

(Good and Bad) Reputation for a Servant of Two Masters*

Heski Bar-Isaac[†]

University of Toronto

Joyee Deb[‡]

New York University

February 2013

Abstract

We consider a model where an agent takes actions to affect her reputation with two heterogeneous audiences with diverse preferences. This framework contrasts with standard models of reputation which treat the audience as homogeneous. A new aspect that arises with heterogeneous audiences is that different audiences may observe outcomes commonly or separately. In our baseline model, if all audiences commonly observe outcomes, reputation concerns are necessarily efficient—the agent’s per-period payoff is higher than in the one-shot play. Instead, when the audiences separately observe different outcomes the opposite result arises—the agent’s per period payoff is lower than in one-shot play. The model involves two strategic types of player: we discuss extending the model to allow for commitment types, the impact of the horizon, and normative implications.

***Acknowledgements.** We thank Sylvain Chassang, Jack Fanning, Bengt Holmstrom, Navin Kartik, Jin Li, Jérôme Mathis, 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, NBER Working Group in Organizational Economics and Bad Homburg Workshop on Reputation.

[†]Email: heski.bar-isaac@rotman.utoronto.ca

[‡]Email: joyee.deb@nyu.edu

1 Introduction

Reputation formation plays an important role in a vast range of economic settings: Firms build reputations for high quality, workers for being smart, politicians for being liberal, and credit rating agencies for being unbiased. Most models of reputation study an economic agent forming a reputation with an audience with homogeneous preferences.¹ Yet, in practice, many settings involve audiences with heterogeneous preferences: a politician faces an electorate with differing ideological preferences; a credit rating agency builds a reputation with both issuers and investors; and a manager cares about the assessment of both her superiors and her subordinates. In this paper, we explore how reputation effects play out in the presence of heterogeneous audiences and highlight new concerns that arise in this context—in particular, that heterogeneous audiences may differ in information as well as preferences. Our results provide clear welfare conclusions and normative implications.

Specifically, we analyze a model with a long-lived agent with incentives to take costly actions, that arise only through reputational concerns, but where there are heterogeneous audiences for her reputation.² We consider an extreme case where audiences have opposed preferences. The agent can be of two types, each of which finds it less costly to “cater” to one audience; that is, to take actions that appeal to one rather than the other audience.³ “Reputation” as is typical in the literature, is a belief about the agent’s type that serves as a guide to predict agent’s behavior in future, or to predict the agent’s behavior on aspects that are hard for an audience to observe.

An issue that arises with heterogeneous audiences—and a key feature of our analysis—is that different audiences might have access to different information. Our framework allows us to explore consequences of different information structures. Specifically, we make a distinction between the case where both audiences commonly observe all outcomes, and the case where different audiences observe the agent’s actions separately.⁴

As an application, consider a middle manager directors commonly or separately observe the man-

¹There are exceptions. A couple of recent papers consider certifiers such as credit-rating agencies (Bouvard and Levy (2010) and Frenkel (2011)) who collect fees from those being certified but must also be mindful of their reputation for honesty or stringency of their certification standard among consumers of the certified product; and show that intermediate reputations can be optimal. In Frenkel (2011) differences in information of the differing audiences plays a key role. Gertner et al. (1988) consider a firm that simultaneously signals to product market rivals and to banks where they show that pooling equilibria may be natural. Somewhat further from our modeling framework, Austen-Smith and Fryer (2005) consider an agent who signals different dimensions of type to audiences who care about distinct dimensions through a single uni-dimensional action; different types of agent pool leading to the “acting white” phenomenon.

²Thus our analysis differs from the literature on common agency following Bernheim and Whinston (1986) which focuses on contractual arrangements (recently reviewed in Martimort (2007)).

³Our model is a little unusual in this respect in allowing for two types of strategic agents; in Section 5.1, we discuss, at some length, the effect of introducing commitment types.

⁴The question of whether an agent prefers that audiences separately observe different outcomes, or commonly observe all outcomes is, therefore, close to literature on cheap talk to multiple audiences explored in Farrell and Gibbons (1989) and Goltsman and Pavlov (2011). In contrast to our results, in the cheap talk literature, it is never the case that an agent is strictly better off by not communicating at all—babbling to one or both audiences is always an equilibrium. 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.

ager’s performance can reflect explicit design choices on reporting structures within the organization, or might describe their ability to assess different aspects of the manager’s work. Similarly, a firm is concerned about its perception among potential employees and clients, who may have conflicting preferences over the firm’s policy on employee compensation, and access to different information.

We show that when both audiences commonly observe all outcomes, reputation is necessarily beneficial: that is, equilibrium welfare is higher in the dynamic game, where the agent might take costly actions to affect the audiences’ beliefs about her type, than in the static, one-shot game where the agent takes no such costly actions. Instead, if different audiences observe the agent’s actions separately, reputation is necessarily harmful—equilibrium per-period welfare of the dynamic game can be lower but never higher than in the static game.

The agent’s behavior in equilibrium is qualitatively different under the two information structures. Under common observations, there may be gains from pooling on a “compromise” action that is somewhat valued by both audiences, but not the preferred action of either. Such behavior arises only when it generates more surplus than costless actions—that is when reputational incentives lead to better outcomes. In such an equilibrium, different types of agent pool on this compromise behavior and the agent benefits from an intermediate reputation that, in effect, allows the agent to commit to such socially valuable compromise in future.

Instead, with separate observations, the agent may be tempted to cater to the preferences of the audience observing that action. Doing so may be costly, but self-defeating since the other audience might (justifiably) anticipate such behavior. Since audiences observe separately, the agent has no means by which to allay such concerns.

Our model thus delivers clear welfare and normative implications. Reputation is good if all outcomes are observed by all audiences, but bad if different audiences separately observe outcomes that affect them both. To the extent that the agent can do so, she would clearly prefer to deal with audiences commonly rather than separately. If this is not possible, there may be other second best solutions—for example, the agent may choose to forego revenues from one of the audiences and focus on the other. Thus, our framework provides support for the perception that it is hard to build reputations with different audiences and it may be better not even to attempt to do so, but instead to (credibly) focus on a single audience.

It is worth noting that, this paper highlights a new reason for distortionary “bad reputation” effects to arise. In the literature on experts, as in Morris (2001) and Ely and Välimäki (2003), bad reputation arises because an audience with homogeneous preferences cannot tell whether the agent has responded appropriately to a private signal that the agent has, and consequently, the agent is tempted to conform to behavior that is anticipated to be “good” even if it is not. (see Section 10.1 of Bar-Isaac and Tadelis (2008) for an overview). Moreover, if reputation hits an absorbing state there is no further learning or opportunity to revise beliefs. Instead, in this paper, bad reputation arises when audiences with opposed preferences cannot commonly observe the agent’s actions. What drives the bad reputation effect is rather

the fact that each audience cannot see all the agent's actions, and reputation concerns with respect to one audience lead the agent to take bad actions with respect to the other.

The rest of the paper is organized as follows: In Section 2, we introduce the model and our solution concept. In Section 3, we characterize equilibrium behavior. In Section 4, we discuss the welfare implications, and the possibility that the agent might renounce one master and choose to serve only one. In our model, the agent may be of two strategic types: In Section 5, we discuss the connection to the literature on reputation with commitment types, and how to extend our model to allow for such commitment types. We also discuss other extensions. In Section 6 we conclude and highlight further issues that arise in reputation-building to heterogeneous audiences.

2 A Model of Multiple Audiences

We consider a discrete-time, infinite horizon model in which one agent interacts with two audiences who have opposed preferences for the agent's actions. As in the standard approach to reputation, knowing the agent's type is helpful for predicting his action choice. 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 the two audiences, whom we denote by L and R respectively. The audiences are uninformed of the agent's type. At the start of the game, both 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 a lower cost, as described below. Similarly, an agent of type θ_R is more favored by R .

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 (a^1, a^2) is simply the sum of the costs of each action and depend on the agent's type; $c((a^1, a^2)|\theta) = c(a^1|\theta) + c(a^2|\theta)$. 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 $\theta \in \{\theta_L, \theta_R\}$,

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

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, or a

single project with different aspects, 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.

Note that in our model both types of agent are strategic; that is, both types have to make decisions about which actions to take. This allows for symmetry in the model and is reasonable in applications, but is in contrast to many standard models of reputation which feature “commitment” types that can only take a single strategy regardless of changes in reputation or other features of the environment.⁵

Signal Structure: The existence of multiple audiences immediately points to the question of what each audience observes about the agent’s actions. We compare two different environments: We consider an environment with “separate observations”, in which, the L -audience observes a^1 and 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.⁶

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 the audience’s expectation of the agent’s action.⁷ 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:

$$\begin{aligned}
w^L(a_R, a_R) &= w^R(a_L, a_L) &= 0 \\
w^L(a_M, a_R) &= w^L(a_R, a_M) = w^R(a_M, a_L) = w^R(a_L, a_M) &= m \\
w^L(a_R, a_L) &= w^L(a_L, a_R) = w^R(a_L, a_R) = w^R(a_R, a_L) &= 1 \\
w^L(a_M, a_M) &= w^R(a_M, a_M) &= 2m \\
w^L(a_M, a_L) &= w^L(a_L, a_M) = w^R(a_M, a_R) = w^R(a_R, a_M) &= 1 + m \\
w^L(a_L, a_L) &= w^R(a_R, a_R) &= 2,
\end{aligned}$$

where $m \in (0, 1)$. The agent’s payoff in any period t is a function of the payments it receives from each

⁵See Section 5.1.1 for a more detailed discussion on commitment types..

⁶We 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 of separate and common observations.

⁷Typically, 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.

audience. We suppose that the agent's per-period utility is given by

$$u_t = w_t^L + w_t^R - c((a^1, a^2)|\theta),$$

where w_t^L and w_t^R are the payments received from the L and R -audiences respectively.⁸ 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 in a given period, which is a function of the ratings that she receives from each of the two Directors.

Strategies and Solution Concept: We consider pure-strategy Markov perfect equilibria. The Markovian assumption, 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. The payoff-relevant state for the agent is given by the beliefs of the two audiences. Note that, in a setting with separate observations, the beliefs held by the two audiences can diverge, and in general, the agent's strategies could depend on higher-order beliefs of the audiences (about the other audience's beliefs). Our restriction to pure strategies gives us tractability, by making higher order beliefs irrelevant.⁹ The relevant state for the agent in this setting is given by the pair of beliefs (λ^L, λ^R) , held by the audiences. 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 $(\lambda^L, \lambda^R) \in [0, 1] \times [0, 1]$.

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 more 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.¹⁰ We make the following standard assumption about off-path beliefs: Once the posterior belief of an audience becomes degenerate, the audience stops updating.¹¹ In particular, under separate observations, if an audience observes a deviation that makes his belief degenerate, he believes that the agent will not exert any costly effort with either audience from now on.

⁸It is a straightforward extension to allow for the audiences to be asymmetric in their influence on the agent's payoffs.

⁹The choice of pure strategy Markov-perfect equilibrium as a solution concept is a restrictive one, and, existence within this class is not guaranteed. However, existence in general is not an issue in this setting (if we allow mixing and non-Markovian strategies). Characterizing all equilibria is not a trivial extension, and is beyond the scope of this paper. In particular, allowing mixed strategies implies that there is non-degenerate learning on equilibrium path, so that equilibrium imposes conditions on all (or many) beliefs simultaneously. Further, the notion of Markovian strategies is much less appealing once we allow mixed strategies coupled with separate observations.

¹⁰An 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.

¹¹This 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.

3 Equilibrium Behavior: Catering, Compromise and Separation

In this section, we describe equilibrium behavior. The overarching finding is that the agent's equilibrium behavior depends critically on the signal structure—whether the audiences observe the agent's actions separately or commonly. We first consider the benchmark of the static one-shot game. We show that repetition of the unique equilibrium (in which the agent never take costly actions) of the static game remains an equilibrium in the repeated game under common observations, but somewhat surprisingly, may fail to be an equilibrium under separate observations. Next, we characterize the equilibria in the repeated game and show that compromise actions do not arise under separate observations, while more extreme actions are harder to sustain under common observations. Below, we present the formal analysis.

At any belief state (λ^L, λ^R) , there are nine different pure strategies (a^1, a^2) for each agent, leading to 81 possible strategy profiles at each state. However, the problem simplifies considerably both in terms of the relevant states to be considered and the possible equilibrium strategy profiles: The first simplification stems from the fact that with pure strategies, the audiences' learning process is very straightforward. When agents choose pure strategies, two cases arise. First, pooling can arise, in which case the state, or, rather, one dimension the (λ^L, λ^R) -state, remains unchanged—no learning occurs in equilibrium. Otherwise, there is separation and beliefs become degenerate so that λ^L , λ^R or both become either 0 or 1.¹² This means that we can restrict attention to equilibrium play in states $(1, 1)$, $(0, 0)$, and (λ_0, λ_0) for both common and separate observation, and $(\lambda_0, 0)$, $(\lambda_0, 1)$, $(0, \lambda_0)$, and $(1, \lambda_0)$, 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) .

¹²Off-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.

- **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) .¹³
- **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) .¹⁴
- **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) .

The first of these is fully separating, the next four are pooling equilibria, and the last involves partial pooling. We refer to 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 characterize parameters under which each type of equilibrium arises; highlighting how these differ based on whether audiences observe the agent’s actions commonly or separately.

3.1 Full separation

We start with considering the existence of a fully separating equilibrium, in which the agent plays only his costless action with both audiences. This is a natural place to start, as this is the unique equilibrium of the one-shot stage-game, and we may expect that it remains an equilibrium in the dynamic game. Our first key result is that while the separating equilibrium always exists under common observations, it may fail to exist under separate observations.

The economic intuition is that under separate observations, the audience cannot always credibly commit to any audience about how he plans to behave with the other audience. In particular, 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 in future. 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.

¹³Again, under common observation we treat (a_R, a_L) as identical.

¹⁴It 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.

Proposition 1 (Full Separation Harder with Separate Observations). *A fully separating equilibrium always exists under common observations, Under separate observations, a fully separating equilibrium exists if and only 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 degenerate off-equilibrium beliefs following any off-equilibrium observation. 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, again, the equilibrium wages would be $w^L(\lambda) = 2\lambda$ and $w^R(\lambda) = 2(1 - \lambda)$. The best deviation would be for a θ_L agent deviating to (a_L, a_R) . This would involve a cost C in the first period but would lead to beliefs $\lambda^L = 1$ and $\lambda^R = 0$ with associated payoffs $w^L(1) = 2$ and $w^R(1) = 2$. Thus the payoff from such a deviation would be $-C + \frac{2\delta}{1-\delta} + \frac{2\delta}{1-\delta}$. 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 $C \geq \frac{2\delta}{1-\delta}$. $\square$

We next turn to reputational equilibria, in which the value of future rewards provide agents the incentive to take costly actions.

3.2 Compromise

Given two audiences with opposed preferences, it is plausible that equilibrium behavior feature compromise. Our second key result is that under separate observations, the agent does not choose the compromise action in any pure strategy MPE. 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 his "favored" audience without adversely affecting the payoff from the other. This in turn implies that the other type of agent would not want to choose the costly compromise action either, since this would reveal his type anyway, thus leading to the erosion of any reputational concern.

It is worthwhile to point out that the intuition can also be extended to a certain extent to a setting with mixed strategy equilibria: We can show that, under separate observations, compromise cannot arise indefinitely, even when we allow for mixing. The next three propositions state these results formally.

Proposition 2 (No Compromise with Separate Observations). *In a setting with separate observations, there is no pure strategy MPE with compromise.*

Proof. Suppose that, there is an equilibrium in which, without loss of generality, an agent of type θ_L plays action a_M to the L -audience; i.e. $a^1 = a_M$ at some state (λ^L, λ^L) . Under separate observations, it cannot be that an agent of type θ_L plays action a_L to the L -audience, since this would reveal the agent's type and the agent could do so costlessly by choosing a_R instead; thus, regardless of the choice of a^2 , deviating to $a^1 = a_L$ will make the L -audience believe that the agent is of type θ_L , and therefore, the continuation payoff from the L -audience will be $\delta V^L(1)$, which is maximal. Moreover, playing a_L is costless for him, and does not affect the future payments from the R -audience (under separate observations). Therefore, the θ_L agent will want 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 $a^2 = a_M$. $\square$

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 3 (Full Compromise with Common Observations). *In a setting with common observations, there exists a pure strategy MPE in which, for all $\lambda \in (0, 1)$, both types of agents play (a_M, a_M) , if and only if*

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

(At $\lambda \in \{0, 1\}$, each type of agent takes her costless action.)

Proof. Necessity: 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

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

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) .¹⁵ 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. Sufficiency is established in a similar way: given the condition in the statement of the proposition, it is clear that an equilibrium can be constructed where both types of agent play (a_M, a_M) for any interior belief, and any off-equilibrium beliefs are degenerate. $\square$

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.

The idea of infeasibility of compromise under separate observations is quite general, and the result can be generalized to a certain extent in a setting with mixed actions. Specifically, we find that, even when we allow for mixed strategy equilibria, it cannot be the case that compromise is played forever. Moreover, if players are patient ($\delta \rightarrow 1$), then compromise does not arise in any mixed strategy equilibrium.

Proposition 4. *In a setting with separate observations, in any mixed strategy equilibrium, there can only be a finite number of periods of compromise.*

Proof. To establish this result, we first prove two intermediate results.

Lemma 1. *Consider the setting with separate observations. In every state, the θ_L agent must play a^L with the L -audience with positive probability.*

Proof. First note, that in any state in which both players play a pure action with the L -audience, we must have the L -agent playing a_L with the L -audience in equilibrium. Otherwise, the L -agent could deviate to a_L —this would make the L -audience believe that he is of the θ_L type and would maximize his continuation payoff from the L -audience without affecting the payoffs from the R -audience.

¹⁵If the audiences observe a deviation to (a_L, a_R) instead of (a_M, a_M) , it is not obvious what beliefs should be at such an off-equilibrium history. In a setting with common observations, it is, however, reasonable to impose that audiences have the same posterior beliefs after any history. It is easy to see that this implies that, at any posterior belief of the audiences, in order for (a_M, a_M) to be optimal, we need $c - \frac{C(1-\delta)}{2} \leq \delta(2m - 1)$. This is already implied by the condition (1).

Next, suppose that agents are supposed to play a mixed action with the L -audience in some state. Any such mixed action on equilibrium path, must involve play of a_L with positive probability (i.e. It cannot be the case that neither type plays $a_1 = a_L$). By an identical argument as above, if a_L was never played on-path with the L -audience, the θ_L -agent would deviate to $a_1 = a_L$.

Further, in particular, an agent of type θ_L must play a_L with positive probability. To see why, note that, if only the θ_R agent were playing A_L with positive probability, then a_L would reveal the agent to be of type θ_L and would give the agent a zero continuation payoff, and moreover a_L is the costliest action for the θ_R -agent - so this is not possible in equilibrium. $\square$

Lemma 2. *Consider the setting with separate observations. If players play a mixed action (say with the L -audience) at any non-degenerate state (interior belief), then types must fully separate at some point on equilibrium path.*

Proof. Suppose that players mixed between actions a_L and a_R at some non-degenerate state in period t with the L -audience, and that separation never occurred subsequently in equilibrium. Lemma 1 now implies that both types are either indifferent between a_L and another action, or strictly prefer playing a_L with the L -audience: Therefore, both θ_L and θ_R types must get the continuation payoff resulting from playing a_L from period $t + 1$ forever into the future. (Recall that the argument for Proposition 2 implies that pooling on action a_M in the future is not possible.) Let W_R denote the future payments (from $t + 1$) that the θ_L agent earns from the L -audience, if he plays a_R in the current period t and then a_L forever after from period $t + 1$.¹⁶ Analogously, let W_L denote the payoff that the θ_L agent earns from the L -audience, if he plays a_L in the current period t followed by a_L forever after. Then mixing between a_R and a_L in period t implies that $-C + \delta W_R = \delta W_L$. For the θ_R agent, playing a_R in period t followed by a_L forever, gives a continuation payoff of $W_R - \frac{C}{1-\delta}$ from the L -audience, while playing a_L in period t followed by a_L forever gives $W_L - \frac{C}{1-\delta}$. Mixing between a_L and a_R for the L -audience in period t implies that $-C + \delta \left(W_L - \frac{C}{1-\delta} \right) = \delta \left(W_R - \frac{C}{1-\delta} \right)$. But, these two indifference conditions are mutually inconsistent, providing a contradiction.

A similar argument can be used to establish that if agents ever mix between a_M and a_L with the L -audience, then this must be followed by separation somewhere on equilibrium path. Suppose, for a contradiction, that players mixed between in a_M and a_L in period t , and separation never occurred on equilibrium path. We can use Lemma 1 above to infer that a_L is played in every period along the path. Let W_M denote the continuation payoff that the θ_L agent earns from the L -audience if he chooses action a_M in current period t and then plays a_L forever after. Let W_L denote the continuation payoff that he earns from the L -audience, if he plays a_L in the current period followed by a_L forever after. Indifference between a_M and a_L in the current period implies $-c + \delta W_M = \delta W_L$. Similarly, indifference for θ_R -

¹⁶Recall that under separate observations, the wage that an agent earns from the R -audience is unaffected by the beliefs of the L -audience.

agent implies that $-c + \delta(W_M - \frac{C}{1-\delta}) = -C + \delta(W_L - \frac{C}{1-\delta})$. But these two indifference conditions are inconsistent for any positive C . $\square$

Now, the two lemmas above together imply that if players ever mix in equilibrium with the L -audience, then there must be either separation somewhere along equilibrium path, or there must be full pooling on action a_L . In particular, the compromise action cannot be played forever. A symmetric argument works for the actions played to the R -audience. $\square$

3.3 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 5 (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 and only 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 $(\lambda^L, \lambda^R) \in (0, 1) \times (0, 1)$,

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

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 $\delta V_L(\lambda^L, 1) = \frac{\delta}{1 - \delta}$. 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. □

The intuition of the results on full catering and full compromise equilibria extend to the other hybrid pooling equilibria. For instance, the infeasibility of full catering under common observation extends to “catering and separation” equilibria: Under common observations, it is not possible for agents cater to one audience and separate with the other in equilibrium. However, such equilibria arise under separate observations (See Proposition A.2 in the appendix). Similarly, the infeasibility of full compromise under separate observations implies that “catering and compromise” equilibria also do not arise under separate observations. However, they can arise under common observations (See Proposition A.1).

Finally, another type of pooling equilibrium that might exist is one 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 appendix in Propositions A.3, that such equilibria do not exist. The intuition 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.

4 Welfare Implications

The equilibrium characterization highlights the fact that reputational incentives are qualitatively different under separate and common observations, and this leads naturally to questions regarding welfare. Is there

a particular environment that is preferred by the agent, and which 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 make welfare comparisons.

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. Comparing all equilibria under separate and common observations, we find that the agent most prefers the equilibrium with full compromise. It is worth noting here that in a setting with separate observations, the agent's per-period payoff is indeed weakly lower than his payoff in the one-shot interaction. In this sense, reputational incentives are harmful under separate observations. On the contrary, under common observations, the agent's per-period expected payoff is necessarily at least as high as that in a one-shot interaction.

Below, we establish these results formally.

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

Corollary 1 (Full Separation Better than Reputational Equilibrium with Separate Observations). *Consider the setting with separate observations. 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. 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 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. First, Corollary 3 states that if full compromise is an equilibrium, then it gives the highest payoff, and this is feasible only under common observation. Second, even if 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 a payoff at least as high as that in the fully-separating equilibrium. We know from Corollary 1, that full separation is the “best” equilibrium under separate observations.

4.1 One or Two Audiences?

The welfare results imply that, if an agent could choose, he would always choose an environment with common observation. However, there may be circumstances where such a choice is infeasible—for example, if separate observation arises because the audiences differ in their ability to assess different aspects of the service that the agent provides. In this case, an alternative that might be feasible in some applications, is for the agent to commit to dealing with only one audiences—that is to commit to collecting fees from only a single audience (for both action choices). This may be easier for audiences to monitor than outcome realizations. It turns out that even when an agent can get fees from both audiences and when serving one does not limit the ability to serve the other, in the setting with separate observations, the agent may still prefer to commit to deal only with a single audience.

To see why, we consider two cases. First, if the agent knew her own type before committing to deal with a single audience, then the choice of audience would signal his type: the θ_L -agent would choose to deal with only the L -audience and would charge a price of 2 in each period (while earning nothing from the R -audience). This would yield a higher payoff than the catering equilibrium under separate observations where the agent can get a fee from both audiences but would earn only $2 - C$ per period. Second, if the agent had to choose whether to raise fees from only a single audience before knowing her type, she may still prefer to do so. She would commit to serving only the L -audience if $\lambda > \frac{1}{2}$ and in this case would expect to earn $2\lambda + (1 - \lambda)(1 - C)$ rather than $2 - C$ whenever the catering equilibrium could be sustained under separate observations when charging to both.¹⁷ The agent would, therefore, prefer to serve only the L -audience in case $\lambda > \frac{1}{1+C}$.

This result about commitment to dealing with a single audience speaks to ideas in marketing and organizational economics that highlight the importance of “focus” in both private and public (Wilson (1989)) settings. For example, Treacy and Wiesema (1995), in a book subtitled “Choose your customers, narrow your focus, dominate your market” highlight the importance of narrow focus: They suggest that

¹⁷It can be verified that, in equilibrium, the L -agent would play (a_L, a_L) in each period and the R -agent would play (a_L, a_R) .

it is impossible to fully cater to all at a high level of quality since excelling in different attributes would require contradictory processes and organization. Our mechanism instead suggests that even if it were possible to do so, customers who are unable to fully assess every aspect of service might (justifiably) be concerned that a broad focus might lead to services that are not optimized for their use. Similarly, our notion of focus might be seen as complementary to that explored in Dewatripont et al. (1999), where focus improves the informational link between performance and talent.

5 Discussion

To conclude, we discuss the how this paper relates to the more standard reputation models with commitment types. We also discuss some natural extensions of the baseline model.

5.1 Relation to Reputation Models with Commitment Types

The 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)), subsequently developed in Fudenberg and Levine (1989), Fudenberg and Levine (1992) and discussed in Mailath and Samuelson (2006) focuses on the incentives of an agent who can take strategic choices over actions, and seeks to develop a reputation for “good” behavior: The canonical setting is one in which there is incomplete information about the type of the agent, and in particular, the audience entertains the (small) possibility that the agent is a “commitment type” that is committed playing a single action.¹⁸ The key result in this setting is that, for sufficiently high discount factors, the introduction of commitment types implies a lower bound on the payoffs that the strategic agent can earn in equilibrium.¹⁹ In canonical examples, along the equilibrium path, the strategic agent can mimic the “Stackelberg” type who is committed to playing the action that is most preferred.

This kind of equilibrium is similar to the compromise equilibria of this paper, that arise under common observation, and arise only if they are efficient. Under common observation with shared audience beliefs, it is sufficient to think of the combined audience payoff, or a representative audience. What is unusual in our setting, as compared to the received literature, is that this “representative audience” has unusual preferences in that neither type is intrinsically “better.” Thus, the agent’s returns-to-reputation are maximized at an interior value leading the agent to want to maintain an interior reputation: 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. Interestingly, and in contrast to standard models with a Stackelberg type, this kind of equilibrium cannot arise in finite

¹⁸For useful overviews of the economic literature on reputation see Bar-Isaac and Tadelis (2008), Cripps (2006), Macleod (2007) and Mailath and Samuelson (2006).

¹⁹See also Gossner (2011) for an elegant derivation based on information theory.

horizons.²⁰

The catering equilibrium that arise in our model, only do so under separate observations. As we describe below, our result that equilibrium payoffs will be lower than in the static game is consistent with the results of Fudenberg and Levine (1989), since under separate observations, the different audiences cannot distinguish different stage-game strategies even in the long run.

Finally, we have separating equilibria: In equilibrium, if the agent successfully separates, he has no further incentive to take costly actions, and so reputation effects die out immediately. This latter observation has been made before 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) show that a competent, strategic agent tries to take an action that separates him from the inept type and reveal him as competent. But once he separates, there is no uncertainty about the agent’s type and, there is no reason left for the agent to take costly actions. Sustaining some uncertainty (or renewal of uncertainty) seems important for long-term incentives: Mailath and Samuelson (2001) and related work of Tadelis (2002) show that the possibility of trading reputation can sustain the uncertainty about an agent’s type. 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 choice to acquire historical observation in Liu (2011).

We contribute to this strand of the reputation literature by presenting a setting, without a commitment (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 an implicit contract that is non-monotonic; that is the rewards associated with reputation or belief are non-monotonic in the reputation. This, in turn, 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.

5.1.1 A Model with Commitment Types

A natural extension is of our model is to allow for the possibility of commitment types.

In addition to the θ_L and θ_R types considered in the baseline model, suppose that the agent could also be of a simple commitment type θ_{a_1, a_2} where an agent of type θ_{a_1, a_2} takes action (a_1, a_2) in every period regardless of history. Further suppose that initially both audiences share a prior on the agent’s type $\lambda(\cdot)$ that assigns non-zero probability to each possible type of agent.

²⁰Proofs are available with the authors and follow a standard backward induction argument.

The analysis under common observation is standard. For sufficiently high discount factors, each strategic types θ_R and θ_L can get arbitrarily close to achieving its Stackelberg payoff—that is the payoff that it could achieve by committing to a particular strategy at the start of the game. Note the different types may have different actions and preferences as to the Stackelberg action. If $2 \geq 4m - 2c$ then the θ_L would achieve its highest payoff by committing to play (a_L, a_L) in each period; in the reputational game with δ close to 1, then following Fudenberg and Levine (1989), Mailath and Samuelson (2006), or Gossner (2011) the lower bound on per-period payoffs for the θ_L in any Nash equilibrium will approach 2. Similarly the θ_R type would achieve its highest payoff by committing to play (a_R, a_R) . Instead, if $4m - 2c > 2$ then both strategic types would mimic the θ_{a_M, a_M} -type.²¹

The analysis under separate observation is somewhat more involved. The difficulty that arises is that each audience cannot observe both the action choices of the agent, and so cannot distinguish the action profile (a_M, a_M) from the action profile (a_M, a_L) or (a_M, a_R) ; thus an L -audience observing a_M in the first period could believe that the agent is strategic (if such action was part of an equilibrium profile for a strategic type), or of type θ_{a_M, a_L} , θ_{a_M, a_M} or θ_{a_M, a_R} . Suppose (in the spirit of the reputation literature) that the likelihood of the agent being of either strategic type is much greater than the likelihood of the agent being a commitment type. Further, it seems reasonable and convenient to suppose that $\lambda(\theta_{a_i, a_i}) \gg \lambda(\theta_{a_i, a_{-i}})$ —that is a commitment type is much more likely to be the type for which $a_1 = a_2$ than $a_1 \neq a_2$ reflecting that this may represent some capability or preference for that kind of action. This assumption also ensures that if an audience observes some action that does not correspond to the equilibrium action of a strategic type, then the audience can be confident that it is from a type that takes the same action to the other audience.

First, it is clear that a strategic agent θ_L need not be able to approximate her Stackelberg payoff. For example, suppose that this was achieved by the action profile (a_M, a_M) then the θ_L would prefer to deviate to (a_L, a_M) which would lead to the agent earning 2 from the left audience and $2m$ from the right audience, yielding an average payoff of $2 + 2m - 2c$ which is strictly greater than the per-period payoff of $4m - 2c$ earned by maintaining the (a_M, a_M) action. Similarly, if the θ_L preferred to commit to action (a_L, a_L) and θ_R to (a_R, a_R) this would not ensure that they can approximate the average payoffs associated with these behaviors. If these were the expected strategies, then the θ_L agent could deviate to (a_L, a_R) and continue to earn 2 perpetually from the L -audience while earning an additional $2 - C$ from the R -audience. Thus, equilibrium may involve the θ_L and θ_R types pooling on (a_L, a_R) in a catering equilibrium. Allowing for commitment types, therefore, does not overturn the result that reputation may lead to worse outcomes than the repetition of the static outcome under separate observation.

²¹Note that in contrast to the baseline model with no commitment types, reputation concerns may lead strategic types to play (a_M, a_M) for some periods even under a finite horizon.

5.2 Purely Information Linkages between Audiences

So far, we assumed that the audiences compensate the agent based on their expectations of both actions of the agent. In some applications this may be reasonable (for example, shareholders take a direct financial interest in all actions that a CEO undertakes, including those with respect to union representatives), whereas in other it may be less so (for example, in a product market application, a consumer pays the firm only for the good that she expects to receive). We may ask what happens in the latter case, in which the linkages between audiences are purely informational, rather than direct payoff linkages.

It turns out that the qualitative results on existence of equilibria are similar to the baseline model insofar as compromise remains impossible under separate observation. The intuition here is identical—one type of agent would prefer to separate by taking his costless action and revealing to the audience that he is of the audience’s preferred type. One substantive difference is that catering equilibria can arise under common observations in this altered setting with pure informational linkages. The intuition is that in the absence of direct payoff linkages, it is no longer true that catering to one audience necessarily implies lowering payoffs from the other. Indeed, catering can even be optimal in this setting.

5.3 Role of Dynamics

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 reader may, rightfully, wonder whether dynamics play any real 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. Indeed, in the two-period model, compromise cannot arise in equilibrium either under separate or common observations. This contrasts strikingly with the possibility of compromise under common observations in the infinite-horizon (Proposition 3). This is surprising, since 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.

The role of dynamics is subtle here: 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.

6 Conclusion

The simplicity of the model and analysis—in particular, the opposed preferences of the two audiences—provides strong conclusions: Reputation-building to agents with diverse preferences leads to better outcomes for the agent than static play (that is, with no scope for reputation effects) if both audiences can observe everything but not if they observe separately.

If the agent cannot commit to take all actions so that all audiences commonly observe them then, to the extent that it is feasible to do so demonstrably, the agent would prefer to give up entirely on collecting rents from one of the audiences and serve only one master. The intuition is that, when there are two audiences who cannot observe everything that the agent does, each audience justifiably fears that the agent panders to the other audience in aspects that the other audience observes more clearly. Consequently, the agent’s actions end up being self-defeating. Beyond this clear intuition and the welfare and normative implications (in cases where the agent can control what audiences observe, or can demonstrably deal with only one audience) the framework raises several other issues.

This paper considers audiences with directly opposed preferences who interact only indirectly through the induced actions of the agent. There are interesting applications where there are more direct interactions; for example in the context of certification, (e.g. in Frenkel (2011) and Bouvard and Levy (2010)) the agent’s action (whether or not to certify) affects the price at which the certified sellers and the buyers trade. There are also applications where audiences are not directly opposed but rather have partially aligned preferences (e.g. consider a firm selling to consumers with both horizontal and vertical quality preferences). In related work, Bar-Isaac and Deb (2013) we consider a setting that allows richer preference heterogeneity, but for tractability, we restrict attention to a two-period career concerns model.

Our results highlight that a key concern for each audience is the presence, knowledge and beliefs of the other audience. By focusing on a simple environment and pure strategies we ensure that the relevant reputation is simply an audience’s belief about the agent’s type; however, it is clear that in a broader environment, beliefs about the other audiences that an agent deals with, together with beliefs about their beliefs, can affect an audience’s expectations of an agent’s behavior. Indeed, an agent may take actions not only to affect an audience’s belief about her “intrinsic type” or capabilities, but also but other characteristics of the environment and other audiences that she faces. In Section 4.1, we consider a rather blunt (though relatively easy to monitor) means of doing so: foregoing trade opportunities with one audience. In ongoing work, we are exploring richer decisions by the agent on the kinds of audiences she should choose to interact with.

References

- 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 Deb, J.: 2013, Career concerns for a servant of two masters and the shape of reputation. Working paper.
- Bar-Isaac, H. and Tadelis, S.: 2008, Seller reputation, *Foundations and Trends in Microeconomics* **4**(4), 273–351.
- Bernheim, B. D. and Whinston, M. D.: 1986, Common agency, *Econometrica* **54**(923-942).
- Bouvard, M. and Levy, R.: 2010, Humouring both parties: a model of two-sided reputation. Working paper.
- Cripps, M. W.: 2006, Reputation, *New Palgrave Dictionary of Economics*, 2nd Edition .
- Dewatripont, M., Jewitt, I. and Tirole, J.: 1999, The economics of career concerns, part ii: Application to missions and accountability of government agencies, *Review of Economic Studies* **66**, 99–217.
- Ekmekci, M.: 2009, Sustainable reputations with rating systems, *Journal of Economic Theory* **146**(2), 479–503.
- Ely, J. C. and Välimäki, J.: 2003, Bad reputation, *Quarterly Journal of Economics* **118**(3), 785–814.
- 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.
- 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.
- 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.
- 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.
- Martimort, D.: 2007, Multi-contracting mechanism, in N. Blundell and Persson (eds), *Advances in Economic Theory Proceedings of the World Congress of the Econometric Society*, Cambridge University Press.
- 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.
- Morris, S.: 2001, Political correctness, *Journal of Political Economy* **109**, 231–265.
- Phelan, C.: 2006, Public trust and government betrayal, *Journal of Economic Theory* **130**, 27–43.
- Tadelis, S.: 2002, The market for reputations as an incentive mechanism, *Journal of Political Economy* **92**(2), 854–882.
- Treacy, M. and Wiesema, F.: 1995, *The Disciple of Market Leaders: Choose Your Customers, Narrow Your Focus, Dominate Your Market*, Addison-Wesley.
- Wilson, J.: 1989, *Bureacracy: What Government Agencies Do and Why They Do It*, New York: Basic Books.

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 (5) and (6) hold, 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 the setting with separate observations. The result follows 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. $\square$