other.

Proof. Consider a strategy profile in which agents cater to the R -audience and compromise with the L -audience; i.e., at all $\lambda \in (0, 1)$, agents pool to play (a_M, a_R) , and at degenerate beliefs, agents choose their respective costless actions. Then, the wages paid by the audiences are as follows: For any $\lambda \in (0, 1)$, we have $w^L(\lambda) = m$, and $w^R(\lambda) = 1 + m$. For $\lambda \in \{0, 1\}$, $w^L(1) = w^R(0) = 2$ and $w^L(0) = w^R(1) = 0$. For interior beliefs, we have

$$V_L(\lambda) = \frac{2m + 1 - c - C}{1 - \delta}, \quad V_R(\lambda) = \frac{2m + 1 - c}{1 - \delta}.$$

At extreme beliefs, $V_R(1) = V_L(1) = V_L(0) = V_R(0) = \frac{2}{1-\delta}$. The most profitable deviation possible is for the θ_L agent to play (a_L, a_L) . For this deviation not to be profitable we need $-C - c + \frac{\delta(2m+1-c-C)}{1-\delta} \geq \frac{2\delta}{1-\delta}$. This reduces to

$$C + c \leq \delta(2m - 1). \quad (2)$$

The most profitable deviation possible is for the θ_R agent to play (a_R, a_R) . For this deviation not to be profitable we need $c + \frac{\delta(2m+1-c)}{1-\delta} \geq \frac{2\delta}{1-\delta}$. This reduces to $c \leq \delta(2m - 1)$, which is implied by (2) above. We can similarly consider the strategy profile in which agents pool on (a_L, a_M) , and check that we get the same condition. \square

Proposition A.2 (No ‘‘Catering and Separation’’ under Common Observations). *i) With common observations, there is no equilibrium in which agents cater to one audience, and separate with the other.*

ii) With separate observations, catering and separation arises in equilibrium. In particular,

- *If $\frac{\delta}{1-\delta} \leq C \leq \delta(2 - \lambda)$, then there exists an equilibrium in which the agents cater to the R -audience and choose their costless actions for the L -audience.*
- *If $\frac{\delta}{1-\delta} \leq C \leq \delta(1 + \lambda)$, then there exists an equilibrium in which the agents cater to the L -audience and choose their costless actions for the R -audience.*

Proof. First consider the setting with common observations. Suppose that there exists an equilibrium in which the agents cater to the R -audience, and separate with the L -audience, i.e. the θ_L -agent chooses

(a_L, a_R) and the θ_R -agent chooses (a_R, a_R) . The equilibrium payoff for the θ_L agent would be $-C + \frac{\delta(1+\lambda+2(1-\lambda))}{1-\delta}$, and that of the θ_R -agent is $\frac{2\delta}{1-\delta}$. In such an equilibrium, the types separate. However, this cannot be optimal, because the θ_L -agent can separate costlessly by playing (a_L, a_L) instead. An identical argument shows that catering to the L -audience and separating with the R audience can also not arise in equilibrium.

Next consider the setting with separate observations. Suppose that there exists an equilibrium in which the θ_L -agent chooses (a_L, a_R) and the θ_R -agent chooses (a_R, a_R) . Here, the best possible deviation for the θ_L agent would be to choose (a_L, a_L) . For this to not be profitable, we require

$$-C + \frac{\delta}{1-\delta} + \frac{\delta(\lambda + 2(1-\lambda) - C)}{1-\delta} \geq \frac{\delta}{1-\delta} \iff C \leq \delta(2 - \lambda). \quad (3)$$

Similarly, for the θ_R agent to not deviate to (a_L, a_R) , we require

$$\frac{\delta(\lambda + 2(1-\lambda))}{1-\delta} \geq -C + \frac{\delta}{1-\delta} + \frac{\delta(\lambda + 2(1-\lambda))}{1-\delta} \iff C \geq \frac{\delta}{1-\delta}. \quad (4)$$

If conditions (3) and (4) are satisfied, catering to the R -audience and separating with the L -audience is an equilibrium. In this environment, it is also possible for the agents to cater to the L -audience and separate with R , i.e., there exists an equilibrium in which the θ_L -agent chooses (a_L, a_L) and the θ_R -agent chooses (a_L, a_R) . Here, the best possible deviation for the θ_L agent would be to choose (a_L, a_R) . For this to not be a profitable deviation, we require

$$\frac{\delta(2\lambda + 1 - \lambda)}{1-\delta} \geq -C + \frac{\delta(2\lambda + 1 - \lambda)}{1-\delta} + \frac{\delta}{1-\delta} \iff C \geq \frac{\delta}{1-\delta}. \quad (5)$$

Similarly, for the θ_R agent to not deviate to (a_R, a_R) , we require

$$-C + \frac{\delta(2\lambda + 1 - \lambda - C)}{1-\delta} + \frac{\delta}{1-\delta} \geq \frac{\delta}{1-\delta} \iff C \leq \delta(1 + \lambda). \quad (6)$$

If conditions (5) and (6) are satisfied, catering to the L -audience and separating with the R -audience is an equilibrium. \square

Proposition A.3 (No Equilibria with Catering to Only One Audience): *There does not exist any equilibrium in which both types of agents cater to only one audience: In other words, pooling on (a_L, a_L) or on (a_R, a_R) cannot be an equilibrium, under either separate or common observations.*

Proof. Consider first the setting with separate observations. The result follows immediately from the fact that an agent can deviate to playing her costless action for the audience that values that action. This will increase her continuation payoff from that audience without affecting continuation payoffs from the other audience.

Next consider the case of common observations. Suppose that there existed an equilibrium in which both agents pool on (a_R, a_R) (at all $\lambda \in (0, 1)$). Then, the payments by the audiences are $w^L(\lambda) = 2$ and $w^R(\lambda) = 0$ for any $\lambda \in (0, 1)$. Any deviation from (a_R, a_R) will change posterior beliefs of both audiences to 1. So, the only deviation we need to check is whether the θ_L agent wants to deviate to (a_L, a_L) . So, we need $-2C + \frac{\delta(2-2C)}{1-\delta} \geq \frac{2\delta}{1-\delta}$, which reduces to $C \leq 0$, which is not possible. \square

Proof of Proposition 5: Proposition 5 states that no compromise is possible in equilibrium in the two-period model, regardless of whether we have separate or common observations. We prove this by directly characterizing all pure strategy equilibria in the two-period model. In particular, the proposition below states that any pure strategy equilibrium must involve either separation or catering.

Proposition A.4 (Equilibria in the two-period model). *In the two-period model. Under separate observations if the initial reputation is interior $\lambda \in (0, 1)$; there is*

- A fully separating equilibrium where θ_L plays (a_L, a_L) and θ_R plays (a_R, a_R) , if $C \geq 2\delta$;
- An equilibrium with catering to both audiences, where both types play (a_L, a_R) if $\delta \geq C$;
- No other pure strategy equilibrium; and,

Under common observations, if the initial reputation is interior then there is

- A fully separating equilibrium where θ_L plays (a_L, a_L) and θ_R plays (a_R, a_R) , if $C \geq \delta$;
- An equilibrium with catering to only one audience (where both types play either (a_R, a_R) or (a_L, a_L)), if $\delta \geq C$;
- No other pure strategy equilibrium.

Proof. We omit this proof. It follows mechanically, applying the off-equilibrium beliefs of Section 2. \square

Proof of Proposition 7.

Proof. The audience's posterior belief about θ following an observation s when the equilibrium action a^* is anticipated, is distributed normally with mean $\nu = \frac{h\mu + s - a^*}{h+1}$ and precision $h + 1$. Therefore,

$$E_{\varepsilon, \theta} [R(\mathbb{E}[\theta|z])] = \int \int_{\theta, \varepsilon} R\left(\frac{h\mu + \theta + \varepsilon + a - a^*}{h+1}\right) \phi_{\theta}(\theta) \phi_{\varepsilon}(\varepsilon) d\theta d\varepsilon \quad (7)$$

$$= \int_Y R\left(\frac{h\mu + Y + a - a^*}{h+1}\right) \phi_Y(Y) dY, \quad (8)$$

where $Y = \theta + \varepsilon$ is normally distributed with mean μ and precision $(\frac{1}{h} + 1)^{-1} = \frac{h}{h+1}$ and $\phi_x(\cdot)$ is the normal density function associated with x . Using a change of variables, (i.e. $x = \frac{Y - \mu}{\sqrt{\frac{h+1}{h}}}$, so $dx = \frac{dY}{\sqrt{\frac{h+1}{h}}}$), we express the above in terms of the standard normal distribution, denoted by $\phi(x)$.

$$\begin{aligned} E_{\varepsilon, \theta} [R(\mathbb{E}[\theta|z])] &= \int_x R \left(\frac{h\mu + x\sqrt{\frac{h+1}{h}} + \mu + a - a^*}{h+1} \right) \frac{1}{\sqrt{2\pi\frac{h+1}{h}}} e^{-\frac{x^2}{2}} \sqrt{\frac{h+1}{h}} dx \\ &= \int_x R \left(\mu + \frac{a - a^*}{h+1} + \frac{x}{\sqrt{h(h+1)}} \right) \phi(x) dx \end{aligned}$$

When the first order condition applies, the agent's maximization is the solution to

$$c'(a) = \frac{1}{h+1} \int R' \left(\mu + \frac{a - a^*}{h+1} + \frac{x}{\sqrt{h(h+1)}} \right) \phi(x) dx \quad (9)$$

Equilibrium effort is correctly anticipated. Setting $a = a^*$ yields the expression in the Proposition. \square

Proof of Proposition 8.

Proof. We can take derivatives of both sides of (1):

$$\begin{aligned} \frac{da^*}{d\mu} &= \frac{1}{c''^*} \frac{d}{d\mu} \left[\frac{1}{h+1} \int_x R' \left(\mu + \frac{x}{\sqrt{h(h+1)}} \right) \phi(x) dx \right] \\ &= \frac{1}{c''^*} \frac{1}{h+1} \int_x R'' \left(\mu + \frac{x}{\sqrt{h(h+1)}} \right) \phi(x) dx, \end{aligned}$$

which is clearly negative if R is concave and positive if R is convex. \square