Strategic Reviews

(Forthcoming in Management Science)

Mohamed Mostagir*

James Siderius†

Abstract

The impact of product reviews on consumer purchasing behavior is empirically well documented. This can create perverse incentives for firms to offer reviewers side payments (”bribes”) in exchange for biased reviews for their products. The presence of bribes distorts the information in reviews and leads to detrimental effects on consumer utility. This paper builds a two-sided reputation model where a reviewer can inflate her reviews in exchange for bribes. If the reviewer accepts bribes and misrepresents her reviews, then she builds her reputation as an inaccurate reviewer and makes consumers less likely to follow her recommendations. This in turn makes firms no longer interested in offering her a bribe. Can the reviewer retain influence over consumers' purchasing decisions while simultaneously accepting bribes and misrepresenting her reviews? We provide a characterization of the environments that allow this kind of manipulation, and show that regulatory policies that aim to reduce bribes can lead to undesirable outcomes. Finally, we show that eliminating bribes from the market can increase the welfare of all market participants, *even* for those firms who would have otherwise bribed in exchange for more favorable reviews.

1 Introduction

Consumers often make purchasing decisions by relying on reviews. This can make it attractive for firms to offer reviewers side payments —“bribes”— in exchange for inflated ratings for their products. A recent example is the case of PewDiePie, a popular YouTube gamer who was at the center of a settlement between the Federal Trade Commission (FTC) and Warner Bros, after the latter was found to have offered thousands of dollars in compensation to reviewers in order to “promote the game in a positive way and not to disclose any bugs or glitches they found”(FTC (2016b)).

This paper focuses on influential reviewers like the example above. We develop a two-sided reputation model where a reviewer with a *private* ability type is situated between consumers and firms. The reviewer is either low or high skilled, and this type determines her accuracy

*University of Michigan Ross School of Business

†Laboratory for Information and Decision Systems, MIT

in evaluating the quality of the firms' products. Consumers observe the reviewer over time and learn how accurate she is by comparing her reviews to the ex-post realized qualities of the products. A more accurate reviewer becomes more influential with consumers, who rely on her reviews to make purchasing decisions. The reviewer obtains utility from this influence, e.g., through direct support from consumers on platforms like Patreon or from advertising revenue on dissemination channels like YouTube.

An influential reviewer is also an appealing target for bribing by firms. In our model, firms are either honest (do not bribe) or strategic (offer bribes if it is profitable) and this attribute is unobservable to consumers. If the reviewer accepts bribes and starts misrepresenting her reviews, her accuracy declines and her influence diminishes. This in turn makes firms no longer interested in bribing her. We identify conditions under which the equilibria of this game admit situations where the reviewer consistently accepts bribes and misleads consumers. We then use our model to discuss the implications to the utility of all players in the market: consumers, firms, and the reviewer.

Contribution and overview of results. Below is a summary of our main results.

Characterization: We first show in Theorem 1 that the above game admits at most two equilibria. In the *separating* equilibrium, both the low and high type reviewers report truthfully, no bribes are offered, and all firms get fair reviews. In the *pooling* equilibrium, the low type reviewer reports truthfully but the high type mimics a low type and posts less accurate reviews. Of the many ways in which she can pretend to be a low type, she chooses a particular one: she injects positive bias into the reviews of firms that bribe her and negative bias into the reviews of firms that do not. In this equilibrium, all the strategic firms bribe and get inflated reviews while all the honest firms get depressed reviews. Consumers do not observe which firms bribe, so they cannot tell if the reviewer is providing bias to some firms over others or is simply inaccurate. As a result, consumers are consistently biased toward the products of bribing firms.

The separating equilibrium *always* exists, but the pooling equilibrium may not. If the pooling equilibrium does not exist, we say the environment is *bribe-proof*. We provide detailed conditions under which the pooling equilibrium (and bribes) can exist.

Comparative statics: Next, we use our model to deliver insights about how the different parameters of the environment affect bribes. Theorem 2 shows that an increase in the proportion of honest firms can lead to *lower* consumer utility. This is because the pooling equilibrium only exists if the high type reviewer can credibly play like a low type. If she biases the reviews of bribing firms up, she needs to bias other firms down so that her overall accuracy is consistent with that of the low type reviewer. The presence of honest firms in the market allows her to do this and maintain a facade of being impartial and fair. A potential implication is that regulations that aim to reduce the proportion of bribing firms in the market (through the threat of auditing and fines, for example) can end up making it easier for the reviewer to accept bribes and manipulate her reviews. These non-monotone effects of regulation on consumer welfare have recently been observed in the empirical work of Ershov and Mitchell (2020).

Consumer utility is also affected by the ex-ante information that the market has about the skill of the reviewer. Propositions 2 and 3 illuminate an important counterintuitive finding: increasing the expected skill of the reviewer may not necessarily lead to better outcomes for consumers. This is because more skilled reviewers have more leeway in inflating their reviews, and this makes the prospect of bribing them more appealing to firms. Thus this improvement in reviewer quality might allow both reviewers and firms to extract greater surplus through bribes and inflated reviews at the cost of hurting consumers.

Total welfare: Finally, Proposition 4 gives a mild condition under which bribe-proof environments improve the total welfare of *all* market participants (consumers, reviewer, and firms). Interestingly, it can be the case that the bribing firms themselves would prefer a world without bribes. In these instances, the reviewer accepts bribes in order to not hurt the firm by giving it a poor review. The bribe that the firm pays to the reviewer becomes a form of extortion or "protection money": it does not benefit the firm in any way, but rather helps it avoid a worse fate.

Commitment versus Fully-dynamic model: Throughout the paper, we present the above findings using a simplified game-theoretic model that assumes the reviewer commits to playing according to a certain (low or high) type. This commitment assumption greatly simplifies the exposition, but importantly, we show through detailed analyses and equilibrium refinement choices that commitment emerges as the unique equilibrium of the full dynamic interaction between players. For the interested reader, the details and analysis of the full dynamic model are included in Appendix B.

Related Literature. The effect of reviews on consumer purchasing behavior is empirically well-documented in Chevalier and Mayzlin (2006); Senecal and Nantel (2004) and Dellarocas et al. (2007). The phenomenon of influential online reviewers is relatively recent, but bears some resemblance to earlier work that documents the effect of movie critics on viewership (e.g. Reinstein and Snyder (2005)). There is also an emerging literature on how consumers learn from reviews, e.g. Ifrach et al. (2019) and Acemoglu et al. (2017), which adopts a Bayesian sequential learning approach where consumers leave reviews of products and these reviews are used by future consumers to determine the quality of the product. Chen and Papanastasiou (2019) study how a firm can manipulate this sequential learning process through planting an initial "fake" purchase, while Dellarocas (2006) is an earlier work that studies a game where firms try to manipulate opinions in online forums and shows when such manipulation is beneficial to consumers or firms. These papers are part of a broader literature that examines manipulation of agents in various online settings (e.g. Candogan and Drakopoulos (2017); Belavina et al. (2018); Papanastasiou (2018); Mostagir et al. (2021)).[1]

Unlike the above-mentioned literature, our paper conceives of the reviewer as a proper agent acting strategically to further her own interests, and so we model all three types of actors. This relates our paper to the literature where an expert communicates information to a buyer on behalf of a seller. Lizzeri (1999) studies the incentives of a monopolistic expert to accurately reveal information and Inderst and Ottaviani (2012) study how firms can compete for market share by trying to influence the expert through kickbacks. In a similar vein, the recent paper of Fainmesser and Galeotti (2020) studies competition between influencers who choose how much 'sponsored' vs. 'organic' content to provide. The equilibria of our simplified commitment model lead to different findings from the above papers. This is a consequence of the fact that the commitment model is the reduced-form outcome of the repeated interactions that allow both sides of the market to learn about the unknown type of the reviewer – an element that is absent in the aforementioned papers. Mitchell (2021) models the dynamic relationship between a reviewer and a consumer but does not model firms, which leads to a different incentive structure and different results from the three-player model.

Finally, our paper is related to the recent work on two-sided reputation, as in Bar-Isaac and Deb (2014) and Bouvard and Levy (2017), where, in the latter, an agent is faced with the problem of providing certification for sellers of different types. Our work shares some similarities, but there are also several important distinctions. Preferences in our paper are aligned (both consumers and firms prefer a high-type to a low-type reviewer) but are conflicting in Bouvard and Levy (2017). In addition, the lack of commitment in Bouvard and Levy (2017), which underscores the fact that commitment arises from the repetition in the dynamic model, leads to different outcomes since the reviewer in that paper can misrepresent herself in the first period in order to take advantage in the second (and final) period. Finally, there are also differences that stem from that fact that while firms elect to be certified (and therefore induce a signaling game),

[1] Some of this literature (e.g. Candogan and Drakopoulos (2017)) also touches on the larger issues of information disclosure, as discussed in Ottaviani and Sørensen (2006) or Dye (1985).

they cannot choose not to be reviewed, as per the Consumer Review Fairness Act of 2016 (FTC (2016a)). Durbin and Iyer (2009) consider a model where an advisor (reviewer in our setting) is more concerned about her reputation for honesty rather than her accuracy, which gives rise to different dynamics as the authors themselves note in their paper.[2]

The next section demonstrates the main ideas of the paper in a reduced-form example. Our formal model and equilibrium characterization are presented in Section 3. Section 4 uses this model to describe what environments admit bribes and examine the effect of the different parameters on consumer utility and total welfare. Section 5 concludes the paper. All proofs are in Appendix A. Interested readers can find the equilibrium analysis that allows us to reduce the full dynamic model to the commitment model used throughout the paper in Appendix B. An optional discussion that connects the main model of the paper to the full dynamic model is in Section 3.5.

2 Reduced-Form Example

We provide intuition for our results through a stylized example. A reviewer repeatedly interacts with a mass of consumers for whom she writes reviews which inform purchasing decisions. Firms arrive sequentially in discrete time, and some of these firms are 'friends' of the reviewer. The reviewer would like to help her friends out, but is impartial to those who are not her friends. We are interested in whether the reviewer can repeatedly steer the *Bayesian* consumers towards products sold by her friends, without those consumers being able to tell that this is what she is doing.

The setup is as follows. Time is discrete $t = 1, 2, \ldots$. Firm t arrives at time t and sells a product whose quality q_t is distributed as $\mathcal{N}(0, 1)$. The reviewer samples the product and observes a signal $s_t = q_t + \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, \sigma_\omega^2)$ and is i.i.d across time and independent of product quality. The reviewer's private type/skill ω is either **H**igh ($\sigma_H = 1$) or **L**ow ($\sigma_L = 2$), and both types are equally likely. After observing her signal, the reviewer posts a review r_t, which may or may not be the same as s_t. A reviewer cannot pretend to be more accurate than she actually is, so type **H** can pretend to be type **L** but the opposite is not possible. If the reviewer always tells the truth regardless of her type, then type **L** reviews are distributed as $r_t^L \sim \mathcal{N}(0, 5)$ and type **H** reviews are distributed as $r_t^H \sim \mathcal{N}(0, 2)$.

The true quality q_t is eventually revealed to every one before the next period. Consumers are aware of the possible reviewer types and the probabilities with which they occur, so they

[2] Quoting Durbin and Iyer (2009): "our advisor does not care about a reputation for accuracy, but rather a reputation for incorruptibility...in general, these two types of reputational incentives do not coincide."

can learn the reviewer's skill by comparing her reviews with the revealed qualities over time. Consumers are also aware that some firms might be friends of the reviewer, *but they do not know which firms are friends and which are not.* If consumers believe that the reviewer is intentionally misrepresenting her reviews, they stop listening to her forever, which is an outcome that she wants to avoid. Note that if the reviewer systematically boosts the reviews of her friends, then she will be behaving in a way that is inconsistent with being either type **L** or type **H**, and this will trigger consumers to ignore her.

Before they make their purchasing decisions, consumers optimally extract information about the quality q_t from the review r_t. They do this through adjusting the review by the precision of what they *believe* the reviewer's type is (this is the standard inverse-variance weighting formula, as in, e.g. Cochran (1954), that we also discuss in Section 3). In particular, if consumers believe that the reviewer's type is $\hat{\omega} \in \{L, H\}$ then the expected quality given the review r_t and the type $\hat{\omega}$ is given by

$$\mathbb{E}[q_t | r_t, \hat{\omega}] = \frac{r_t}{1 + \sigma_{\hat{\omega}}^2}$$

Now suppose that in this sequence of firms, with probability $1/2$ a firm is a friend of the reviewer and with probability $1/2$ it is not. Let Z_t be the random variable indicating whether firm t is a friend ($Z_t = +1$) or not ($Z_t = -1$). Assume that the reviewer is indeed type **H** but that no one else knows this. As mentioned, if she reports truthfully, then her review for any firm (friend or not) is distributed as $r_t \sim \mathcal{N}(0, 2)$. Consumers, having learned that the reviewer is type **H**, would believe that the expected quality given a review r_t is equal to $\mathbb{E}[q_t | r_t, \hat{\omega} = \mathbf{H}] = r_t/2$ and therefore the average quality belief for *any* firm is

$$\mathbb{E}_{q_t}[\mathbb{E}[r_t | q_t]/2] = 0.$$

i.e. as expected, when the reviewer behaves truthfully, consumers do not prefer friends to non-friends, as both types get the same reviews on average. Instead of reporting truthfully however, the reviewer can pretend to be type **L**. She can do this in a specific way that benefits her friends as follows: she takes her signal s_t, injects noise ε_t', and writes a review $r_t' = s_t + \varepsilon_t'$, where $\varepsilon_t' = |X_t| \cdot Z_t$ and $X_t \sim \mathcal{N}(0, \sigma_L^2 - \sigma_H^2) = \mathcal{N}(0, 3)$. It is easy to see that $\varepsilon_t' \perp s_t$ and $\varepsilon_t' \sim \mathcal{N}(0, 3)$. Therefore, $r_t' \sim \mathcal{N}(0, 5)$. Importantly, r_t' is statistically unidentifiable from the random variable $r_t^L \sim \mathcal{N}(0, 5)$, which is the distribution of the reviews written by a truthtelling reviewer of type **L**. Consumers cannot distinguish whether the reviewer is indeed type **L**, or is a type **H** pretending to be type **L** (and biasing her friends in the process). All they can observe is a review distribution that looks like the bottom curve in Figure 1.

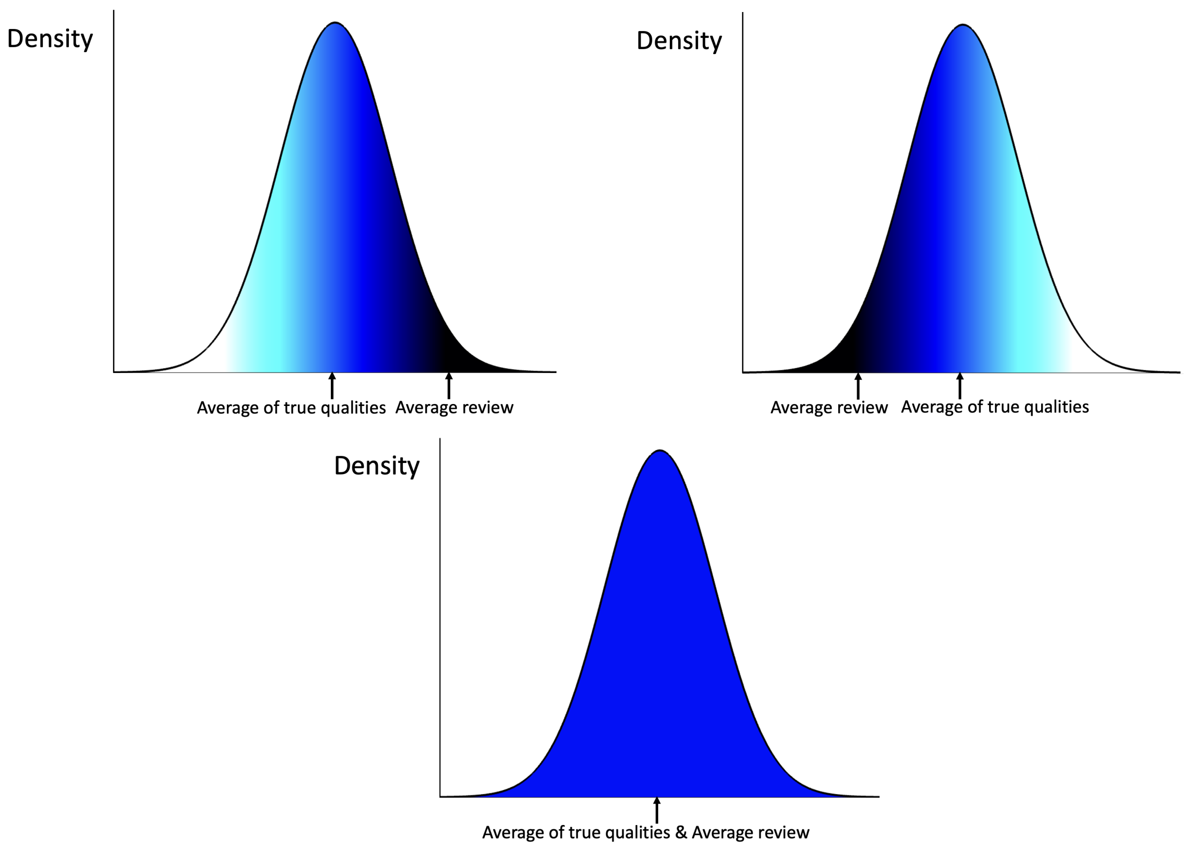

Figure 1. The top two curves display the review distributions for friends (left curve) and non-friends (right curve). The darker the shading, the more likely the review is to come from this part of the curve (compared to a fair review). Consumers cannot tell reviews are coming from the top curves – instead they can only observe the bottom distribution, which shows no bias.

Now, notice that even though the reviews over all firms are normally distributed, the reviews for friends are distributed as in the top- left curve of Figure 1. The average review for friends of quality q_t is therefore:

$$\mathbb{E}[r_t'|q_t] = \mathbb{E}[s_t + \varepsilon_t'|q_t] = \mathbb{E}[s_t + |X_t|\ |q_t] = q_t + \sqrt{\frac{6}{\pi}}$$

whereas the reviews for non-friends are distributed as in the right curve of Figure 1. The average review for non-friends of quality q_t is:

$$\mathbb{E}[r_t'|q_t] = \mathbb{E}[s_t + \varepsilon_t'|q_t] = \mathbb{E}[s_t - |X_t|\ |q_t] = q_t - \sqrt{\frac{6}{\pi}}$$

From the consumer's perspective, having learned that the reviewer is playing as type $\mathbf{L}$, she believes the expected quality given a review r_t' is equal to $\mathbb{E}[q_t|r_t', \hat{\omega} = \mathbf{L}] = \frac{r_t'}{1+\sigma_L^2} = r_t'/5$. This means that the average quality belief for friends is:

$$\mathbb{E}_{q_t}[\mathbb{E}[r_t'|q_t]/5] = \sqrt{\frac{6}{25\pi}}$$

Similarly, consumers (on average) believe the quality of non-friends to be $-\sqrt{\frac{6}{25\pi}}$. Because consumers cannot tell which firms are friends or not friends, they will inherently be biased toward friends of the reviewer, unless the reviewer's true type is actually **L**. $\square$

The rest of the paper generalizes the above example. In particular, being a friend arises endogenously through those firms who choose to offer the reviewer a bribe (and hence become friends). In the model, firms are myopic and bribes are not contractible, so it is possible that the reviewer collects the bribe and not follow through with a biased review. Firms, anticipating this, might not want to offer a bribe to begin with. However, we show that this game can sometimes admit an equilibrium where firms offer bribes in exchange for bias from reviewers. We then use our model to deliver additional insights about how consumer utility changes as a function of the parameters of the problem.

3 Model

We consider the following dynamic model. Firms are agents that produce and sell a product to a unit mass of consumers on $[0, 1]$. Firms arrive sequentially in discrete time, and market their product to consumers through a reviewer. The reviewer broadcasts her experiences of products to consumers and might do so in a strategic way to influence purchasing decisions.

3.1 Timing

At $t = 0$, the skill of the reviewer, ω, is drawn from a known Bernoulli distribution over the finite type space $\Omega = \{L, H\}$ (for Low and High-skill, respectively), where $\omega = H$ is drawn with probability p. These skill types correspond to the reviewer's precision, σ_ω, in assessing the firm's product, with $\sigma_H < \sigma_L$. The reviewer knows her actual type ω, but firms and consumers only know the type distribution.

The following sequence of events happens at every time $t \geq 1$:

(a) Firm t arrives with a product of random quality q_t, which is drawn from a standard normal distribution $q_t \sim \mathcal{N}(0, 1)$. With probability θ, the firm is honest and offers no bribe. With probability $1-\theta$ the firm is strategic and decides to either **Bribe** (by offering $b_t > 0$) or **Not Bribe** (by offering $b_t = 0$).

(b) The reviewer samples the product and receives an unbiased, noisy signal $s_t = q_t + \varepsilon_t$ of the quality, where $\varepsilon_t \sim \mathcal{N}(0, \sigma_\omega^2)$ and is i.i.d. across time and independent of q_t. The reviewer posts a review $r_t \in \mathbb{R}$ of the product, where r_t is not necessarily equal to s_t.

(c) Every consumer $i \in [0,1]$ observes the review r_t and then, based on r_t, chooses whether to purchase the good $x_{i,t} \in \{0,1\}$ at unit price.

We assume the true quality q_t is revealed through a public signal after each period and that the history of reviews $\{r_1, \ldots, r_{t-1}\}$ and true qualities $\{q_1, \ldots, q_{t-1}\}$ is publicly known. The history of bribes $\{b_1, \ldots, b_{t-1}\}$ on the other hand is private to the reviewer, and is not visible to consumers or firms. We refer to the above model as the *full dynamic model.*

3.2 Actions

Throughout the paper, we focus on a version of that model where all players have binary actions. Consumers either buy $(x_i = 1)$ or not $(x_i = 0)$, strategic firms either offer bribe $b_t = b^*$ or $b_t = 0$,[3] and the reviewer plays according to the commitment model we outline below. This is without loss of generality: for the interested reader, the connection to the more general strategy spaces of the full dynamic model is discussed in Section 3.5 and the formal equivalence is shown in Online Appendix B.

Commitment. We assume the reviewer commits at time 0 to playing an *effective* type $\hat{\omega} \in \{L, H\}$. This effective type is possibly distinct from her true type ω. A low-skilled reviewer $(\omega = L)$ must choose $\hat{\omega} = L$, since she cannot pretend to be more skilled than she actually is. The more interesting case is the high-skilled reviewer $(\omega = H)$ who can choose either $\hat{\omega} = H$ or $\hat{\omega} = L$.[4] If she chooses the latter, then her reports are less accurate compared to the case when she reports truthfully.

Firms and consumers observe the choice of effective type $\hat{\omega}$. While we take this observation as an assumption, the fact that a reviewer commits to an effective type $\hat{\omega} \in \{L, H\}$ means that agents will learn this type (i.e., how she actually plays over time) after observing enough history

[3]Here, b^* is taken as exogenous, but a endogenous choice of b_t (over a continuous action space) does not affect the main insights of the model, as explained informally in Section 3.5 and formally shown in Online Appendix B.

[4]The assumption that a reviewer can mimic lower types but not higher types is reminiscent of the reputation papers of Milgrom (1981) and Grossman (1981). We show the connection formally as an outcome of the fully dynamic model in Lemma B7, stated and proven in Online Appendix B. We discuss this connection informally in Section 3.5.

of past play. For this reason, we will use *effective type* and *reputation* interchangeably when we refer to $\hat{\omega}$.

If the type-H reviewer commits to $\hat{\omega} = H$, she writes (truthful) reviews $r_t = s_t$. However, if she commits to $\hat{\omega} = L$, then she distorts her reviews in a way that makes her statistically indistinguishable from a true type-L. The set of functions mapping s_t to r_t such that r_t is statistically identical to a type-L reviewer is not unique. We focus on a specific mapping, which is a generalization of the strategy in Section 2. In this mapping, which is parameterized by the proportion of honest firms θ, the type-H reviewer writes reviews $r_t = s_t + \varepsilon'_t$ in the following way:

- If the firm chooses **Not Bribe**, then the reviewer adds noise ε'_t from the bottom θ percentile of the $\mathcal{N}(0, \sigma_L^2 - \sigma_H^2)$ distribution.

- If the firm chooses **Bribe**, then the reviewer adds noise ε'_t from the top $(1-\theta)$ percentile of the $\mathcal{N}(0, \sigma_L^2 - \sigma_H^2)$ distribution.

Under the above mapping, bribing firms are given the most upward bias possible and non-bribing firms are given the most downward bias, subject to the reviewer having to appear as a type L.[5] This leads to the biggest gain from bribing out of all maps from s_t to r_t, and therefore provides the most incentive for firms to bribe.[6] Hence, to determine if bribes can exist in the system, it is enough to check if firms are willing to bribe under this mapping.

3.3 Payoffs

We now describe the payoffs of the three types of agents:

Consumers. Recall there is a continuum of consumers on $[0, 1]$ with consumer i having outside option $\phi(i)$. For simplicity, we make the following assumption:

Assumption 1. $\phi(i)$ is increasing in i, with $\lim_{i\to 0} \phi(i) = -\infty$ and $\lim_{i\to 1} \phi(i) = \infty$.

Assumption 1 guarantees that for all t, a positive fraction of consumers purchase the time-t product and a positive fraction do not. In particular, the map $\phi^{-1}(\alpha)$ is well-defined for any $\alpha \in \mathbb{R}$ and corresponds to the measure of consumers with $\phi(i) \leq \alpha$. Our results hold for any outside option distribution that satisfies Assumption 1.

[5]Formally, r_t must be statistically identical to one that is drawn from $r_t = q_t + \varepsilon_t$, with ε_t independent of q_t and distributed as $\mathcal{N}(0, \sigma_L^2)$, given that q_t is observed at the end of each period and can be compared to r_t.

[6]This is shown formally in Lemma A1.

Consumers who purchase the product experience the true quality q_t.[7] Consumers are myopic and maximize their current-period utility given the posted review r_t and the reputation of the reviewer $\hat{\omega}$ (which provides information about the accuracy of the reviewer); i.e., consumer i chooses her consumption $x_{i,t} \in \{0,1\}$ according to:

$$x^*_{i,t}(r_t, \hat{\omega}) = \arg \max_{x_{i,t} \in \{0,1\}} \mathbb{E}[(q_t - \phi_i)x_{i,t}|r_t, \hat{\omega}]$$

It is easy to see that each consumer will employ a cutoff strategy $x^*_{i,t}(r_t, \hat{\omega}) = 1$ iff $\mathbb{E}[q_t|r_t, \hat{\omega}] \geq \phi_i$. Therefore, we can define $X^*_t(r_t, \hat{\omega}) = \mu\left(\{i : x^*_{i,t}(r_t, \hat{\omega}) = 1\}\right)$, where μ is the Lebesgue measure, as the *total* consumption of the product. Note that $X^*_t(r_t, \hat{\omega}) = \phi^{-1}(\mathbb{E}[q_t|r_t, \hat{\omega}])$.

Reviewer. The reviewer is a far-sighted agent who discounts her payoff over time at rate $\delta \in (0,1)$. The reviewer's payoff reflects a desire to obtain (i) influence on consumer purchasing decisions and (ii) any potential bribes she can solicit.

(i) *Influence:* The influence I_t of the reviewer measures how much consumption reacts to her review on average. Observe that $\phi^{-1}(0)$ is the consumption when no review is written and $\phi^{-1}(\mathbb{E}[q_\tau|r_\tau, \hat{\omega}])$ is the consumption when review r_τ is written at time τ, then

$$I_t = \frac{1}{t}\sum_{\tau=1}^{t}(\phi^{-1}(\mathbb{E}[q_\tau|r_\tau, \hat{\omega}]) - \phi^{-1}(0))^2 \tag{1}$$

The above equation implies that the influence of the reviewer depends on how she was able to change average consumption patterns in the past.

We assume that the reviewer cares about her influence with propensity $\beta \geq 0$. The case $\beta > 0$ accounts for the possibility that the reviewer derives direct revenue and/or utility from her influence. This may come, for example, from advertisements on her dissemination channel or direct payments from her patrons.

(ii) *Bribes*: Reviewers may accept (non-contractual) bribes b_t from firms. These bribes are only observed by the reviewer and the bribing firm.

Reviewers maximize their long-term average payoff:

$$V = (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}(\beta \cdot I_t + b_t)$$

[7] Equivalently, we can assume that consumers, like the reviewer, get a noisy experience that is centered around the true quality, i.e., they get experience $e_{i,t} = q_t + \eta_{i,t}$ for some unbiased noise term $\eta_{i,t}$. The noise in the consumer's experience allows for the possibility that consumers do not all have identical ex-post utility from using the product. This does not change any of our results.

Firms. Recall that firms are either honest (never bribe) or strategic, with the former type occurring with probability θ. Firms receive a payoff proportional to the measure of active consumers $0 \leq X_t^* \leq 1$ (i.e., those who purchase the good) less any bribes they pay to the reviewer, b_t. Thus, we can write the total payoff of the time-t firm as $X_t^* - b_t$. Firms cannot observe bribes from previous periods, but can make inferences about how receptive the reviewer is to bribes based on her reputation $\hat{\omega}$. For example, a reviewer with reputation $\hat{\omega} = H$ is committed to reporting high-quality signals and therefore will not be receptive to bribes. Thus, a (strategic) firm at time t chooses b_t to maximize $\mathbb{E}[X_t^* - b_t | \hat{\omega}]$. Firms also have an outside option $\gamma > 0$ that they give up if they choose to enter the market.

3.4 Equilibrium Characterization

We consider the Bayesian Nash equilibria (BNE) of the aforementioned game, which may or may not be in mixed strategies. Recall that the type-H reviewer plays according to the strategy in Section 3.2, which maps signals s_t and bribes b_t to reviews r_t.

In the following result, we characterize the equilibrium strategies of the three agent types:

Theorem 1. *All equilibria are in pure strategies and take one of two forms:*

(i) Separating equilibrium*: The reviewer chooses $\hat{\omega} = \omega$, every firm bribes $b_t = 0$, and consumer i purchases if and only if*

$$\phi(i) \leq \frac{r_t}{1 + \sigma_\omega^2}$$

(ii) Pooling equilibrium*: The reviewer chooses $\hat{\omega} = L$, strategic firms bribe $b_t = b^*$, and consumer i purchases if and only if*

$$\phi(i) \leq \frac{r_t}{1 + \sigma_L^2}$$

under generic conditions.[8] *Moreover, the separating equilibrium always exists.*

Note that there is either just one equilibrium (separating) or two equilibria (separating and pooling). We call environments where only the separating equilibrium exists *bribe-proof*, because the equilibrium admits no bribes and both reviewer types report their signals truthfully. This is the first-best outcome relative to environments where the pooling equilibrium also exists. With the pooling equilibrium, consumers do not get first-best outcomes because a type-H reviewer distorts her signals by playing effective type L.

[8]By "generic conditions," we mean that for the set of initial parameters in the model, the subset where there are exceptions is a null set, i.e., has measure zero. Therefore, while other knife-edge cases might exist, perturbing the parameters by an arbitrarily small amount is guaranteed to eliminate these exceptions. Hence, these non-generic cases are not stable outcomes of the model, and thus will not be a primary focus.

We decompose the equilibrium strategies of Theorem 1 by agent type:

(i) *Consumers*: Based on the reputation of the reviewer, $\hat{\omega}$, the consumer obtains an estimate of the expected quality given the period-t review using the inverse-variance estimator:

$$\mathbb{E}[q_t|r_t, \hat{\omega}] = \frac{r_t/\hat{\sigma}^2}{1+1/\hat{\sigma}^2} = \frac{r_t}{1+\hat{\sigma}^2}$$

where $\hat{\sigma}$ is used as short-hand for $\sigma_{\hat{\omega}}$. Note that the 1 in the denominator is a consequence of assuming $q_t \sim \mathcal{N}(0,1)$.[9] The consumer thus tempers the review on both extremes: she believes very good reviews overhype the quality, whereas very negative reviews are overly harsh. Moreover, when the reviewer plays a more-skilled effective type (low $\hat{\sigma}$), all consumers place greater weights on the review and lower weights on their priors.

(ii) *Reviewer*: A type-L reviewer plays as her true type in both the separating and pooling equilibria. On the other hand, a type-H reviewer has a choice to play either $\hat{\omega} = H$ or $\hat{\omega} = L$. Based on her choice, either the separating or pooling equilibrium survives. This choice is in turn determined by whether the bribe b^* is sufficiently high to compensate for her reduced influence. The reduction in influence makes consumers pay less attention to reviews, which hurts the the reviewer in two ways: directly, through the decreased payoffs she gets from the loss in influence, and indirectly, through the reduction in the firms' willingness to bribe for biased reviews.

(iii) *Firms*: Strategic firms observe the reputation $\hat{\omega}$ of the reviewer and decide whether the demand boost they get from a biased review is enough to offer $b_t = b^*$. When $\hat{\omega} = H$, the firms recognize that the reviewer has a strong reputation for honest, high-quality reviews. This makes them abstain from bribing (the separating equilibrium). Conversely, when $\hat{\omega} = L$, the firm weighs the possibility that this is a type-H reviewer who accepts bribes and biases reviews (the pooling equilibrium) against the possibility that this an actual type-L reviewer who cannot misreport to help bribing firms (the separating equilibrium).

3.5 Discussion: Commitment Assumption and the Full Dynamic Model

A key assumption of our model is that the reviewer commits upfront (at $t = 0$) to a strategy mapping signals and bribes to reviews. This is a common assumption when the agent has

[9] This is a classic signal extraction problem (e.g. Cochran (1954)); take $Z = W + Y$ where $W \sim \mathcal{N}(0,1)$ and $Y \sim \mathcal{N}(0, \sigma_Y^2)$. Then $\mathbb{E}[W|Z=z] = \frac{z/\sigma_Y^2}{1+1/\sigma_Y^2}$ (i.e., the prior of $\mathcal{N}(0,1)$ regresses the estimate of W closer to zero).

reputational concerns over how the accuracy of her reviews are perceived.[10] We use the commitment assumption in our model for parsimony and to most transparently demonstrate our main ideas.

That being said, the main findings of Theorem 1 are robust to relaxing the commitment assumption In Online Appendix B.1, we present a more general dynamic model where the reviewer has full autonomy over how she writes reviews. While there may be many more equilibria in the fully dynamic setting, we present mild equilibrium refinements in Online Appendix B.2 that lead to the same equilibria of Theorem 1 (as shown in Appendices B.3 and B.4). A key finding is that the reviewer must mimic an effective type. This is because consumers *will* learn how the reviewer is playing and will stop listening to her if they realize she is inconsistent with a type-L or type-H reviewer playing truthfully (see Lemma B3). Moreover, a low-skilled reviewer ($\omega = L$) must play effective type $\hat{\omega} = L$ because she cannot remove noise from her signal to improve her estimate and make it on par with that of a type-H reviewer (Lemma B7). Thus, similar to the commitment model, it is the type-H reviewer who must choose whether to play $\hat{\omega} = H$ or $\hat{\omega} = L$.

Similarly, the assumption that the firm has a binary choice to either **Bribe** ($b_t = b^*$) or **Not Bribe** ($b_t = 0$), instead of choosing a continuous bribe $b_t \in \mathbb{R}_+$, is also without loss of generality. The equilibria of the full dynamic model of Appendix B.1, under the refinements of Appendix B.2, lead to all firms (who elect to bribe) choosing the same bribe b^* (see Lemmas B5 and B6).

Finally, the equivalence of the full dynamic model and the commitment model depends on the assumption of a patient reviewer (i.e., δ close to 1). This is a standard assumption in the reputation literature (e.g., Fudenberg and Levine (1989)), and to retain consistency with this model, we assume $\delta \to 1$ throughout the remainder of the paper. For the same reason, we look at the long-run outcomes in the commitment model, despite the fact that this is not necessary to derive our results.

4 How do Bribes Impact Consumer Utility?

We now examine the effect of bribes on consumer utility. A consumer's decision to purchase in period t is informed by her outside option $\phi(i)$, the review r_t, and the reviewer's reputation $\hat{\omega}$. The review induces a set of active consumers $\mathcal{A}_t$, and consumer i's utility is non-negative if she purchases the product whenever the quality q_t is at least equal to her outside option $\phi(i)$.

[10] See, for example, Section 1.C in Kamenica and Gentzkow (2011), where the reviewer in our paper is the "Sender" and each consumer is a "Receiver".

Definition 1. Consumer utility at time t is given by the average consumer utility of *active consumers* conditional on reviews r_t and reputation $\hat{\omega}$, $CU_t = \int_{i \in \mathcal{A}_t} (q_t - \phi_i)\, di$.[11] The average consumer utility at time t is the time-average of consumer utility up until time t, $\bar{CU}_t = \frac{1}{t} \sum_{\tau=1}^{t} CU_\tau$.

The following result establishes that average consumer utility and the reviewer's influence (from Equation (1)) converge to values that only depend on the effective accuracy $\hat{\sigma}$ of the reviewer.

Proposition 1. *The influence of the reviewer as well as average consumer utility converge (almost surely) as $t \to \infty$ to $I_\infty(\hat{\sigma})$ and $\bar{CU}_\infty(\hat{\sigma})$, respectively. These quantities only depend on and are strictly decreasing in $\hat{\sigma}$.*

From the previous section, an H-type reviewer plays either $\hat{\sigma} = \sigma_H$, which maximizes both her influence and consumer utility, or she plays $\hat{\sigma} = \sigma_L$, in which case both of these quantities are strictly less than their first-best levels.

An H-type reviewer therefore weighs the loss of influence from playing $\hat{\omega} = L$ against the bribes she receives from firms. This will determine whether the pooling equilibrium of Theorem 1 survives or if only the separating equilibrium exists. In the latter case, the environment is bribe-proof and delivers first-best levels of consumer utility. For the pooling equilibrium, consumer utility is below first-best because information about product quality is obscured by the reviewer.

4.1 Bribe-Proof Environments

We now characterize bribe-proof environments as a function of the parameters of the problem: the proportion θ of honest firms and the information the market has about the comparative skills of the reviewers and the likelihood of the different reviewer types. We study each of these parameters in the next three subsections.

4.1.1 Proportion of Honest Firms

We define the *fair consumption* (denoted by $\bar{X}^*$) as the ex-ante expected consumption for a firm when the reviewer is of low skill but guarantees a truthful review, i.e., $\mathbb{E}_{r_t \sim \mathcal{N}(0, 1+\sigma_L^2)} \left[\phi^{-1} \left(\frac{r_t}{1+\sigma_L^2} \right) \right]$. Firms enter the market if the expected consumption of their product is at least equal to their outside option γ. To avoid trivial cases where no firms enter, we assume that the fair

[11] Equivalently, one can measure total utility by summing over the outside option of inactive consumers and adding $q_t \cdot \mu(\mathcal{A}_t)$ (i.e., the utility of active consumers). These two quantities differ only by the constant $\int_0^1 \phi_i \, di$.

consumption of the firm is higher than the outside option.[12] We capture how appealing bribing is to the reviewer in the following definition:

Definition 2. The *bribing propensity*, ψ, is the ratio of the fair consumption (less the outside option) to the change in influence when a type H reviewer mimics type L, i.e., $\psi = (\bar{X}^* - \gamma)/(I_\infty(\sigma_H) - I_\infty(\sigma_L))$.

The bribing propensity ψ serves as a benefit/cost proxy for the reviewer: the numerator $\bar{X}^* - \gamma$ is the amount the strategic firms would gain from bribing (and therefore how much they are willing to pay in bribes) when all firms are strategic and the denominator $I_\infty(\sigma_H) - I_\infty(\sigma_L)$ is how much the reviewer will lose in influence by playing as a lower instead of a higher type. Thus, the higher this ratio is, the more appealing bribing becomes to the reviewer.

Note that when the reviewer has no intrinsic value for influence, i.e., $\beta = 0$, then bribes will always exist, since they are the only source of payoff for the reviewer. The next result shows that bribes can still exist even when the reviewer values her influence highly enough. This depends on the proportion of honest firms θ. In particular, the property of being bribe-proof is non-monotone in θ. When most firms are either strategic or honest (i.e., when θ is near 0 or 1), the environment is bribe-proof. However, intermediate values of θ may admit bribes:

Theorem 2. *There exists $0 < \underline{\theta} < \bar{\theta} < 1$ such that:*

(a) For all $\beta > \psi$ and $\theta < \underline{\theta}$ or $\theta > \bar{\theta}$, the setting is bribe-proof (i.e., only the separating equilibrium exists);

(b) For some $\beta > \psi$, probability of high type p, and bribe b^*, there exists $\underline{\theta} < \theta^* < \bar{\theta}$, such that the pooling equilibrium exists and consumer utility is strictly below first-best levels.

Theorem 2 describes a phase transition where, as the proportion of honest firms increases, the environment transitions from being bribe-proof to being susceptible to bribes and back again. Recall from Section 3.2 that the adjustment noise that the type H reviewer adds (to mimic type L) is drawn from a normal distribution that is parametrized by the proportion of honest firms θ. As θ becomes small enough and less than some $\underline{\theta}$, the bribing firm still receives a mostly fair review, because the reviewer cannot inject much bias when nearly all firms are strategic. This provides the firm with close to fair consumption less the outside option γ. Naturally, a bribe

[12]Even if that was not the case, a more complex (but equivalent) model guarantees that some non-bribing firms will still enter. In an effort to not over-complicate our model, we abstract away from the detail that firms may have *heterogeneous* outside options. This makes some non-bribing firms opt out of entering if they believe they will get reviews that are not biased in their favor, but still permits those with lower outside options to enter. We incorporate this into our model by assuming that the proportion θ is already calibrated so that it designates only those (non-bribing) firms that still find it profitable to enter and not bribe, regardless of whether these firms expect a bad review.

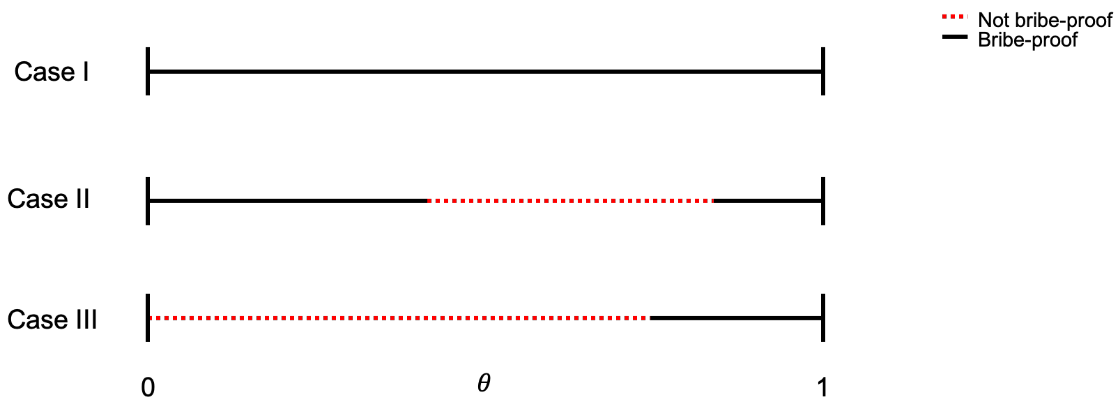

Figure 2. Bribe-proof environments as a function of proposrtion of honest firms θ.

that only provides a fair review (and fair consumption) does not offer much benefit to the firm. Consequently, in equilibrium, all firms offer $b_t = 0$ and the reviewer reports truthfully. Hence, the environment is bribe-proof.

At the other end, when $\theta > \bar{\theta}$, the reviewer can inject a lot of positive bias into the reviews of bribing firms. This makes firms willing to offer a bribe. However, because there are only few of these firms, it becomes unprofitable for the reviewer to trade off her influence with the bribes she (infrequently) collects. Instead, she opts to maximize her influence by reporting truthfully. This again leads to a bribe-proof environment.

The intermediate region of θ is where the reviewer and the strategic firms can come to an agreement. The reviewer incurs the reduced influence from playing type L, and the strategic firms find that the bias they receive is worth offering a bribe. Because θ is in an intermediate region, there are enough strategic firms, and the reviewer collects this bribe fairly often. This makes bribing profitable for both the reviewer and firms. As a result, the reviews misreports and the environment admits bribes.

Figure 2 presents the quintessential cases that can arise. When $\beta > \psi$, Theorem 2(a) applies and we know that either Case I or Case II hold (i.e., the extreme ends of θ are bribe-proof). In some cases, as identified in Theorem 2(b) and shown in Case II, there exists an intermediate range of θ where the environment may not be bribe-proof, even though the extremes are. In environments where $\beta < \psi$ (and Theorem 2 does not apply), other situations (e.g., Case III) can occur. For example, the property of being bribe-proof can be monotone in the fraction of honest firms θ, although Cases I and II can still occur in these environments.

Theorem 2 offers a warning for policymakers. It suggests that policy interventions that aim to reduce the number of firms that are willing to bribe (e.g., through threats of audits and punishments) can move the whole system from a bribe-proof environment (low θ) to a regime that admits bribes (intermediate θ). This reduces consumer utility and accomplishes the opposite goal of what the intervention is designed for. Similarly, even when the proportion of honest firms is high enough and the environment is bribe-proof, a small slip that reduces that proportion of these firms below $\bar{\theta}$ is enough for the environment to revert back to being vulnerable to bribes.

4.1.2 Disparity between Skill Types

Having established that consumer utility is non-monotone in the proportion of honest firms, we next show that it is also non-monotone in the difference of possible reviewer skills. Recall that σ_H and σ_L are the precisions of the high-skilled and low-skilled reviewers, respectively. Assume that the accuracy of the low-skilled reviewer improves (i.e., σ_L decreases) while σ_H remains the same. Does this make consumers better or worse off? We address this in the following result:

Proposition 2. *Let $\Delta = \sigma_L - \sigma_H$. The environment is bribe-proof as $\lim_{\Delta \to 0}$ or $\lim_{\Delta \to \infty}$.*

Proposition 2 shows that the answer to this question is not immediately obvious. When the skill of the low-type reviewer improves, the environment can transition from being bribe-proof to admitting bribes. This occurs because firms do not have enough incentives to bribe when the difference in skills between the reviewer types is too small ($\Delta \to 0$) or too large ($\Delta \to \infty$), but may have incentives to do so for intermediate values of Δ. We next describe the details for these three cases.

When the difference is too large (i.e., $\Delta \to \infty$), the type H reviewer can inject a lot of bias by mimicking a type L. Firms however are unwilling to bribe because even though bribing would deliver heavily favorable reviews, consumers pay very little attention to what the reviewer has to say because the type L reviewer being mimicked is extremely inaccurate. Similarly, when σ_H is fixed and $\sigma_L \to \infty$, the reviewer must sacrifice substantial influence in order to accept bribes, which requires a more significant bribe from the firm.

Assume that the skill of the type L reviewer increases (i.e., Δ shrinks) and she becomes more accurate. As before, the H-type reviewer can still help the bribing firms by mimicking type L and injecting bias into her reviews. The difference with the previous case is that consumers now pay more attention to the reviews because the type L being mimicked has reasonable accuracy. This makes bribing attractive to firms and the environment transitions into admitting bribes.

If the skill of the L-type reviewer increases further such that the difference between the types becomes too small (i.e., $\Delta \to 0$), then the bias an H-type reviewer can inject by mimicking an L-type is minimal. This is because she is mimicking a skill that is not that different from her own.

Firms are not interested in offering a bribe for a review that is too similar to the one they would get without a bribe, and the environment transitions back to being bribe-proof.

Proposition 2 highlights the fact that the property of being bribe-proof is more a matter of relative rather than absolute skill levels. The existence of bribes in the system depends on whether the reviewer can successfully imitate a less informative reviewer without losing her influence over consumer decisions altogether. It establishes that improvement in skill levels — which would normally improve the accuracy of information about quality— need not improve consumer utility. In fact, this skill improvement can lead to more information distortion when bribes exist in equilibrium.

4.1.3 Prior Beliefs about Skill

We conclude our discussion of bribe-proof environments by discussing how the incentives of players respond to the likelihood of the reviewer being highly-skilled. In this section, we show that a reviewer who is more likely to be low skilled can be better for consumers than a reviewer who is more likely to be high skilled.

Recall that the probability p indicates the likelihood that the true type of the reviewer is H. We consider what happens to consumer utility as p increases. When there are no bribing firms (i.e., $\theta = 1$), increasing p always improves consumer utility. Because all firms are honest, reviewers report truthfully (i.e., the separating equilibrium is unique) and consumers get more accurate recommendations. However, with the introduction of firms that might bribe, an increase in the likelihood of the type H reviewer also increases the potential for bribes and information distortion. This is because an actual type H reviewer is able to inflate the reviews of the bribing firms while an actual type L cannot. Thus, as p increases, the incentives for firms to bribe become stronger and the environment may no longer be bribe-proof. This is summarized in the following result:

Proposition 3 (Prior Beliefs)**.** *There exists p^* such that when $p < p^*$ the environment is bribe-proof (unique separating equilibrium) and consumer utility is increasing in p. However, when $p > p^*$ consumer utility is minimized (pooling equilibrium exists) and strictly below the first-best levels when $p < p^*$.*

Proposition 3 shows that an increase in the prior probability p of a high-type reviewer introduces two competing effects: it increases the potential for more precise reviewers, but it also increases the incentives of the firm to bribe and corrupt the reviews. Thus, while increasing p improves the quality of the information the reviewer receives (on average), it also encourages

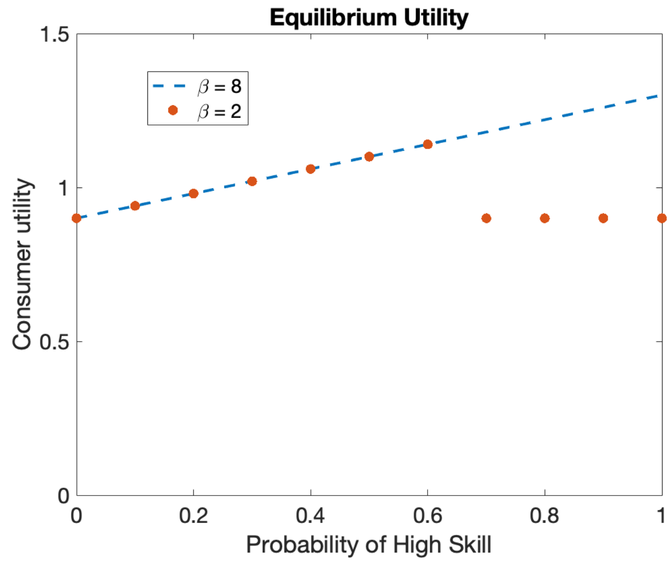

Figure 3. Equilibrium utility for varying ex-ante skill.

bribery and therefore can reverse the anticipated improvement in consumer utility.

The property that increasing the expected skills of the reviewer can lead to a *reduction* in (average) consumer utility can be seen in Figure 3. In this figure, p varies from 0 to 1 for different values of the influence propensity β of the reviewer. When $p = 0$, the reviewer is always low skill and no bribes are offered because firms have nothing to gain from bribing. When p is close to 1, the firm knows that the reviewer is likely a type H reviewer. Observing this reviewer playing $\hat{\omega} = L$ makes firms almost certain that this is a reviewer who is biasing reviews in exchange for bribes, which makes firms more likely to offer them.

Whether the environment is bribe-proof is determined by the value of β. When β is large (blue dashed line in Figure 3), the reviewer places a very high value on her influence and cannot be swayed by bribes. In this case, increasing p (and the expected accuracy of the reviewer) can only increase consumer utility. This is because consumers receive better information about product qualities and this information is never distorted.

However, when β is not too large (red dotted line), there is a more subtle effect from increasing p. As p increases, firms are more willing to bribe knowing there is a higher likelihood that the reviewer is a true type H playing as effective type L and biasing her reviews. Eventually (around $p = 0.6$ in the figure), the firm is willing to pay b^* as a bribe, and with smaller β, the bribe b^* is enough to entice the reviewer to "sell out" and choose $\hat{\omega} = L$. Before that point, increasing p increases the accuracy of reviews (as in the case of high β) and increases consumer utility. But once p hits the threshold, the firm becomes willing to bribe knowing the reviewer likely provides biased reviews in exchange for bribes ($\hat{\omega} = L$ even though $\theta = H$, in the pooling equilibrium).

This causes a precipitous drop to the lowest consumer utility at this threshold, as both skill types play identically as the low-type in equilibrium.

Figure 3 illuminates an important counterintuitive finding: improving the quality of information and efficacy of reviews may not necessarily lead to better outcomes for consumers. This is because firms can also leverage this improvement for their own gain, and together with reviewers, extract greater surplus through bribes at the cost of hurting consumers.

4.2 Total Welfare Analysis

Up until now, we have used consumer welfare as our metric for comparing good versus bad outcomes. In certain circumstances, it might be prudent to look at the *total* welfare of all market participants, i.e., the sum of utilities of consumers, reviewers, and firms. Note that the net utility of the reviewer is $\beta \cdot I_t$, in addition to any bribes that she may receive. Likewise, the net utility of the firm is the consumption X_t^* less any amount that is paid out to the reviewer in bribes. Thus, the bribes cancel out, and in addition to the consumer utility CU_t studied up until now, we also include the payoffs for the reviewers and the firms. Formally, the total welfare is given by:

$$W \equiv \frac{1}{t} \sum_{\tau=1}^{\infty} \left[\underbrace{\beta \cdot I_\tau + b_\tau}_{\text{Reviewer}} + \underbrace{X_\tau^* - b_\tau}_{\text{Firm}} + \underbrace{CU_\tau}_{\text{Consumer}} \right]$$

Given this, we obtain a generalization of the welfare analysis presented:

Proposition 4. *If ϕ is symmetric around the average consumer at $1/2$ (i.e., $\phi(1/2+z) = -\phi(1/2-z)$), then an equilibrium with bribes (i.e., pooling) has strictly lower welfare for firms, consumers, and in total than an equilibrium with no bribes (i.e., separating).*

In other words, Proposition 4 states that under mild conditions on ϕ, bribe-proof is best not only for consumers and the market as a whole, but for *the firms themselves.* This means that firms, on average, are hurt by the existence of bribes in the system. The reason is that when ϕ is symmetric, total firm consumption across both honest and strategic firms remains unaffected when the reviewer increases the variance of her reviews. Consumers simply temper their interpretation of the review in their purchasing decision and there is no net effect in the expected consumption over time. In a bribe-proof environment, firm welfare is given by consumption of its product. When there are bribes, firm welfare is given by consumption minus the bribe they have to pay. Because bribes are positive and average consumption (across all firms) remains the same, firm welfare decreases.

It is clear that honest firms are hurt when the environment is not bribe-proof, since their reviews are consistently biased down. Interestingly, the *bribing* firms themselves could also be hurt, and would prefer if bribes were eliminated from the system altogether. To see why this is the case, consider the following two settings where half the firms are honest and the other half are strategic (similar to the setting of Section 2):

- *Setting A*: This is the typical setting where firms can bribe in exchange for biased reviews (as in the pooling equilibrium of Theorem 1). A firm that bribes receives a review that is biased up (distributed as the top-left figure of Figure 1), whereas a firm that does not bribe receives a review that is biased down (distributed as the top-right figure of Figure 1). The difference in expected consumption from these two reviews determines whether a strategic firm will pay b^* to receive a biased review.

- *Setting B*: For the sake of demonstration, suppose all firms agree upfront not to bribe and that this agreement is enforceable. Firms receive an honest review from the distribution given in the bottom figure of Figure 1. This review provides a tighter estimate of the quality q_t for consumers, as the reviewer commits to $\hat{\omega} = H$ instead of injecting bias. Moreover, the firm does not pay a bribe to receive this review.

In Setting A, the firm may bribe simply to avoid a bad review. It may be the case that a fair review plus not having to pay a bribe (i.e., Setting B) is more profitable than the biased review that the firm solicits with bribe b^* (Setting A). However, the threat of the bad review is enough to force the strategic firms to bribe in equilibrium. In these cases, the bribe that the firm pays to the reviewer is a form of extortion: it offers no benefit to the firm, but rather saves it from being punished with a bad review. As a result, eliminating bribes can be a Pareto improvement for *all* firms in the market.

5 Final Remarks

Review manipulation is a problem that takes on multiple forms. In this paper, we focused on the case where an influential reviewer can misrepresent her reviews in order to favor specific firms. This decreases the utility of consumers as they purchase products of inferior quality and avoid products of high quality. The reviewer achieves this through a novel mechanism that takes advantage of the statistical learning process of consumers. This mechanism allows us to derive insights into the long-run nature of these interactions. For example, we show that an increase in the number of firms that are not willing to bribe (e.g., through a stringent auditing policy) might make the market more vulnerable to bribes and review distortion. We also show that improving the quality of the reviewer can result in worse outcomes for consumers.

An important factor that prevents bribes in our model is that the reviewer can generate payoffs from consumers via her influence (captured by the parameter β). This is a channel that is becoming more common in practice, as many online reviewers receive recurring payments directly from consumers through platforms like Patreon. Thus, a combination of low auditing and higher "wages" can improve welfare. This conclusion is similar to the literature on corruption, as in the work of Becker and Stigler (1974) and the empirical documentation in Di Tella and Schargrodsky (2003).

Our model can be extended in several directions. In an online appendix (Appendix B attached to this submission), we show that our results hold even when the reviewer does not commit to playing a specific strategy but can instead deploy any dynamic strategy she chooses. Under mild assumptions, we show that the reviewer must still play using the same commitment strategy as in the model in this paper, and all the insights from that model continue to hold.

Another extension is competition between reviewers. Our setup considered a monopolist reviewer. When there are multiple reviewers, it can be shown that a firm will not bribe more than a single reviewer in the same period. This is because the long-run dynamics will eventually show that biased reviews are correlated over time. Consequently, this will reveal which reviewers are accepting bribes and which firms are offering them. An unexpected consequence of competition is that a market with a monopolist reviewer may be bribe-proof but starts admitting bribes when another reviewer is added. This is because the new reviewer "steals" some influence from the monopolist reviewer, who now starts accepting bribes to make up for this decreased influence.[13]

Finally, we remark that our work is applicable beyond the setup of online reviews. Think tanks and research institutions that obtain funding from outside sources face similar tradeoffs: they want to produce reliable results for the public, but they risk cutting off their funding if these results do not advance the interests of the funding sources. Another area of application is the inspection and certification market, which is projected to be worth more than $400 Billion by 2025 (Bloomberg (2019)). These markets cover a wide range of industries (manufacturing, agriculture, healthcare, etc.) and can be modeled using the framework developed in this paper. Despite the importance of these interactions and their impact on all parties involved, their dynamic nature have not been analyzed or perfectly understood. We hope that the model in this paper serves as a first step in this direction.

[13] We provide intuition for the discussion in this paragraph in Appendices B.5.4 and B.5.5 of Online Appendix B.5.

Appendix

A Proofs

Throughout the proofs we will use the following shorthand notation: Let $\mathcal{N}(0, \sigma^2, \underline{z}, \bar{z})$ denote the conditional normal distribution with mean 0 and variance σ^2, but where the value is drawn between the $\underline{z}$ and $\bar{z}$ *percentiles* of the distribution.

A.1 Main Body

Proof of Theorem 1. Because the reviewer commits to and announces her effective type $\hat{\omega}$, her strategy of mapping signals and bribes to reviews is fixed. This means that the game is stationary. Thus, it is sufficient to consider one period in the equilibrium analysis.

The roadmap for the proof is as follows. First, we derive the consumers' best response, which depends only on the reviewer's choice of effective type $\hat{\omega}$ (her "effective accuracy"). Second, we show that the a high-type reviewer must give up influence when she elects $\hat{\omega} = L$ instead of $\hat{\omega} = H$. Of the pure strategy equilibria, we call the ones where $\hat{\omega} = \omega$ the *separating* case, and the ones with $\hat{\omega} = L$ for both types the *pooling* case. Third, we show that of the two possible separating cases (corresponding to **Not Bribe** and **Bribe**), only the **Not Bribe** case can be supported in equilibrium. Whereas of the two possible pooling cases, only **Bribe** can be supported. Fourth, we show that generically, no mixed-strategy equilibria can exist, so only the unique separating and unique pooling pure-strategy equilibria can exist. We conclude by proving that the separating equilibrium is in fact *always* an equilibrium (whereas the pooling one may or may not be).

Because consumers cannot observe the bribes and are myopic, they simply purchase the good if $\phi(i) \leq \mathbb{E}[q_t|r_t, \hat{\omega}]$. We note that consumers can compute $\mathbb{P}[q_t|r_t]$ (where $\mathbb{P}[.]$ represents the density) using Bayes' rule given that the reviewer writes $r_t = q_t + \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, \hat{\sigma}^2)$:[14]

$$\begin{aligned}
\mathbb{P}[q_t|r_t] = \frac{\mathbb{P}[r_t|q_t]\mathbb{P}[q_t]}{\int_{-\infty}^{\infty} \mathbb{P}[r_t|q]\mathbb{P}[q]\, dq} &= \frac{\exp\left(-\frac{(r_t-q_t)^2}{2\hat{\sigma}^2}\right)\exp\left(-\frac{q_t^2}{2}\right)}{\int_{-\infty}^{\infty} \exp\left(-\frac{(r_t-q)^2}{2\hat{\sigma}^2}\right)\exp\left(-\frac{q^2}{2}\right)\, dq} \\
&= \frac{1}{\sqrt{2\pi}\hat{\sigma}} \frac{\exp\left(-\frac{(r_t-q_t)^2}{2\hat{\sigma}^2}\right)\exp\left(-\frac{q_t^2}{2}\right)}{\exp\left(-\frac{r_t^2}{2(\hat{\sigma}^2+1)}\right)/\sqrt{\hat{\sigma}^2+1}} \\
&= \frac{1}{\sqrt{2\pi}}\sqrt{\frac{\hat{\sigma}^2+1}{\hat{\sigma}^2}} \exp\left(-\frac{\left(q_t - \frac{r_t}{1+\hat{\sigma}^2}\right)^2}{2\hat{\sigma}^2/(1+\hat{\sigma}^2)}\right)
\end{aligned}$$

which is the normal distribution density function with mean $\frac{r_t}{1+\hat{\sigma}^2}$ and variance $\frac{\hat{\sigma}^2}{1+\hat{\sigma}^2}$. Thus, $\mathbb{E}[q_t|r_t, \hat{\omega}] = \frac{r_t}{1+\hat{\sigma}^2}$, identifying the consumer's strategy stated in Theorem 1.

We now show that a type-H reviewer loses some influence when she plays $\hat{\omega} = L$. Let $\hat{\sigma}$ be the effective precision of the reviewer. Given the consumer's strategy we had just identified and

[14] Note that because the consumer cannot observe the bribes, this distribution depends only on the effective type $\hat{\omega}$ (with corresponding variance $\hat{\sigma}^2$) and not on the true type ω (or the true signal variance σ_ω^2). This is shown in Lemma A1 in Appendix A.2.

recalling that it depends on $\mathbb{E}[q_t|r_t, \hat{\omega}]$, we can write the reviewer's expected influence as

$$\mathbb{E}_{r_t}[I_t] = \frac{1}{\sqrt{2\pi(1+\hat{\sigma}^2)}} \int_{-\infty}^{\infty} \left(\phi^{-1}\left(\frac{r_t}{1+\hat{\sigma}^2}\right) - \phi^{-1}(0)\right)^2 \cdot \exp\left(-\frac{r_t^2}{2(1+\hat{\sigma}^2)}\right) \, dr_t$$

Let us make the substitution $\alpha = r_t/\sqrt{1+\hat{\sigma}^2}$:

$$\mathbb{E}_{r_t}[I_t] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \left(\phi^{-1}\left(\frac{\alpha}{\sqrt{1+\hat{\sigma}^2}}\right) - \phi^{-1}(0)\right)^2 \cdot \exp(-\alpha^2/2) \, d\alpha \tag{2}$$

Note that $\phi^{-1}(\cdot)$ is an increasing function, and $\left|\alpha/\sqrt{1+\hat{\sigma}^2}\right|$ is decreasing in $\hat{\sigma}$, so the integrand is decreasing in $\hat{\sigma}$ pointwise for all α. Thus, worse precision $\hat{\sigma}$ (i.e., higher $\hat{\sigma}$) leads to lower influence.

We use this result to analyze the game between the reviewer and firms. Note that the type-L reviewer has no choice in her strategy and the type-H reviewer has two choices. Thus, there are two potential (pure-strategy) maps from reviewer types ω to strategies $\hat{\omega}$. We call the map where $\hat{\omega} = \omega$ the *separating* case and $\hat{\omega} = L$ for all ω the *pooling* case. In the former, all reviewer types report truthfully and do not differentiate based on bribes, so **Not Bribe** is a best response for the firm. In the latter, all reviewer types play according to L. If the firm were to play **Not Bribe** with probability 1 this would not be an equilibrium: because $\beta > 0$ (generic conditions rule out $\beta = 0$), the type-H reviewer gives up influence when choosing $\hat{\omega} = L$, but does not get compensated with any bribes. Therefore, there is no equilibrium where the firms choose **Not Bribe** with probability 1 and the type-H reviewer plays $\hat{\omega} = L$. Thus, the pure pooling equilibrium is the following: the firm chooses **Bribe**, both types play $\hat{\omega} = L$, and the consumer buys if $\phi(i) \leq r_t/(1+\sigma_L^2)$.

Next, we show that generically, no mixed strategy equilibria exist. If the firm totally mixes between **Bribe** and **Not Bribe**,[15] it must be indifferent between **Bribe** and **Not Bribe**; in other words:

$$\mathbb{E}_{q_t,\varepsilon_t,\varepsilon_t' \sim \mathcal{N}(0,\sigma_L^2-\sigma_H^2,0,\theta)}\left[\phi^{-1}\left(\frac{r_t}{1+\hat{\sigma}^2}\right)\right] = \mathbb{E}_{q_t,\varepsilon_t,\varepsilon_t' \sim \mathcal{N}(0,\sigma_L^2-\sigma_H^2,\theta,1)}\left[\phi^{-1}\left(\frac{r_t}{1+\hat{\sigma}^2}\right)\right] - b^*$$

which occurs on a zero measure set. So generically, in the pooling case, the firm must choose **Bribe** with probability 1.

Next, we show that the type-H reviewer does not mix between $\hat{\omega} = H$ and $\hat{\omega} = L$ either. Suppose the $\omega = H$ reviewer plays $\hat{\omega} = H$ with probability $\zeta \in (0,1)$ and $\hat{\omega} = L$ with probability $1-\zeta$. From the same reasoning as before, when the reviewer plays $\hat{\omega} = L$, it must be the case that firms play **Bribe** with probability 1. If the reviewer is randomizing, then her expected payoff from playing $\hat{\omega} = L$ and $\hat{\omega} = H$ should be equal:

$$\mathbb{E}_{r_t}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\beta I_t + b_t \,\middle|\, \hat{\omega} = L\right] = \mathbb{E}_{r_t}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\beta I_t + b_t \,\middle|\, \hat{\omega} = H\right]$$
$$\implies \beta\mathbb{E}_{r_t}[I_t \,|\, \hat{\omega} = L] + b^* = \beta\mathbb{E}_{r_t}[I_t \,|\, \hat{\omega} = H]$$

where the expressions $\mathbb{E}_{r_t}[I_t \,|\, \hat{\omega} = L]$ and $\mathbb{E}_{r_t}[I_t \,|\, \hat{\omega} = H]$ are time-independent because of the stationarity of the game (and given by the expression in Equation (2)). This equality holds only

[15] By total mixing, we mean that the firm chooses each action with positive probability.

on a set of measure zero, so the reviewer has a strict preference for either $\hat{\omega} = H$ or $\hat{\omega} = L$, and only the two pure-strategy equilibria (separating and pooling) can exist.

Finally, we show that the separating case is always an equilibrium. Suppose that firms believe $\hat{\omega} = \omega$ is played in equilibrium. When the firm observes $\hat{\omega} = H$, it knows that this is a true type-H reviewer who reports truthfully. When the firm observes $\hat{\omega} = L$, it knows it is a true type-L reviewer who reports truthfully. In either case, it is a best response for the firm to select **Not Bribe**. Given the firm picks **Not Bribe**, the true type-H reviewer who selects $\hat{\omega} = L$ gives up influence but does not receive any bribes. This is not a profitable deviation, and hence the separating case is always an equilibrium. $\square$

Proof of Proposition 1. Because $\mathbb{E}[q_\tau | r_\tau, \hat{\omega}]$ are i.i.d. random variables for the consumers, we know by the strong law of large numbers that:

$$\frac{1}{t}\sum_{\tau=1}^{t}(\phi^{-1}(\mathbb{E}[q_\tau|r_\tau,\hat{\omega}]) - \phi^{-1}(0))^2 \overset{a.s.}{\to} \mathbb{E}_{q_\tau, r_\tau}\left[(\phi^{-1}(\mathbb{E}[q_\tau|r_\tau,\hat{\omega}]) - \phi^{-1}(0))^2\right] = I_\infty(\hat{\sigma})$$

where the final equality can be seen by simply integrating out q_τ and noting that $I_\infty(\hat{\sigma})$ for effective type $\hat{\omega}$ playing with precision $\hat{\sigma}$ is given exactly by $\mathbb{E}_{r_t}[I_t|\hat{\omega}]$ from Theorem 1. Thus, this implies $I_\infty(\hat{\sigma})$ is decreasing in $\hat{\sigma}$.
Similarly, for consumer utility, we can write:

$$\bar{CU}_t = \frac{1}{t}\sum_{\tau=1}^{t} CU_\tau = \frac{1}{t}\sum_{\tau=1}^{t}\int_0^1 (q_\tau - \phi(i)) \cdot \mathbf{1}_{\mathbb{E}[q_\tau|r_\tau,\hat{\omega}] \geq \phi(i)} \; di$$

which are i.i.d. random variables, so by the strong law of large numbers:

$$\bar{CU}_t \overset{a.s.}{\to} \int_0^1 \mathbb{E}_{q_\tau, r_\tau}\left[(q_\tau - \phi(i)) \cdot \mathbf{1}_{\mathbb{E}[q_\tau|r_\tau,\hat{\omega}] \geq \phi(i)}\right] \; di \equiv \bar{CU}_\infty(\hat{\sigma})$$

Any consumer i with outside option $\phi(i)$ is an active consumer at time t if and only if $\mathbb{E}[q_t|r_t,\hat{\omega}] = r_t/(1+\hat{\sigma}^2) \geq \phi(i)$, or in other words, $r_t \geq (1+\hat{\sigma}^2)\phi(i)$. The average utility as $t \to \infty$ can be measured as:

$$\bar{CU}_i = \frac{1}{2\hat{\sigma}\pi}\int_{-\infty}^{\infty}\int_{(1+\hat{\sigma}^2)\phi(i)}^{\infty} (q - \phi(i)) \exp(-q^2/2) \exp\left(-\frac{(r-q)^2}{2\hat{\sigma}^2}\right) \; dr \; dq$$

By Fubini's theorem, let us reverse the order of integration and evaluate:

$$\begin{aligned}\bar{CU}_i &= \frac{1}{2\hat{\sigma}\pi}\int_{(1+\hat{\sigma}^2)\phi(i)}^{\infty}\int_{-\infty}^{\infty} (q - \phi(i)) \exp(-q^2/2) \exp\left(-\frac{(r-q)^2}{2\hat{\sigma}^2}\right) \; dq \; dr \\ &= \frac{1}{\sqrt{2\pi}(1+\hat{\sigma}^2)^{3/2}}\int_{(1+\hat{\sigma}^2)\phi(i)}^{\infty} (r - (1+\hat{\sigma}^2)\phi(i)) \exp\left(-\frac{r^2}{2(1+\hat{\sigma}^2)}\right) \; dr\end{aligned}$$

Let us make the same change of variables $\alpha = r/\sqrt{1+\hat{\sigma}^2}$:

$$\begin{aligned}
\bar{CU}_i &= \frac{\sqrt{1+\hat{\sigma}^2}}{\sqrt{2\pi}(1+\hat{\sigma}^2)^{3/2}} \int_{\sqrt{1+\hat{\sigma}^2}\phi(i)}^{\infty} \left(\sqrt{1+\hat{\sigma}^2}\alpha - (1+\hat{\sigma}^2)\phi(i)\right) \exp\left(-\alpha^2/2\right) \, d\alpha \\
&= \frac{1}{\sqrt{2\pi(1+\hat{\sigma}^2)}} \int_{\sqrt{1+\hat{\sigma}^2}\phi(i)}^{\infty} \left(\alpha - \sqrt{1+\hat{\sigma}^2}\phi(i)\right) \exp(-\alpha^2/2) \, d\alpha
\end{aligned}$$

Making a final change of variables to $\kappa = \alpha - \sqrt{1+\hat{\sigma}^2}\phi(i)$, we have that:

$$\begin{aligned}
\bar{CU}_i &= \frac{1}{\sqrt{2\pi(1+\hat{\sigma}^2)}} \int_0^{\infty} \kappa \exp\left(-\frac{(\kappa + \sqrt{1+\hat{\sigma}^2}\phi(i))^2}{2}\right) \, d\kappa \\
&= \frac{1}{\sqrt{2\pi(\hat{\sigma}^2+1)}} \exp\left(-\frac{\phi(i)^2(1+\hat{\sigma}^2)}{2}\right) - \frac{1}{2}\phi(i)\mathbf{erfc}\left(\frac{\phi(i)\sqrt{1+\hat{\sigma}^2}}{\sqrt{2}}\right)
\end{aligned}$$

where erfc$(\cdot)$ is the complementary error function, erfc$(z) = \frac{2}{\sqrt{\pi}} \int_z^{\infty} e^{-x^2} dx$. Let us write $\zeta = \sqrt{1+\hat{\sigma}^2}$, so that

$$\bar{CU}_i = \frac{1}{\sqrt{2\pi}\zeta} \exp\left(-\frac{\phi(i)^2\zeta^2}{2}\right) - \frac{1}{2}\phi(i)\mathbf{erfc}\left(\frac{\phi(i)\zeta}{\sqrt{2}}\right)$$

Differentiating with respect to ζ, we obtain:

$$[\partial\zeta] : -\left(\sqrt{\frac{2}{\pi}} + \frac{1}{\sqrt{2\pi}\zeta^2}\right) \exp\left(-\frac{\phi(i)^2\zeta^2}{2}\right) < 0$$

Therefore, $\bar{CU}_i$ is decreasing in ζ for all i, which implies it is decreasing in $\hat{\sigma}$. Hence, the average consumer utility $\bar{CU}_\infty$ is also decreasing in $\hat{\sigma}$. $\square$

Proof of Theorem 2. We present the proofs of part (a) and part (b) separately:

(a) The firm's bribe is bounded above by $1-\gamma$, so the expected bribe (i.e., $(1-\theta)b^*$) received by the reviewer is at most $(1-\theta)\cdot(1-\gamma)$ in any pooling equilibrium. Because $\beta > 0$ and $I_\infty(\sigma_H) > I_\infty(\sigma_L)$ (see Proposition 1), it is clear that for $(1-\theta) < (1-\bar{\theta}) \equiv \frac{1-\gamma}{\beta(I_\infty(\sigma_H)-I_\infty(\sigma_L))}$, the high-type reviewer would prefer being truthful to mimicking $\hat{\omega} = L$ and receiving bribes (which holds for $\theta > \bar{\theta}$).

Similarly, when $\theta \to 0$, if the reviewer is high-skill mimicking low-skill, we see that the firm receives (net) expected consumption from bribing equal to $\bar{X}^* - \gamma$, since by Theorem 1, the firm receives review $r_t \sim q_t + \varepsilon_t + \varepsilon'_t$ where $\varepsilon'_t \sim \lim_{\theta\to 0} \mathcal{N}(0, \sigma_L^2 - \sigma_H^2, \theta, 1) = \mathcal{N}(0, \sigma_L^2 - \sigma_H^2)$, which implies that $r_t \sim \mathcal{N}(0, 1+\sigma_L^2)$ and total expected consumption is given simply by $\bar{X}^*$. On the other hand, if the firm does not bribe, it receives $r_t \sim q_t + \varepsilon_t + \varepsilon'_t$ where $\varepsilon'_t \sim \lim_{\theta\to 0} \mathcal{N}(0, \sigma_L^2 - \sigma_H^2, 0, \theta) = Dirac(-\infty)$, so total expected consumption is 0, which falls below the firm's outside option γ. Thus, the firm does not enter the market and receives a payoff of the outside option γ. The difference in these gives the maximal bribe of $\bar{X}^* - \gamma$. If $b^* > \bar{X}^* - \gamma$, the environment is trivially bribeproof, so assume that $b^* < \bar{X}^* - \gamma$. When the reviewer supplies bias in exchange for bribes, she loses an (average) influence payoff of:

$$\begin{aligned}
\beta(I_\infty(\sigma_H) - I_\infty(\sigma_L)) &> \psi(I_\infty(\sigma_H) - I_\infty(\sigma_L)) \\
&= \bar{X}^* - \gamma
\end{aligned}$$

As $\theta \to 1$ the environment is bribe-proof because the loss in influence payoff (strictly) exceeds the $\bar{X}^* - \gamma > (1-\theta)b^*$ for all θ. By continuity we can find a lower bound on θ, $\underline{\theta}$, such that the environment is bribe-proof for all $\theta < \underline{\theta}$.

(b) We will find θ^* for which the environment is not bribe-proof (i.e., a pooling equilibrium exists). For this, we find some β, p, b^* where this pooling equilibrium exists. It is enough to find β and p such that the highest bribe the "average" firm (i.e., $(1-\theta)b_t$ for bribing firms offering b_t) is willing to pay is lower than lowest bribe the reviewer is willing to accept (as then there exists some b^* in-between where a pooling equilibrium does exist).

For ease of notation, we call $\nu \equiv \frac{1}{2\pi\sqrt{(1+\sigma_H^2)(\sigma_H^2-\sigma_L^2)}}$. When the proportion of truthful firms is θ, note the total expected consumption obtained from bribing is given by:

$$\frac{1}{(1-\theta)\nu}\int_{-\infty}^{\infty}\int_{\sqrt{\sigma_H^2-\sigma_L^2}\cdot\Phi^{-1}(\theta)}^{\infty}\phi^{-1}\left(\frac{s+\varepsilon'}{1+\sigma_L^2}\right)\cdot\exp\left(-\frac{s^2}{2(1+\sigma_H^2)}\right)\cdot\exp\left(-\frac{(\varepsilon')^2}{2(\sigma_L^2-\sigma_H^2)}\right)\,d\varepsilon'\,ds$$

which we denote as $\kappa(\theta)$. On the other hand, the expected consumption from not bribing is given by:

$$\frac{1}{\theta\nu}\int_{-\infty}^{\infty}\int_{-\infty}^{\sqrt{\sigma_H^2-\sigma_L^2}\cdot\Phi^{-1}(\theta)}\phi^{-1}\left(\frac{s+\varepsilon'}{1+\sigma_L^2}\right)\cdot\exp\left(-\frac{s^2}{2(1+\sigma_H^2)}\right)\cdot\exp\left(-\frac{(\varepsilon')^2}{2(\sigma_L^2-\sigma_H^2)}\right)\,d\varepsilon'\,ds$$

which we denote by $\zeta(\theta)/\theta$. Note, the firm is willing to pay up to the bribe $\left(p\kappa(\theta)+\min\{-p\zeta(\theta)/\theta,(1-p)\bar{X}^*-\gamma\}\right)$, which means the average bribe from firms is $\bar{\mathcal{B}}(\theta) \equiv (1-\theta)\left(p\kappa(\theta)+\min\{-p\zeta(\theta)/\theta,(1-p)\bar{X}^*-\gamma\}\right)$. Differentiating with respect to θ and applying the fundamental theorem of calculus:

$$\begin{aligned}\frac{\partial\bar{b}^*(\theta)}{\partial\theta} &= -p\eta(\theta)+p\left[\frac{1}{\theta^2}\zeta(\theta)-\frac{1-\theta}{\theta}\eta(\theta)\right]\cdot\mathbf{1}_{\zeta(\theta)/\theta\geq\gamma/p-(1-p)/p\cdot\bar{X}^*}+(\gamma-(1-p)\bar{X}^*)\mathbf{1}_{\zeta(\theta)/\theta<\gamma/p-(1-p)/p\cdot\bar{X}^*}\\ &= \frac{p}{\theta}\cdot(\zeta(\theta)/\theta-\eta(\theta))\cdot\mathbf{1}_{\zeta(\theta)/\theta\geq\gamma/p-(1-p)/p\cdot\bar{X}^*}+(\gamma-(1-p)\bar{X}^*-\eta(\theta))\cdot\mathbf{1}_{\zeta(\theta)/\theta<\gamma/p-(1-p)/p\cdot\bar{X}^*}\end{aligned}$$

where

$$\eta(\theta)\equiv\frac{\partial\Phi^{-1}(\theta)/\partial\theta}{2\pi\sqrt{1+\sigma_H^2}}\cdot\exp\left(-(\Phi^{-1}(\theta))^2/2\right)\cdot\int_{-\infty}^{\infty}\phi^{-1}\left(\frac{s+\sqrt{\sigma_H^2-\sigma_L^2}\cdot\Phi^{-1}(\theta)}{1+\sigma_L^2}\right)\cdot\exp\left(-\frac{s^2}{2(1+\sigma_H^2)}\right)\,ds$$

and noting that $[(1-\theta)\kappa(\theta)]' = -\eta(\theta)$ and $\zeta'(\theta)=\eta(\theta)$. By the inverse function theorem, one can see that $\partial\Phi^{-1}(\theta)/\partial\theta = \sqrt{2\pi}\exp((\Phi^{-1}(\theta))^2/2)$, so the above reduces to:

$$\eta(\theta)=\frac{1}{\sqrt{2\pi(1+\sigma_H^2)}}\cdot\int_{-\infty}^{\infty}\phi^{-1}\left(\frac{s+\sqrt{\sigma_H^2-\sigma_L^2}\cdot\Phi^{-1}(\theta)}{1+\sigma_L^2}\right)\cdot\exp\left(-\frac{s^2}{2(1+\sigma_H^2)}\right)\,ds$$

Note that since $\lim_{x\to-\infty}\mathbb{E}[\phi^{-1}(y)|y\leq x]=0$, we know that $\lim_{\theta\to 0}\zeta(\theta)/\theta=0$. Similarly since $\lim_{\theta\to0}\Phi^{-1}(\theta)=-\infty$, we see $\lim_{\theta\to0}\eta(\theta)=0$. Therefore, fixing some small $\gamma>0$, we obtain on an open interval $\theta\in(0,\underline{\theta})$ where $\partial\bar{\mathcal{B}}(\theta)/\partial\theta>0$ for p sufficiently close to 1 (as then $\zeta(\theta)/\theta<\gamma/p-(1-p)/p\cdot\bar{X}^*$). This implies there exists $\theta^*\in(0,\underline{\theta})$ such that $\bar{\mathcal{B}}(\theta^*)>$

$\lim_{\theta\to 0}\bar{\mathcal{B}}(\theta)$. Choosing β such that:

$$\psi = \frac{\bar{X}^* - \gamma}{I_\infty(\sigma_H) - I_\infty(\sigma_L)} \equiv \frac{\lim_{\theta\to 0}\bar{\mathcal{B}}(\theta)}{I_\infty(\sigma_H) - I_\infty(\sigma_L)} < \beta < \frac{\bar{\mathcal{B}}(\theta^*)}{I_\infty(\sigma_H) - I_\infty(\sigma_L)}$$

implies the reviewer receives a large enough bribe under θ^* to compensate her for her loss in influence. Thus, the environment with θ^* is not bribe-proof. $\square$

Proof of Proposition 2. Recall the adjustment noise the reviewer injects is given by $\varepsilon'_t = \mathcal{N}(0, \sigma_L^2 - \sigma_H^2, 0, \theta)$ when she does not get a bribe and $\varepsilon'_t = \mathcal{N}(0, \sigma_L^2 - \sigma_H^2, \theta, 1)$ when she does. Observe that both of these converge to a Dirac-delta function at 0 as $\Delta \to 0$ (i.e., as $\sigma_H \to \sigma_L$). In particular, they converge to the same distribution, and thus $\mathbb{E}[X_t^*]$ converges to the fair consumption $\bar{X}^*$ regardless of the firm's action of **Bribe** or **Not Bribe**. Hence, for any fixed $b^* > 0$, for σ_H sufficiently close to σ_L, the firm has a best response to choose **Not Bribe**.

When the reviewer mimics low-type, we know the payoff from a bribe is given by:

$$\mathbb{E}_{q_t,\varepsilon_t,\varepsilon'_t\sim\mathcal{N}(0,\sigma_H^2-\sigma_L^2,\theta,1)}\left[\phi^{-1}\left(\frac{q_t+\varepsilon_t+\varepsilon'_t}{1+\sigma_L^2}\right)\right] - \mathbb{E}_{q_t,\varepsilon_t,\varepsilon'_t\sim\mathcal{N}(0,\sigma_H^2-\sigma_L^2,0,\theta)}\left[\phi^{-1}\left(\frac{q_t+\varepsilon_t+\varepsilon'_t}{1+\sigma_L^2}\right)\right]$$

If $\Delta \to \infty$, then it must be that $\sigma_L \to \infty$, and by Lebesgue's dominated convergence theorem (since ϕ^{-1} is bounded on $[0,1]$), we have that:

$$\begin{aligned}\lim_{\sigma_L\to\infty}\mathbb{E}_{q_t,\varepsilon_t,\varepsilon'_t\sim\mathcal{N}(0,\sigma_H^2-\sigma_L^2,\theta,1)}\left[\phi^{-1}\left(\frac{q_t+\varepsilon_t+\varepsilon'_t}{1+\sigma_L^2}\right)\right] &= \lim_{\sigma_L\to\infty}\mathbb{E}_{q_t,\varepsilon_t,\varepsilon'_t\sim\mathcal{N}(0,\sigma_H^2-\sigma_L^2,0,\theta)}\left[\phi^{-1}\left(\frac{q_t+\varepsilon_t+\varepsilon'_t}{1+\sigma_L^2}\right)\right]\\ &= \phi^{-1}(0)\end{aligned}$$

which implies that for any bribe $b^* > 0$, the firm prefers to take action **Not Bribe**. There are two cases. The first is that σ_H is bounded above as $\Delta \to \infty$. Since $\beta > 0$, $I_\infty(\sigma_H)$ is non-vanishing as $\Delta \to \infty$, so the high-type reviewer mimics her own type as $\sigma_L \to \infty$, which makes the environment bribe-proof and therefore attains first-best consumer utility. The second is that $\sigma_H \to \infty$ in which case the consumption is given by $X_t^* = \phi^{-1}(0)$ regardless of the firm's bribe, so for any $b^* > 0$, the firm's best response is **Not Bribe**. $\square$

Proof of Proposition 3. We show there exists p^* such that the environment is bribe-proof when $p < p^*$ and is not bribe-proof when $p > p^*$. This establishes the claim because consumer utility in the bribe-proof environment is $p\bar{CU}_\infty(H) + (1-p)\bar{CU}_\infty(L)$ (which is increasing in p) and consumer utility is just $\bar{CU}_\infty(L)$ in an environment that is not bribe-proof, with $\bar{CU}_\infty(H) > \bar{CU}_\infty(L)$ by Proposition 1.

Note the value of bribing for a firm is given by the difference between the firm entering and bribing and its next best option (either entering and not bribing, or staying out). Let $\underline{X}$ and $\overline{X}$ denote expected consumption when the reviewer biases the firm down and up, respectively, when she is actually the high-type reviewer but mimics the low type. If the firm enters and bribes, it receives $p\overline{X} + (1-p)\bar{X}^*$, whereas if the firm enters and does not bribe, it receives $p\underline{X} + (1-p)\overline{X}^*$, and if it stays out it receives γ. Therefore, the value of the bribe is given by:

$$\min\{p\overline{X} + (1-p)\bar{X}^* - (p\underline{X} + (1-p)\bar{X}^*), p\overline{X} + (1-p)\bar{X}^* - \gamma\} = p\overline{X} + \min\{-p\underline{X}, (1-p)\bar{X}^* - \gamma\}$$

The environment is bribe-proof if $p\overline{X} + \min\{-p\underline{X}, (1-p)\bar{X}^* - \gamma\} < b^* < \beta(I_\infty(\sigma_H) - I_\infty(\sigma_L))$ and otherwise admits bribes (per Theorem 1). In the event that $\beta(I_\infty(\sigma_H) - I_\infty(\sigma_L)) > b^*$, the reviewer never relinquishes her influence for a bribe (regardless of p), so the environment is always bribe-proof and taking any $p^* > 1$ establishes the claim.

Otherwise, $b^* < \beta(I_\infty(\sigma_H) - I_\infty(\sigma_L))$ and it is enough to check if $p\overline{X} + \min\{-p\underline{X}, (1-p)\bar{X}^* - \gamma\} < b^*$ to determine if the environment is bribe-proof. Since b^* is constant with respect to p but $p\overline{X} + \min\{-p\underline{X}, (1-p)\bar{X}^* - \gamma\}$ is increasing in p (as $\overline{X} > \bar{X}^* > \underline{X}$), if the inequality holds with a given p, it still holds with any $p' < p$, so the environment is still bribeproof for any $p' < p$ if it is bribeproof with p. By monotonicity, there exists a cutoff p^* where $p < p^*$ is bribe-proof but $p > p^*$ is not. $\square$

Proof of Proposition 4. The first-best level of welfare for the consumer was established in Proposition 1. Note that because ϕ is symmetric around 1/2, $\phi^{-1}(z) = 1 - \phi^{-1}(-z)$ for $z \geq 0$, or $\phi^{-1}(z) + \phi^{-1}(-z) = 1$ for all $z \geq 0$. Similar to Proposition 1 by strong LLN, we know that for the firms, the average consumption satisfies:

$$
\begin{aligned}
\lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} X_\tau^* &\overset{a.s.}{=} \mathbb{E}_{r_t\sim\mathcal{N}(0,1+\hat{\sigma}^2)}\left[\phi^{-1}\left(\frac{r_t}{1+\hat{\sigma}^2}\right)\right] \\
&= \frac{1}{\sqrt{2\pi}\hat{\sigma}}\int_{-\infty}^{\infty} \phi^{-1}\left(\frac{\alpha}{1+\hat{\sigma}^2}\right)\cdot \exp\left(-\frac{\alpha^2}{2\hat{\sigma}^2}\right) d\alpha \\
&= \frac{1}{\sqrt{2\pi}\hat{\sigma}}\int_{0}^{\infty} \left(\phi^{-1}\left(\frac{\alpha}{1+\hat{\sigma}^2}\right) + \phi^{-1}\left(\frac{-\alpha}{1+\hat{\sigma}^2}\right)\right)\cdot \exp\left(-\frac{\alpha^2}{2\hat{\sigma}^2}\right) d\alpha \\
&= \frac{1}{\sqrt{2\pi}\hat{\sigma}}\int_{0}^{\infty} \exp\left(-\frac{\alpha^2}{2\hat{\sigma}^2}\right) d\alpha \\
&= 1/2
\end{aligned}
$$

independent of $\hat{\sigma}$. Thus, regardless of which type the reviewer mimics, the average consumption of all firms is a constant.

In the expected welfare calculation for the firm, we have $\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t} X_\tau^*$ when the environment is bribe-proof (separating equilibrium). When the environment admits bribes (i.e., pooling equilibrium), firm welfare is given by $\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t} X_\tau^* - (1-\theta)b^*$, where $\lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t} X_\tau^*$ is the same as before by the previous argument. Because $b^* > 0$, firm welfare strictly decreases with the introduction of bribes.

Any bribes exchanged between the reviewer and the firm offset in the welfare calculation. Moreover, the reviewer obtains the most influence in a bribe-proof environment because she mimics her true type and not a lower type (with less influence). So the first-best welfare is obtained for the firm and the consumer in the bribe-proof environment (with strictly lower welfare for both in the pooling equilibrium). The bribe-proof environment always maximizes total welfare: the firm's expected consumption is unchanged, the consumer does worse, the reviewer's influence decreases, and the bribes are neither net positive or negative, because they are simply transfers between agents. $\square$

A.2 Supplemental Results

We show that the strategy the reviewer uses throughout the paper is the one that leads to the highest review bias for bribing firms. Let m be the map introduced in Section 3 for the type-H reviewer mimicking type-L. A type H reviewer who commits to playing $\hat{\omega} = L$ writes reviews $r_t = s_t + \varepsilon_t'$, where the additional noise ε_t' she injects into her signal is determined in the following way:

- If the firm chooses **Not Bribe**, then the reviewer adds noise $\varepsilon_t' \sim \mathcal{N}(0, \sigma_L^2 - \sigma_H^2, 0, \theta)$
- If the firm chooses **Bribe**, then the reviewer adds noise $\varepsilon_t' \sim \mathcal{N}(0, \sigma_L^2 - \sigma_H^2, \theta, 1)$

Lemma A1 (Mimicking Lemma)**.** *The difference in expected payoff from consumption (i.e., $\mathbb{E}[X_t^*]$) between actions* **Bribe** *and* **Not Bribe** *is maximized under mapping m, subject to the constraint that $\theta \cdot r_t($**Not Bribe**$) + (1-\theta) \cdot r_t($**Bribe**$)$ has the same joint distribution with q_t as a truthful type-L reviewer (i.e., $(r_t, q_t) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ where $\mathbf{\Sigma} = \begin{pmatrix} 1 & 1 \\ 1 & 1+\sigma_L^2 \end{pmatrix}$).*

Proof of Lemma A1. Step 1: Let $R : (s_t, b_t) \mapsto r_t$ be *any* reviewer strategy that maps signals s_t and bribes b_t to reviews r_t and satisfies the distributional constraints, i.e., $\tilde{R}(s_t) \equiv \theta R(s_t, 0) + (1-\theta)R(s_t, b^*)$ is constrained to be jointly normal with q_t, satisfying (1) $\mathcal{N}(0, 1+\hat{\sigma}^2)$ and (2) $\text{cov}(\tilde{R}(s_t), q_t) = 1$. Constraint (1) ensures the reviewer matches a truthful reviewer of some type $\hat{\omega}$ and condition (2) ensures it has the proper covariance with q_t (i.e., jointly normal variables with correlation $1/\sqrt{1+\hat{\sigma}^2}$). Because the consumers do not observe the bribes b_t, given that θ firms bribe 0 and $(1-\theta)$ firms bribe b^*, $\tilde{R}$ is the distribution that consumers observe over time. While consumers do not observe signals either, they can measure how $\tilde{R}$ compares to a truthful reviewer who trivially maps signals $s_t \mapsto R_t$ (but ignores any b_t).

First, note that $\tilde{R}(s_t) - s_t$ is normally distributed: it is the difference of two jointly normal random variables. Similarly, observe that $\text{cov}(\tilde{R}(s_t) - s_t, q_t) = \text{cov}(\tilde{R}(s_t), q_t) - \text{cov}(s_t, q_t) = 1 - 1 = 0$, so $\tilde{R}(s_t) - s_t$ is uncorrelated with q_t. Because both are normally distributed, this implies that they are in fact independent as well.

Hence, for any strategy $\tilde{R}(s_t) - s_t$ to be independent of q_t, it must be the case that $\tilde{R}(s_t) - s_t$ is independent of s_t, because all of these are jointly normal, but s_t and q_t are correlated. In summary, $\tilde{R}(s_t) - s_t$ must have the distribution $\mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2)$ and be orthogonal to s_t, for *any* choice of R for the reviewer. Therefore, we can write $\tilde{R}(s_t) = s_t + \varepsilon_t'$, where ε_t' is orthogonal to s_t and has distribution $\mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2)$. Rewritten:

$$\tilde{R}(s_t) = s_t + \theta(R(s_t, 0) - s_t) + (1-\theta)(R(s_t, b^*) - s_t)$$

Step 2: Next, we consider the map $m(s_t, 0) = s_t + \varepsilon_t'$ with $\varepsilon_t' \sim \mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2, 0, \theta)$ and $m(s_t, b^*) = s_t + \varepsilon_t'$ with $\varepsilon_t' \sim \mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2, \theta, 1)$ (and $\tilde{m}(s_t) \equiv \theta m(s_t, 0) + (1-\theta) m(s_t, b^*)$). It is easy to see $\tilde{m}$ satisfies the required distributional constraint (1):

$$\mathbb{P}[\tilde{m}(s_t) \le a] = \begin{cases} \theta \mathbb{P}[m(s_t, 0) \le a], & \text{if } a \le \sqrt{\hat{\sigma}^2 - \sigma_\omega^2}\,\Phi^{-1}(\theta) \\ \theta + (1-\theta)\mathbb{P}[m(s_t, b^*) \le a] & \text{if } a \ge \sqrt{\hat{\sigma}^2 - \sigma_\omega^2}\,\Phi^{-1}(\theta) \end{cases}$$

Note that $\mathbb{P}[m(s_t, 0) \le a] = \frac{1}{\theta}\mathbb{P}[\mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2) \le a]$ when $a \le \sqrt{\hat{\sigma}^2 - \sigma_\omega^2}\,\Phi^{-1}(\theta)$ and $\mathbb{P}[m(s_t, b^*) \le a] = \frac{1}{1-\theta}(\mathbb{P}[\mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2) \le a] - \theta)$ when $a \ge \sqrt{\hat{\sigma}^2 - \sigma_\omega^2}\,\Phi^{-1}(\theta)$. Thus, $\mathbb{P}[\tilde{m}(s_t) \le a] = \mathbb{P}[\mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2) \le a]$. It is also trivially satisfies distributional constraint (2) because ε_t' is drawn independently from s_t (and q_t).

Suppose that $m(s_t, b^*) - s_t$ first-order stochastically dominates all other distributions $R(s_t, b^*) - s_t$ subject to the distributional constraints that $\tilde{R}$ satisfy: (i) $\tilde{R}(s_t) - s_t$ is orthogonal to s_t; (ii) $\tilde{R}(s_t) - s_t \sim \mathcal{N}(0, \hat{\sigma}^2 - \sigma_\omega^2)$. Note this also implies that $m(s_t, 0)$ is first-order stochastically dominated by all other $R(s_t, 0) - s_t$ given distributional restrictions (i) and (ii). This means that $\mathbb{E}[X_t^* | b_t = b^*]$ under the reviewer strategy of m is <u>larger</u> than under any other strategy R that satisfies the constraints *and* $\mathbb{E}[X_t^* | b_t = 0]$ under the reviewer strategy of m is <u>smaller</u> than under any other strategy R satisfying the constraints. Thus, their difference is maximized amongst all such R with mapping m.

Step 3: It remains to prove FOSD of $m(s_t, b^*) - s_t$ over any other $R(s_t, b^*) - s_t$, as we do next. Consider some arbitrary $\tilde{R}(s_t)$ satisfying conditions (1) and (2) as before (where recall $\tilde{R}(s_t) = \theta R(s_t, 0) + (1-\theta)R(s_t, b^*)$). Note that for any $a < (\hat{\sigma}^2 - \sigma_\omega^2)\Phi^{-1}(1-\theta)$, then $\mathbb{P}[m(s_t, b^*) - s_t \geq a] = 1$ (i.e., for any a that is below the $(1-\theta)$th percentile of the normal distribution with variance $\hat{\sigma}^2 - \sigma_\omega^2$, $\varepsilon_t' = m(s_t, b^*) - s_t$ is necessarily drawn above a, by construction). Hence, $\mathbb{P}[m(s_t, b^*) - s_t \geq a] \geq \mathbb{P}[R(s_t, b^*) - s_t \geq a]$ trivially, because the latter is a probability measure between 0 and 1.

On the other hand, for any $a \geq \sqrt{\hat{\sigma}^2 - \sigma_\omega^2}\Phi^{-1}(1-\theta)$ note that:

$$\mathbb{P}[m(s_t, b^*) - s_t \geq a] = \frac{1 - \Phi\left(\frac{a}{\sqrt{\hat{\sigma}^2 - \sigma_\omega^2}}\right)}{1-\theta}$$

as the bottom $(1-\theta)$ fraction of the distribution occurs with probability 0, and the top $(1-\theta)$ fraction of the distribution is drawn proportional to the normal density, so the tail end of the distribution is amplified by $1/(1-\theta)$. Because of distributional constraints (i) and (ii) for $\tilde{R}(s_t) - s_t$, we know that $1 - \Phi\left(\frac{a}{\sqrt{\hat{\sigma}^2 - \sigma_\omega^2}}\right) = \mathbb{P}[\tilde{R}(s_t) - s_t \geq a]$, thus:

$$\begin{aligned}
&= \frac{1}{1-\theta}\mathbb{P}[\tilde{R}(s_t) - s_t \geq a] \\
&= \frac{1}{1-\theta}\left(\theta\mathbb{P}[\tilde{R}(s_t) - s_t \geq a | b_t = 0] + (1-\theta)\mathbb{P}[\tilde{R}(s_t) - s_t \geq a | b_t = b^*]\right) \\
&= \frac{\theta}{1-\theta}\mathbb{P}[\tilde{R}(s_t) - s_t \geq a | b_t = 0] + \mathbb{P}[\tilde{R}(s_t) - s_t \geq a | b_t = b^*] \\
&\geq \mathbb{P}[\tilde{R}(s_t) - s_t \geq a | b_t = b^*] \\
&= \mathbb{P}[R(s_t, b^*) - s_t \geq a]
\end{aligned}$$

where the second equality follows from the law of total probability, the third equality distributes the fraction $1/(1-\theta)$, the inequality follows from the fact that $\frac{\theta}{1-\theta}\mathbb{P}[\tilde{R}(s_t) - s_t \geq a | b_t = 0] \geq 0$ (i.e., probabilities are non-negative), and the final equality is by definition. This establishes that $m(s_t, b^*) - s_t$ first-order stochastically dominates *any* arbitrary reviewer strategy $R(s_t, b^*) - s_t$ satisfying the distributional constraints required, which completes the proof of Lemma A1. □

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar (2017), "Fast and slow learning from reviews." Technical report, National Bureau of Economic Research.

Bar-Isaac, Heski and Joyee Deb (2014), "(good and bad) reputation for a servant of two masters." *American Economic Journal: Microeconomics*, 6, 293–325.

Becker, Gary S and George J Stigler (1974), "Law enforcement, malfeasance, and compensation of enforcers." *The Journal of Legal Studies*, 3, 1–18.

Belavina, Elena, Simone Marinesi, and Gerry Tsoukalas (2018), "Designing crowdfunding platform rules to deter misconduct." *Available at SSRN 3093437*.

Bloomberg (2019), "Testing, inspection and certification market size worth $404.7 billion by 2025."

Bouvard, Matthieu and Raphael Levy (2017), "Two-sided reputation in certification markets." *Management Science*, 64, 4755–4774.

Candogan, Ozan and Kimon Drakopoulos (2017), "Optimal signaling of content accuracy: Engagement vs. misinformation."

Chen, Li and Yiangos Papanastasiou (2019), "Seeding the herd: Pricing and welfare effects of social learning manipulation." *Available at SSRN 3456139*.

Chen, Ying, Navin Kartik, and Joel Sobel (2008), "Selecting cheap-talk equilibria." *Econometrica*, 76, 117–136.

Chevalier, Judith A and Dina Mayzlin (2006), "The effect of word of mouth on sales: Online book reviews." *Journal of marketing research*, 43, 345–354.

Cochran, William G (1954), "The combination of estimates from different experiments." *Biometrics*, 10, 101–129.

Crawford, Vincent P and Joel Sobel (1982), "Strategic information transmission." *Econometrica: Journal of the Econometric Society*, 1431–1451.

Cripps, Martin W., George J. Mailath, and Larry Samuelson (2004), "Imperfect Monitoring and Impermanent Reputations." *Econometrica*, 72, 407–432. Publisher: [Wiley, Econometric Society].

Dellarocas, Chrysanthos (2006), "Strategic manipulation of internet opinion forums: Implications for consumers and firms." *Management science*, 52, 1577–1593.

Dellarocas, Chrysanthos, Xiaoquan Michael Zhang, and Neveen F Awad (2007), "Exploring the value of online product reviews in forecasting sales: The case of motion pictures." *Journal of Interactive marketing*, 21, 23–45.

Di Tella, Rafael and Ernesto Schargrodsky (2003), "The role of wages and auditing during a crackdown on corruption in the city of buenos aires." *The Journal of Law and Economics*, 46, 269–292.

Durbin, Erik and Ganesh Iyer (2009), "Corruptible advice." *American Economic Journal: Microeconomics*, 1, 220–42.

Dye, Ronald A. (1985), "Disclosure of Nonproprietary Information." *Journal of Accounting Research*, 23, 123–145, URL http://www.jstor.org/stable/2490910. Publisher: [Accounting Research Center, Booth School of Business, University of Chicago, Wiley].

Ershov, Daniel and Matthew Mitchell (2020), "The Effects of Influencer Advertising Disclosure Regulations: Evidence From Instagram." In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, 73–74, Association for Computing Machinery, New York, NY, USA, URL https://doi.org/10.1145/3391403.3399477.

Fainmesser, Itay P and Andrea Galeotti (2020), "The market for online influence." *American Economic Journal: Microeconomics*.

FTC (2016a), "Consumer review fairness act: What businesses need to know." URL https://www.ftc.gov/tips-advice/business-center/guidance/consumer-review-fairness-act-what-businesses-need-know.

FTC (2016b), "Warner bros. settles ftc charges." URL https://www.ftc.gov/news-events/press-releases/2016/07/warner-bros-settles-ftc-charges-it-failed-adequately-disclose-it.

Fudenberg, Drew and David Levine (1989), "Reputation and equilibrium selection in games with a patient player." *Econometrica*, 57, 759–78.

Fudenberg, Drew and David Levine (1992), "Maintaining a reputation when strategies are imperfectly observed." *The Review of Economic Studies*, 59, 561–579.

Grossman, Sanford J. (1981), "The Informational Role of Warranties and Private Disclosure about Product Quality." *The Journal of Law & Economics*, 24, 461–483, URL https://www.jstor.org/stable/725273. Publisher: [University of Chicago Press, Booth School of Business, University of Chicago, University of Chicago Law School].

Harsanyi, John C, Reinhard Selten, et al. (1988), "A general theory of equilibrium selection in games." *MIT Press Books*, 1.

Ifrach, Bar, Costis Maglaras, Marco Scarsini, and Anna Zseleva (2019), "Bayesian social learning from consumer reviews." *Operations Research*, 67, 1209–1221.

Inderst, Roman and Marco Ottaviani (2012), "Competition through commissions and kickbacks." *American Economic Review*, 102, 780–809.

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian Persuasion." *American Economic Review*, 101, 2590–2615.

Kreps, David M and Robert Wilson (1982), "Reputation and imperfect information." *Journal of economic theory*, 27, 253–279.

Lizzeri, Alessandro (1999), "Information revelation and certification intermediaries." *The RAND Journal of Economics*, 214–231.

Mailath, George J, J George, Larry Samuelson, et al. (2006), *Repeated games and reputations: long-run relationships*. Oxford university press.

Milgrom, Paul and John Roberts (1982), "Predation, reputation, and entry deterrence." *Journal of economic theory*, 27, 280–312.

Milgrom, Paul R. (1981), "Good News and Bad News: Representation Theorems and Applications." *The Bell Journal of Economics*, 12, 380–391, URL https://www.jstor.org/stable/3003562. Publisher: [RAND Corporation, Wiley].

Mitchell, Matthew (2021), "Free ad (vice): internet influencers and disclosure regulation." *The RAND Journal of Economics*, 52, 3–21.

Mostagir, Mohamed, Asu Ozdaglar, and James Siderius (2021), "When is society susceptible to manipulation?" *Management Science, forthcoming*.

Ottaviani, Marco and Peter Norman Sørensen (2006), "Reputational cheap talk." *The Rand journal of economics*, 37, 155–175. Publisher: Wiley Online Library.

Papanastasiou, Yiangos (2018), "Fake news propagation and detection: A sequential model."

Reinstein, David A and Christopher M Snyder (2005), "The influence of expert reviews on consumer demand for experience goods: A case study of movie critics." *The journal of industrial economics*, 53, 27–51.

Sen, Pranab K. and Julio M. Singer (1994), *Large Sample Methods in Statistics: An Introduction with Applications*. CRC Press.

Senecal, Sylvain and Jacques Nantel (2004), "The influence of online product recommendations on consumers' online choices." *Journal of retailing*, 80, 159–169.

Sobel, Joel (1985), "A theory of credibility." *The Review of Economic Studies*, 52, 557–573.

Sobel, Joel (2009), "Signaling games." *Encyclopedia of Complexity and Systems Science*, 19, 8125–8139.

B Online Appendix

We analyze the fully dynamic model that gives rise to the commitment model of Section 3. In this model, the reviewer has full autonomy over her strategy of reporting reviews and need not commit ahead of time to mimicking a truthful reviewer. Our work builds on classic models of reputation (e.g., Kreps and Wilson (1982); Milgrom and Roberts (1982); Sobel (1985); Fudenberg and Levine (1989, 1992); Mailath et al. (2006)), but unlike this literature, the reviewer in our setup plays two simultaneous reputation games on both sides of the market.

The presentation is as follows. In Appendix B.1, we layout the timing of the model in the same way as in Section 3.1. In Appendix B.2, we provide our equilibrium concept and the refinements that allow us to relate the equilibria of the fully dynamic model to that of Theorem 1. We conclude with Appendix B.4, which provides both intuitive proof "sketches" and the formal proofs relating the equilibria of the two settings.

B.1 Microfoundations: Full Dynamic Model

As before, at $t = 0$, the skill of the reviewer, ω, is drawn from a known Bernoulli distribution over the finite type space $\Omega = \{L, H\}$ (for Low and High-skill, respectively), where $\omega = H$ with probability p. With small probability $\pi_0 > 0$, the reviewer is a truthful (as opposed to strategic) reviewer who is behavioral and always reports $r_t = s_t$. The reviewer knows her own precision σ_ω, but no other agent (i.e. firm or consumer) does.

The following sequence of events happens at every time $t \geq 1$:

(a) Firm t arrives with a good of random quality q_t, which is drawn from a standard normal distribution $q_t \sim \mathcal{N}(0, 1)$ i.i.d. across time. The firm offers a bribe $b_t \in \{0\} \cup \mathbb{R}^+$ after observing the distribution of reviews $\{r_1, \ldots, r_{t-1}\}$

(b) The reviewer samples the product and receives an unbiased, noisy signal $s_t = q_t + \varepsilon_t$ of the quality, where $\varepsilon_t \sim \mathcal{N}(0, \sigma_\omega^2)$ and is i.i.d. across time and independent of q_t. The reviewer posts a review $r_t \in \mathbb{R}$ of the product, where r_t is not necessarily equal to s_t.

(c) Every consumer $i \in [0, 1]$ observes the review r_t and then, based on r_t and H_t (explained below), chooses whether to purchase the good $x_{i,t} \in \{0, 1\}$ at unit price. Consumers who elect $x_{i,t} = 1$ are *active* consumers.

Consumers and firms observe the *distribution* of public history H_t which includes reviews $\{r_1, \ldots, r_{t-1}\}$ and realized qualities $\{q_1, \ldots, q_{t-1}\}$ (but not bribes $\{b_1, \ldots, b_{t-1}\}$). The revelation of past q_t's can be thought of as a critical or consumer consensus that emerges after the product has been out in the marketplace for a while.

B.2 Microfoundations: Equilibrium Concept

We start by defining some of the notation used throughout the Online Appendix. As noted before, we let b_t, s_t, and H_t denote the bribe, signal, and public history at time t, respectively. We denote by π_t the belief of the consumes at time t about whether the reviewer is truthful. We denote by R_t the strategy of the reviewer, which is a stochastic map from (b_t, s_t, H_t) to a review r_t. We denote by R^H and R^L the strategy of a truthful reviewer who is type H and L, respectively. Similarly, we denote by F_t the strategy of the firm, which is a stochastic map from H_t to a bribe b_t, and $\mathcal{C}_{i,t}$ the strategy of consumer i, which is a stochastic map from (H_t, r_t) to a consumption decision $x^*_{i,t}$.

B.2.1 Perfect Bayesian Equilibrium

We focus on the Perfect Bayesian Equilibria (PBE) of two simultaneous games: one that occurs between the firm and the reviewer, and another that occurs between the reviewer and consumers. Recall that, by assumption, θ proportion of the firms are honest and do not bribe whereas $(1-\theta)$ proportion of firms are strategic. Formally, we consider the following equilibrium concepts:

1. *Consumer-reviewer equilibrium*: Consider a (fixed) strategy for a strategic firm $F_t(H_t)$, which may or may not be in pure strategies. Let $\tilde{F}_t(H_t)$ be the strategy of a firm who is ex-ante not known to be honest or truthful by the consumers, i.e.,

$$\tilde{F}_t(H_t) = \begin{cases} F_t(H_t), & \text{if firm at } t \text{ is strategic} \\ 0, & \text{if firm at } t \text{ is honest} \end{cases}$$

   The reviewer has a private history h_t that includes the realized bribes b_t. A *consumer-reviewer equilibrium* is then a perfect Bayesian equilibrium of the dynamic game between the reviewer and the consumer with bribes given by $\tilde{F}_t(H_t)$.

2. *Full equilibrium*: Taking as given the consumer-reviewer equilibrium, each (strategic) firm at time t chooses its strategy $F_t(H_t)$ to maximize its (current-period) profit.

The full equilibrium includes the (possibly time-dependent) strategies for the firms, the reviewer, and the consumers. For simplicity, we denote the full equilibrium strategies as a sequence of tuples $\chi_t \equiv \{F_t, R_t, \mathcal{C}_{i,t}\}$, with $F_t, R_t, \mathcal{C}_{i,t}$ denoting the (possibly mixed) strategies for the firm, reviewer, and consumer, respectively, at every point in time t.[16] In Appendix B.3, we show that such an equilibrium always exists for any discount factor δ.

B.2.2 Equilibrium Refinement

There is a multiplicity of PBEs for this dynamic game. For the purpose of refinement, we impose three conditions: (i) risk-dominance in the consumer-reviewer game, (ii) an indifference condition for firms to break ties in a consistent way across all t, and (iii) consumers face a very small cost for "peeking" at reviews. We discuss each of these in detail.

(i) Risk Dominance. The consumer-reviewer game is a cheap-talk game, which is a subset of signaling games. Unlike other signaling games, cheap talk games are more difficult to refine, with a sizable literature that proposes different selection arguments (see Chen et al. (2008); Sobel (2009)). We use *risk dominance* (Harsanyi et al. (1988)) as a refinement criterion for the consumer-reviewer game. Risk dominance is used to select among multiple equilibria. If a long-run player (in our case the reviewer) would prefer to play equilibrium A if short-run players (in our case the consumers) play according to equilibrium B than play equilibrium B if short-run players play according to equilibrium A, then equilibrium A is the risk-dominant equilibrium. This is still an equilibrium: all players play best responses to each other in Equilibrium A. However, the long-run player would prefer to err on the risk-dominant equilibrium, because if other players play according to a different equilibrium, it is still the "safest" option.

We next define formally the risk dominance concept for our consumer-reviewer game:

[16] Each mixed strategy is a distribution over the action space of that agent, as defined in Section 3.

Definition 3. An equilibrium profile $\{R_t, \mathcal{C}_{i,t}\}$ of the consumer-reviewer game is *risk-dominant* if the expected payoff for the reviewer under $\{R_t, \mathcal{C}'_{i,t}\}$ is more than the payoff under $\{R'_t, \mathcal{C}_{i,t}\}$ for all other equilibria $\{R'_t, \mathcal{C}'_{i,t}\}$.

"Babbling" (i.e., the reviewer reporting useless information, and the consumer ignoring it) is always an equilibrium in a cheap-talk game when there is no chance of a truthful reviewer (Crawford and Sobel (1982)). Thus, it can be used to threaten the reviewer with an eventual babbling equilibrium if she is detected as strategically manipulating reviews. Thus, once there is sufficient evidence that the reviewer is untruthful, consumers can revert to this *babbling-trigger* equilibrium and stop listening to her forever. This results in the reviewer eventually losing all influence, which in turn makes firms stop offering bribes, yielding the reviewer a long-term payoff of zero. The risk-dominant equilibria are the ones where the reviewer can ensure that she avoids such an outcome. Therefore, as we will show, in a risk-dominant sense, the only equilibria involve reviewers mimicking a truthful type of skill type H or L, as otherwise they risk being babble-triggered and losing all future payoffs.

Examples of equilibria that are not risk-dominant include cases where the reviewer manipulates her reviews, the consumers realize this, but still decide to listen to her. For instance, if the high-type has $\sigma_H = 1$ and the low-type has $\sigma_L = 3$, there may be an equilibrium (that is not risk-dominant) where the reviewer mimics some $\sigma' = 2$. These equilibria are less appealing because they require that consumers accept the fact they are being manipulated; that is, they continue to listen to recommendations that they know, with certainty, are advantaging some firms and not others. In addition, reviewers who are openly manipulative might get punished by regulatory bodies through large fines and bans, which can be equivalent to consumers ignoring the reviewer forever. Finally, these equilibria are significantly harder to analyze, and ex-ante offer no clear benefit of obtaining additional insights.

(ii) Indifference Condition. We assume that all firms break ties in a consistent way, for example, by choosing the lower of two bribes when indifferent. This is because it is possible under the current equilibrium concept for reviewers to offer bribe "schedules" (e.g., $10 for a 4-star review, $20 for a 5-star review), so that firms are indifferent between all bribes on the schedule, and so they could, in theory, randomize using the "correct" proportions in order to keep the reviewer safe from detection by the consumer.

However, these mixed equilibria are not robust to small amounts of noise in the firm payoffs. If some firms benefit (even minutely) from different bribes on the bribe schedule, or the reviewer's schedule contains small amounts of noise, then the firm will (almost) never randomize in equilibrium. Thus, small perturbations that differentiate the firms' preferences for reviews and bribes will make the reviewer unable to set a bribe schedule that always guarantees perfect indifference. This kills off the broader space of review manipulation possible in equilibrium.

(iii) Peeking Cost. We assume consumers face a small cost $\varepsilon > 0$ for looking at a review. This eliminates equilibria where consumers continue to return to read reviews even when they know for a fact the review contains no pertinent information about the product's quality. We send $\varepsilon \to 0$ so that this cost is not a substantial factor for the consumer, but is not zero, so it comes into play when consumers realize reviews are useless. Implicitly, this implies that consumers who stop reading product reviews also no longer observe the public history H_t.

Robust Equilibrium. We say an equilibrium is *robust* if (i) it is a full equilibrium, and (ii) it satisfies the aforementioned three conditions.

B.2.3 Time-Invariant (Limit) Strategy Profiles

Because the equilibrium strategies of the agents may be time-varying, characterizing the entire history of equilibrium play can be complex. Thus, we investigate whether there exist any time-invariant strategy profiles that closely approximate the on-path play of the agents in equilibrium. This provides a cleaner expression of how agents play the dynamic game. Moreover, it offers more straightforward intuition about our results, without having to worry about minor deviations from this profile.

We define this formally. Let $\chi_t[\delta] \equiv \{F_t[\delta], R_t[\delta], \mathcal{C}_{i,t}[\delta]\}$ be a sequence of (full) equilibrium strategies for the sequence of parameters δ converging to 1. Such a sequence is guaranteed to exist per the existence result in Appendix B.3, Theorem B1. A time-invariant strategy profile is given by $\chi^* \equiv (F^*, R^*, \mathcal{C}_i^*)$, which does not have any dependence on time t or discount factor δ. Informally, we say a time-invariant profile *closely approximates* an equilibrium sequence if, as $\delta \to 1$ and $t \to \infty$ (in that order), the distance between the equilibrium strategy profile $\chi_t[\delta]$ and the time-invariant profile χ^* (as defined by the ∞-norm[17]) converges to 0. More precisely:

Definition 4. We say the equilibrium strategy profile χ^* is a *time-invariant (limit) equilibrium* if $\lim_{t\to\infty} \lim_{\delta\to 1} ||\chi_t[\delta] - \chi^*||_\infty = 0$.

We can also define the notion of a "unique" time-invariant strategy profile, where all full equilibrium sequences have the same time-invariant (limit) equilibrium:

Definition 5. We say there is a unique time-invariant (limit) equilibrium if any sequence of full equilibria $\chi_t[\delta]$ converges (in the $\lim_{t\to\infty} \lim_{\delta\to 1}$ order) to the same limit equilibrium χ^*.

A unique time-invariant (limit) equilibrium profile χ^* suggests that we can pin down the play of firms, consumers, and the reviewer when the reviewer is very patient and we look far into the future. Note we can also apply Definition 5 to robust equilibria: we say there is a unique robust time-invariant (limit) equilibrium if any sequence of robust equilibria $\chi_t[\delta]$ converges to the same limit equilibrium χ^*.

While we focus on the long-term play of the agents in our model, the short-term dynamics are also interesting, but less tractable to characterize. It is possible that the reviewer might solicit large bribes for extremely biased reviews in the short-term or "showboat" to demonstrate to firms she is an actual type H reviewer. As $\delta \to 1$, however, we note that the additional payoff the reviewer gets from these deviations from the limit profile χ^* vanishes. For this reason, we focus entirely on the profile χ^*.

B.3 Microfoundations: Equilibrium Existence and Uniqueness

In this section, we discuss conditions under which the consumer-reviewer equilibrium, full equilibrium, and robust equilibrium exist and are unique. We prove the following results in Appendix B.4. By standard existence results, we have:

Theorem B1 (Existence)**.** *For any δ and any fixed firm strategy $F_t(H_t)$, a consumer-reviewer equilibrium exists; moreover, a full equilibrium and robust equilibrium always exist.*

[17] Recall that strategies are probability distributions over actions (or equivalently, random variables from actions to real numbers), and the ∞-norm measures the differences in probability distributions. For random variables Y_1 and Y_2 defined on probability space $(\Omega, \mathcal{F}, P)$, the ∞-norm is defined by $||Y_1 - Y_2||_\infty \equiv \sup_{B\in\mathcal{F}} |P(Y_1^{-1}(B)) - P(Y_2^{-1}(B))|$. In our case, $\Omega = \mathbb{R}^2 \times ([0,1] \times \{-1,1\})$ and $\mathcal{F} = \mathcal{B}^2 \times (\mathcal{B} \cap [0,1] \times 2^{\{0,1\}})$.

The first part of Theorem B1 can be proven using standard existence arguments, as per the reputation papers of Fudenberg and Levine. The second part of Theorem B1 can be proven by construction, as there is always a bribe-free equilibrium. To see this, note that there is a *lack of commitment* by the reviewer: if the firm bribes the reviewer, she is not committed to giving any particular review. This gives rise to a multiplicity of equilibria, but in particular, a trivial equilibrium where there are no bribes and all reviewers report truthfully. On the other hand, we say a full equilibrium is *bribing* if there exist bribes (almost surely) at some point along the history of play. Under certain conditions, there may be no equilibria other than the trivial one, in which case we say the environment is *bribe-proof.* As our central focus is on determining whether an environment is bribe-proof or not (i.e., if there are bribing equilibria too), we concentrate on bribes which are maximally-supported between the firm and the reviewer, as we define next:

Definition 6. (Maximal Equilibrium) We say a time-invariant (limit) equilibrium $(F^*, R^*, \mathcal{C}_i^*)$ is *maximal* if F^* has the largest expected bribe[18] of any time-invariant (limit) equilibrium.

In other words, maximality captures the largest bribe amount firms would be willing to give to reap the benefits of the reviewer's favoritism. It is then both necessary and sufficient to check whether the maximal equilibrium has bribes in order to know whether the setting is entirely bribe-proof. One can also apply Definition 6 to *maximal robust* equilibria, in which case the condition needs only hold for the set of robust equilibria.

Lastly, we obtain our main existence and uniqueness result:

Theorem B2 (Existence and Uniqueness of a Maximal Robust Equilibrium)**.** *Under generic conditions, there is a unique maximal robust time-invariant (limit) equilibrium* $(F^*, R^*, \mathcal{C}_i^*)$.

When the environment is bribe-proof, the unique maximal time-invariant robust equilibrium in the limit coincides with the trivial equilibrium. This equilibrium corresponds to the separating equilibrium of Theorem 1 when the environment is bribe-proof, and otherwise corresponds to the largest b^* such that a pooling equilibrium exists.

B.4 Microfoundations: Proofs

We first provide a sequence of intermediate results that we use throughout the remainder of the proofs. In the proof of each result, we first provide a clearly-marked intuition section before launching into the formal mathematical proof.

Lemma B1 (Cripps et al. (2004) Generalization)**.** *Let $R_t[\delta]$ be the strategy of a (strategic) reviewer with discount factor δ at time t. Then the strategic reviewer always either* mimics *an effective type $\hat{\omega} \in \{L, H\}$ or her truthfulness type is discovered, in the sense that in every consumer-reviewer equilibrium,* $\lim_{t\to\infty} \lim_{\delta\to 1} \pi_t(1-\pi_t)||R_t[\delta] - R^{\hat{\omega}}||_\infty = 0$ *almost surely for all* $\hat{\omega} \in \{L, H\}$.

Intuition. Over time, one of two things will happen: either consumers learn, almost surely, that the reviewer is (not) truthful (corresponding to $\pi_t \to 0$ or $\pi_t \to 1$), or the reviewer plays arbitrarily close to a truthful type's strategy (thus making the two types statistically indistinguishable). This is because the reviewer cannot simultaneously misrepresent herself by consistently playing

[18] Note that by the indifference condition refinement, the maximal bribe always involves a pure strategy for the firm, and so the maximal equilibrium bribe is just the maximum bribe that can be maintained in equilibrium. Thus, "expected" is superfluous when considering only robust equilibria, but makes the definition well-posed for any set of equilibria (e.g., all full equilibria).

differently from the truthful type *and* sustain her reputation as playing exactly like the truthful type.

Proof. We use the same merging-style argument presented in Cripps et al. (2004). Let $\pi_t^{\hat{\omega}}$ denote the belief of (Bayesian) consumers that the reviewer is a truthful type with accuracy type $\hat{\omega}$. Note that $\pi_t = \pi_t^L + \pi_t^H$. We prove that regardless of δ, $\lim_{t\to\infty} \pi_t^{\hat{\omega}}(1-\pi_t^{\hat{\omega}})||R_t[\delta] - R^{\hat{\omega}}||_\infty = 0$ for both $\hat{\omega} \in \{L, H\}$ which proves the claim. Fix a realization of the review written at time t as r_t. By Bayes' rule, we have:

$$\pi_{t+1}^{\hat{\omega}}(r_t) = \frac{\pi_t^{\hat{\omega}}\mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}]}{\pi_t^{\hat{\omega}}\mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}] + (1-\pi_t^{\hat{\omega}})\mathbb{P}[r_t|\text{ reviewer plays } R_t]}$$

Rearranging,

$$\frac{\pi_{t+1}^{\hat{\omega}}(r_t)}{\pi_t^{\hat{\omega}}}\left(\pi_t^{\hat{\omega}}\mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}] + (1-\pi_t^{\hat{\omega}})\mathbb{P}[r_t|\text{ reviewer plays } R_t]\right) = \mathbb{P}[r_t|\text{ reviewer plays } R_t^{\hat{\omega}}]$$

Similarly, we also have:

$$\frac{1-\pi_{t+1}^{\hat{\omega}}(r_t)}{1-\pi_t^{\hat{\omega}}}\left(\pi_t^{\hat{\omega}}\mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}] + (1-\pi_t^{\hat{\omega}})\mathbb{P}[r_t|\text{ reviewer plays } R_t]\right) = \mathbb{P}[r_t|\text{ reviewer plays } R_t]$$

Taking the difference of these expressions yields:

$$\left|\frac{\pi_{t+1}^{\hat{\omega}}(r_t)}{\pi_t^{\hat{\omega}}} - \frac{1-\pi_{t+1}^{\hat{\omega}}(r_t)}{1-\pi_t^{\hat{\omega}}}\right|\left(\pi_t^{\hat{\omega}}\mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}] + (1-\pi_t^{\hat{\omega}})\mathbb{P}[r_t|\text{ reviewer plays } R_t]\right)$$
$$= \left|\mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}] - \mathbb{P}[r_t|\text{ reviewer plays } R_t]\right|$$

Note that $\left(\pi_t^{\hat{\omega}}\mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}] + (1-\pi_t^{\hat{\omega}})\mathbb{P}[r_t|\text{ reviewer plays } R_t]\right) \leq 1$, therefore:

$$\max_{r_t} |\pi_{t+1}^{\hat{\omega}}(r_t) - \pi_t^{\hat{\omega}}| \geq \pi_t^{\hat{\omega}}(1-\pi_t^{\hat{\omega}})\left|\mathbb{P}[r_t|\text{ reviewer plays } R_t] - \mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}]\right|$$

Note that $\pi_t^{\hat{\omega}}$ is a bounded martingale (consumers update their beliefs about the truthfulness of the reviewer in a Bayesian way by observing the full public history H_t), so it converges almost surely by the martingale convergence theorem. In particular, $\pi_t^{\hat{\omega}}$ is a Cauchy sequence almost surely. Thus, the left-hand side converges to zero, which implies that the right-hand side must converge to zero almost surely. Thus, there exists no set of positive measure B such that for all $r_t \in B$, $\limsup_{t\to\infty} \pi_t^{\hat{\omega}}(1-\pi_t^{\hat{\omega}})\left|\mathbb{P}[r_t|\text{ reviewer plays } R_t] - \mathbb{P}[r_t|\text{ reviewer plays } R^{\hat{\omega}}]\right| > 0$, and the expression we want to prove follows immediately from the definition of the ∞-norm. □

Lemma B2. *In every robust equilibrium, there exists some π^* such that if $\pi_t < \pi^*$, the reviewer receives a payoff of zero for all periods after t.*

Intuition. In cheap talk games, babbling is always an equilibrium between a sender and receiver: the sender sends useless information, the receiver ignores it, and neither has an incentive to deviate. Therefore, in the consumer-reviewer game, once consumers becomes sufficiently convinced that the reviewer is strategic (via review manipulation), it is an equilibrium for them to *babble-trigger* the reviewer by ignoring the review altogether. The reviewer then has no influence and cannot solicit bribes. This property is true in every risk-dominant equilibrium.

To see this, note that a strategic reviewer gets a strictly positive payoff (through her influence) by mimicking a truthful reviewer of her same accuracy, *even if consumers would not play babbling-trigger*. However, if the reviewer exposes herself as (likely) strategic (through the public belief of her being truthful becoming sufficiently small, i.e., $\pi_t < \pi^*$), then babbling-trigger gives her an average-discounted payoff near 0 when she is patient. This is the worst payoff possible for the reviewer; thus, the risk-dominant equilibrium involves consumers playing babbling-trigger.

Proof. First, we show that there exists some π^* and a full equilibrium where once $\pi_t < \pi^*$, consumers do not pay the peeking cost to read the review anymore (i.e., the "babbling-trigger" equilibrium). Consider an equilibrium where once $\pi_t < \pi^*$, the strategic reviewer writes an orthogonal review (e.g., $r_t = 0$ always) regardless of any bribe b_t, her signal s_t, or any public history H_t. When reading a review from a truthful reviewer, the consumer benefits by at most:

$$(\mathbb{E}\left[q_t | q_t \geq \phi(i)\right] - \phi(i))\, \mathbb{P}[q_t \geq \phi(i)]$$

which would occur if a truthful reviewer could provide the most informative review of $r_t = q_t$ precisely. This is maximized when $\phi(i) = 0$ (i.e., consumers "on-the-fence" about purchasing benefit the most from reading the review). Thus, an upper bound for reading a review from a truthful reviewer is $\mathbb{E}[q_t | q_t \geq 0]\mathbb{P}[q_t \geq 0] = 1/\sqrt{2\pi}$. The benefit of reading a review then is upper bounded by $\pi_t/\sqrt{2\pi}$, which is less than the peeking cost ε whenever $\pi_t < \varepsilon\sqrt{2\pi}$. Hence, setting $\pi^* = \varepsilon\sqrt{2\pi}$, no consumer reads the reviews at time t when $\pi_t < \pi^*$, and moreover, consumers do not update π_t thereafter, so they never read reviews from that point on. In turn, the reviewer is indifferent between posting all reviews because none of them influence consumer purchasing decisions, so a (strategic) reviewer cannot benefit from a review that is not orthogonal to her signal s_t (for instance, by always writing $r_t = 0$). After $\pi_t < \pi^*$, consumers no longer update their beliefs about whether the reviewer is a truthful type, so this babble equilibrium persists forever. Thus, this is indeed a full equilibrium.

Second, it is clear that the reviewer obtains zero payoff for all periods after t where $\pi_t < \pi^*$: her influence is now zero by definition, and because no consumer changes their consumption based on her reviews, no firm would be willing to bribe her. Thus, in a full equilibrium, the reviewer must get zero payoff forever once $\pi_t < \pi^*$.

Finally, we show the babbling-trigger equilibrium must be the consumers' strategy profile in every risk-dominant equilibrium of the consumer-reviewer game, which establishes that any robust equilibrium has the property claimed in the lemma. Consider the sequence of full equilibrium profiles of the consumer-reviewer game $\{R_t[\delta], \mathcal{C}_{i,t}[\delta]\}$ with babbling-trigger at π^* (with existence established in the previous paragraphs) and any another sequence of full equilibrium profiles of the consumer-reviewer game $\{R'_t[\delta], \mathcal{C}'_{i,t}[\delta]\}$. Recall, by definition of risk-dominant equilibrium, it suffices to show that the (average-discounted) payoff the reviewer obtains under $\{R_t[\delta], \mathcal{C}'_{i,t}[\delta]\}$ is higher than the payoff she obtains under $\{R'_t[\delta], \mathcal{C}_{i,t}[\delta]\}$.

Consider reviewer strategy $R'_t[\delta]$. For every $\epsilon > 0$ we have either (i) $\liminf_{t\to\infty} \liminf_{\delta\to 1} \pi_t < \pi^*$ with probability (at least) $1 - \epsilon$, or (ii) $\liminf_{t\to\infty} \liminf_{\delta\to 1} \pi_t \geq \pi^*$ with probability (at least) ϵ in the full equilibrium sequence. Note that $\beta I_t + 1 \leq \beta + 1 < \infty$ is the maximum stage-game payoff the reviewer can obtain because no firm ever bribes more than \$1 (this is the maximum profit of the firm) and $I_t \leq 1$ because the maximum influence attainable for the reviewer is if she manages to convince the entire unit mass of consumers to switch their consumption decision as a result of her review. Under case (i) we have an upper bound on the average-discounted payoff of the reviewer: $(1+\beta)(1-\delta^{t+1}) + (1+\beta)\epsilon\delta^{t+1}$. For fixed t, as we take $\delta \to 1$, this payoff converges to $(1+\beta)\epsilon$; because ϵ can be made arbitrarily close to 0, the payoff of a patient reviewer is arbitrarily close to 0 under this profile. Under case (ii), we know by Lemma B1 that $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t[\delta]-R^{\hat{\omega}}||_\infty=0$ a.s. for some $\hat{\omega}\in\{L,H\}$, so the payoff to the strategic reviewer is the same (as $\delta\to 1$) as that of a truthful reviewer who plays $R^{\hat{\omega}}$. To see this, recall that by Lemma B1, the reviewer's strategy R_t must be indistinguishable from $R^{\hat{\omega}}$ (in the ∞-norm sense). Thus, at $t\to\infty$, (i) she receives the same influence from the consumers as if she is the truthful type $R^{\hat{\omega}}$ and (ii) because (myopic) firms see the distribution of reviews and realized qualities identical to the strategy $R^{\hat{\omega}}$, their bribes must be based entirely on their prior distributions about the type of the reviewer (who may or may not actually be strategic).

Now let us consider the profile $\{R_t[\delta],\mathcal{C}'_{i,t}[\delta]\}$. We know that because $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t[\delta]-R^{\hat{\omega}}||_\infty=0$ for some $\hat{\omega}\in\{L,H\}$, the strategic reviewer's payoff is the same under case (ii) from before. Under case (i), however, it is strictly more because $\liminf_{t\to\infty}\liminf_{\delta\to 1}\pi_t\geq\pi^*$ under the $R_t[\delta]$ full equilibrium (as $\delta\to 1$) sequence with high probability.[19] Because the peeking cost ε is positive, some proportion of consumers must obtain value from reading the review because it (in expectation) improves (and thus influences) their consumption decisions. Thus, I_t is lower bounded by a positive constant for all t, which implies the reviewer has a positive (average-discounted) payoff. Moreover, under, case (i), the reviewer receives a better payoff under $\{R_t[\delta],\mathcal{C}'_{i,t}[\delta]\}$ than under $\{R'_t[\delta],\mathcal{C}_{i,t}[\delta]\}$ (again as $\delta\to 1$), and the same payoff under case (ii). Hence, the total payoff is better under $\{R_t[\delta],\mathcal{C}'_{i,t}[\delta]\}$ than $\{R'_t[\delta],\mathcal{C}_{i,t}[\delta]\}$, establishing this is property holds in every risk-dominant equilibrium. $\square$

Lemma B3. *There exists $\hat{\omega}\in\{L,H\}$ such that $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t[\delta]-R^{\hat{\omega}}||_\infty=0$ in every robust full equilibrium.*

Intuition. A strategic reviewer always has the option to copy a truthful reviewer of the same precision as herself; this means the consumer will heed her recommendations exactly the same as if she was actually that truthful type. This provides a lower bound on the payoff attainable (through her influence) by a strategic reviewer. By Lemma B2, if a patient reviewer does not play close to a truthful strategy, her (average-discounted) payoff will be close to 0, and so she does not get this lower bound utility. Thus, she must mimic in equilibrium (in other words, in terms of Lemma B1, we know precisely which term must converge to 0).

Proof. By Lemma B1, for all $\epsilon>0$ and $\delta<1$, there exists T^* such that $\mathbb{P}[\pi_t\cdot||R_t(\delta)-R^{\hat{\omega}}||_\infty\geq\epsilon]\leq\epsilon$ for all $t>T^*$, for some $\hat{\omega}\in\{H,L\}$. Thus, for any $\epsilon>0$, we know that $\mathbb{P}[\limsup_{t\to\infty}\limsup_{\delta\to 1}\pi_t\cdot||R_t(\delta)-R^{\hat{\omega}}||_\infty\geq\epsilon]=0$. This implies that for any $\epsilon>0$, either (i) $\mathbb{P}[\limsup_{t\to\infty}\limsup_{\delta\to 1}\pi_t\geq\epsilon]=0$ or (ii) $\mathbb{P}[\limsup_{t\to\infty}\limsup_{\delta\to 1}||R_t(\delta)-R^{\hat{\omega}}||_\infty\geq\epsilon]=0$. To prove the lemma, it is sufficient to show that (i) cannot hold for any $\epsilon>0$.

Recall that ω is the reviewer's true type (not effective type). Note that by mimicking a truthful type exactly, the reviewer eventually obtains an influence of $I_\infty(\omega)$ almost surely (see Proposition 1). By Fudenberg and Levine (1992), in every consumer-reviewer equilibrium where $\delta\to 1$, the reviewer receives at least $\beta I_\infty(\omega)\geq\beta I_\infty(L)>0$, since $\beta>0$ (i.e., her Stackelberg payoff). By Lemma B2, if (i) were to hold, then for any $\epsilon>0$, there exists T^{**} such that $\mathbb{P}[\pi_t<\pi^*]\geq 1-\epsilon$. In this case, the reviewer receives a payoff of at most $(1+\beta)(1-\delta^{T^{**}+1})+(1+\beta)\epsilon\delta^{T^{**}}=(1+\beta)-(1+\beta)(1-\epsilon)\delta^{T^{**}}$. For sufficiently small $\zeta>0$, we can choose T^{**} sufficiently large such that $\epsilon<1-\sqrt{1-\frac{\beta I_\infty(\omega)-\zeta}{1+\beta}}$. For this value of T^{**}, let us pick δ^* such that $\left(1-\frac{\beta I_\infty(\omega)-\zeta}{1+\beta}\right)^{1/((1+\beta)T^{**})}<\delta^*<1$. Then $(1+\beta)-(1+\beta)(1-\epsilon)\delta^{T^{**}}<\beta I_\infty(\omega)-\zeta$ for all

[19] This is a consequence of the peeking cost vanishing (after sending $t\to\infty$ and then $\delta\to 1$), so the "mimicking" strategy always guarantees that π_t never falls below the trigger threshold.

$\delta > \delta^*$. This contradicts the fact that the reviewer receives at least her Stackelberg payoff in every equilibrium, so in fact (i) cannot hold, and (ii) must. $\square$

Lemma B4. *For a truthful reviewer with precision $\hat{\sigma}$, every consumer has belief distribution about the quality of the product (conditional on the review) given by $q_t | r_t \sim \mathcal{N}\left(\frac{r_t}{1+\hat{\sigma}^2}, \frac{\hat{\sigma}^2}{1+\hat{\sigma}^2}\right)$.*

Proof. The proof involves a straightforward application of Bayes' rule and is shown in the proof of Theorem 1 in Appendix A.

Lemma B5. *The (strategic) reviewer's "average" strategy, $\tilde{R}_t$,*[20] *can only depend on the public history in a trivial way, i.e., for any two public histories H_t, H_t', $\lim_{t\to\infty} \lim_{\delta\to 1} ||\tilde{R}_t(b_t, s_t, H_t) - R_t(b_t, s_t, H_t')||_\infty = 0$. Moreover, in every maximal robust equilibrium, $R_t(0, s_t, H_t)$ and $\int_{b_t\in F_t(H_t)} R_t(b_t, s_t, H_t)\, db_t$ can also only depend on the public history in a trivial way.*

Intuition. Lemma B5 makes two claims. First is that the reviewer cannot differentiate her average strategy solely based on the public history H_t. One example of such a strategy is that the reviewer favors firms arriving in even periods and disadvantages firms arriving in odd periods. Because consumers can observe the public history they will be able to detect this differentiation over time.

The second claim of Lemma B5 is that in a maximal equilibrium, the reviewer cannot differentiate how she handles bribing and non-bribing firms based on public history. To see this, note that if bribes depend on H_t, then there exist two (infinite) sequences of timelines $\{H_\tau', H_\tau''\}$ where the bribes differ for bribing firms and for honest firms depending on which timeline is realized, and where bribing firms get less in expectation under H_τ' than under H_τ'' for every τ. But this does not admit the maximal bribe: the reviewer should implement her strategy for the H_τ'' timeline sequence after all of the histories in H_τ' to solicit a larger bribe.

Proof. For a bribing firm t, we let $F_t(H_t)$ be its mixed strategy over bribes, as a function of the public history H_t. Note the "average" (over b_t) strategy employed by the reviewer is in fact:

$$\tilde{R}_t(s_t, H_t) = \theta R_t(0, s_t, H_t) + (1-\theta) \int_{b_t\in F_t(H_t)} R_t(b_t, s_t, H_t)\, db_t$$

for some $\hat{\omega}$ per Lemma B3 (i.e., $\lim_{t\to\infty} \lim_{\delta\to 1} ||R_t - R^{\hat{\omega}}||_\infty = 0$ for some $\hat{\omega}$). Note because the consumer can also observe H_t, it must be true that for all H_t, $\lim_{t\to\infty} \lim_{\delta\to 1} ||\tilde{R}_t(s_t, H_t) - R^{\hat{\omega}}||_\infty = 0$. But, recall that $R^{\hat{\omega}}$ depends only on s_t and not b_t or H_t. Therefore, by the triangle inequality, $\tilde{R}_t$ cannot depend on H_t in the sense that $\lim_{t\to\infty} \lim_{\delta\to 1} ||\tilde{R}_t(s_t, H_t) - \tilde{R}_t(s_t, H_t')||_\infty = 0$.

Next, consider any equilibrium where $\limsup_{t\to\infty} \limsup_{\delta\to 1} ||R_t(0, s_t, H_t) - R_t(0, s_t, H_t')||_\infty > \zeta$ for some $\zeta > 0$ (the alternate case is that $\lim_{t\to\infty} \lim_{\delta\to 1} ||R_t(0, s_t, H_t) - R_t(0, s_t, H_t')||_\infty = 0$). Then there exists an infinite sequence of t, $\{\tau_1, \tau_2, \ldots, \tau_k, \ldots\}$, such that there exist two histories H_{τ_k}, H_{τ_k}' at every time τ_k where:

$$\mathbb{E}_{r_{\tau_k} \sim R_{\tau_k}(0, s_{\tau_k}, H_{\tau_k})} \left[\phi^{-1}\left(\frac{r_{\tau_k}}{1+\hat{\sigma}^2}\right)\right] < \mathbb{E}_{r_{\tau_k} \sim R_{\tau_k}(0, s_{\tau_k}, H_{\tau_k}')} \left[\phi^{-1}\left(\frac{r_{\tau_k}}{1+\hat{\sigma}^2}\right)\right] \tag{3}$$

and the expression $\phi^{-1}\left(\frac{r_{\tau_k}}{1+\hat{\sigma}^2}\right)$ denotes the consumption after review r_{τ_k} for large t, as a consequence of Lemma B3, Lemma B4, and Lebesgue's dominated convergence theorem (ϕ^{-1}

[20] The average review incorporates both the honest and strategic firms: $\tilde{R}_t(s_t, H_t) = \theta R_t(0, s_t, H_t) + (1 - \theta) \int_{b_t\in F_t(H_t)} R_t(b_t, s_t, H_t)\, db_t$

is bounded). Instead, suppose the reviewer picks a different strategy $R'_{\tau_k}(0, s_{\tau_k}, H_{\tau_k})$ which takes $R'_{\tau_k}(0, s_{\tau_k}, H'_{\tau_k}) \leftarrow R_{\tau_k}(0, s_{\tau_k}, H_{\tau_k})$ and $R'_{\tau_k}(b_{\tau_k}, s_{\tau_k}, H'_{\tau_k}) \leftarrow R_{\tau_k}(b_{\tau_k}, s_{\tau_k}, H_{\tau_k})$ at all times τ_k. It is clear that this can still be supported in equilibrium because $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t - R'_t||_\infty = 0$ by the observation in the previous paragraph. Moreover, to offset the inequality in (3) it must be the case that for bribing firms:

$$\mathbb{E}_{b_{\tau_k}\sim F_{\tau_k}(H_{\tau_k})}\left[\mathbb{E}_{r_{\tau_k}\sim R_{\tau_k}(b_{\tau_k}, s_{\tau_k}, H_{\tau_k})}\left[\phi^{-1}\left(\frac{r_{\tau_k}}{1+\hat{\sigma}^2}\right)\right]\right] > \mathbb{E}_{b_{\tau_k}\sim F_{\tau_k}(H_{\tau_k})}\left[\mathbb{E}_{r_{\tau_k}\sim R_{\tau_k}(b_{\tau_k}, s_{\tau_k}, H'_{\tau_k})}\left[\phi^{-1}\left(\frac{r_{\tau_k}}{1+\hat{\sigma}^2}\right)\right]\right]$$

which implies that there exists an equilibrium where $b'_{\tau_k} > b_{\tau_k}$ for all such τ_k (i.e., by playing R'_t), so the equilibrium of R_t is not maximal. This implies that $R_t(0, s_t, H_t)$ cannot depend on H_t as $t \to \infty$, which immediately shows that $\int_{b_t\in F(H_t)} R_t(b_t, s_t, H_t)\, db_t$ may not depend on H_t either. □

Lemma B6. *Consider the strategy $R^{eq}_t(b_t, s_t, H_t)$ where the reviewer writes reviews according to the strategy given in Section 3 (i.e., map m of Lemma A1).*[21] *Then $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t(b_t, s_t, H_t) - R^{eq}_t(b_t, s_t, H_t)||_\infty = 0$ in every robust equilibrium.*

Intuition. Lemma B3 establishes that the reviewer must play as one of the two effective types, and by Condition (ii) of Appendix B.2.2, firms break ties in a consistent manner. Therefore, the problem that strategic firms face upon entry (whether and how much to bribe) are identical, and the maximal bribe b^* they are willing to give is time-invariant. The reviewer's strategy space consists of all the strategies that obey the distributional constraints of making her appear unbiased (and reporting truthfully). Over that space, and from the bribing firm's perspective, the reviews generated from the strategy in the statement of the lemma first-order stochastically dominate the reviews generated from all other strategies, and is therefore the strategy for which firms are willing to pay the maximal bribe.

Proof. In a robust equilibrium, every firm at time t chooses the bribe that maximizes its expected utility (given the distribution of reviews from H_t), but breaks indifferences in a deterministic (and consistent) way via condition (ii) (see Appendix B.2.2). Therefore, the firm always plays a pure strategy, and $\int_{b_t\in F_t(H_t)} R_t(b_t, s_t, H_t)\, db_t$ immediately reduces to $R_t(F_t(H_t), s_t, H_t)$. By Lemma B5, $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t(F_t(H_t), s_t, H_t) - R_t(F_t(H_t), s_t)||_\infty = 0$, for some $R_t(F_t(H_t), s_t)$ (which does not depend directly on H_t).

Recall that by Lemma B3 the reviewer mimics a truthful strategy eventually almost surely. Because firms observe the distribution of the public history H_t, firms' beliefs about the actual type and the effective type of the reviewer converge over time, the former to the prior p and the latter to a point mass on either R^H_t or R^L_t (the strategy chosen by the reviewer). Since the payoff of the firm at t is a function only of b_t and X^*_t, which is entirely determined by R_t, all strategic firms are identical, and the outcome of review manipulation is the same for all strategic firms. Thus, observe in a maximal equilibrium, b_t is given by the supremum over all bribes supported

[21] Recall: The reviewer samples the product and receives an unbiased, noisy signal s_t of the quality. If the reviewer's effective type is H, she reports truthfully, $r_t = s_t$. If the reviewer's effective type is L, then:

- if her actual type is L, she collects the bribe if one is offered, but still reports truthfully $(r_t = s_t)$, or;
- if her actual type is H, she differentiates her review based on whether she was offered a bribe:
  - (i) If $b_t = 0$, the reviewer writes the review $r_t = s_t + \varepsilon'_t$, where $\varepsilon'_t \sim \mathcal{N}(0, \sigma^2_H - \sigma^2_L, 0, \theta)$.
  - (ii) If $b_t = b^*$, the reviewer writes the review $r_t = s_t + \varepsilon'_t$, where $\varepsilon'_t \sim \mathcal{N}(0, \sigma^2_H - \sigma^2_L, \theta, 1)$.

by some equilibrium $R_t(b_t(H_t), s_t)$ as $t \to \infty$, which we can denote by b^* (note such a maximal bribe cannot depend on t, because the reviewer's strategy does not depend on H_t, including t itself, as shown in Lemma B5). The unique equilibrium outcome for the firms (in a maximal equilibrium) is given by strategic firms bribing b^* and honest firms bribing 0. Hence, we can denote the strategy $R_t(0, s_t)$ for honest firms and $R_t(b^*, s_t)$ for strategic firms in the limiting reviewer strategy, with the "average" strategy denoted by $\tilde{R}_t(s_t) = \theta R_t(0, s_t) + (1-\theta) R_t(b^*, s_t)$. If $\lim_{t\to\infty} \lim_{\delta\to 1} ||\tilde{R}_t(s_t) - R^{\hat{\omega}}||_\infty = 0$ for some $\hat{\omega}$ (as shown in Lemma B3), then $\tilde{R}_t(s_t)$ must converge in distribution to a normal distribution $\mathcal{N}(0, 1+\hat{\sigma}^2)$ with $\text{cov}(\tilde{R}_t(s_t), q_t) = 1$. Since this distribution is entirely determined by its covariance matrix, $\tilde{R}_t(s_t)$ converges (in distribution) to a unique distribution.

Therefore, we can write the *limiting* strategy (i.e., the one $\tilde{R}_t$ converges to in the ∞-norm) as simply $\tilde{R}$. Based on the previous paragraph, we know that $\tilde{R}(s_t)$ is constrained to have the distributional constraints shown in Lemma A1. Because $m(s_t, b^*) - s_t$ first-order stochastically dominates $\tilde{R}(s_t, b^*) - s_t$ by Lemma A1, the expected review under $\tilde{R}^*(s_t, b^*)$ is always at least as high as under any other reviewer strategy $\tilde{R}(s_t, b^*)$ satisfying the constraints, which implies expected consumption is maximized (relative to $\tilde{R}(s_t, 0)$), and so generates the maximal bribe. □

Lemma B7 (Entropy Lemma)**.** *If an agent's type has precision σ, then it is impossible to mimic a precision $\hat{\sigma} < \sigma$, i.e., $\limsup_{t\to\infty} \limsup_{\delta\to 1} ||R_t[\delta] - R^{\hat{\omega}}||_\infty > 0$ if $\hat{\omega}$ is more precise that the reviewer's true type (regardless of whether the reviewer is strategic or honest).*

Intuition. A reviewer cannot improve her signal-to-noise ratio by applying a clever mapping from her signal to what she reports. A type H reviewer can inject additional noise into her review (i.e., add entropy) but a type L reviewer cannot remove noise (remove entropy) from her observed signal.

Proof. Consider an agent who mimics some precision $\hat{\sigma} < \sigma_\omega$. Let $f_t : s_t \mapsto r_t$ be any stochastic function (possibly time-varying) mapping the reviewer's signals into reviews (which can depend on bribes, the public history, etc). Since s_t is a sufficient statistic for q_t (given r_t), by the Fisher-Neyman factorization theorem we can represent $f_t(s_t) = g_t(s_t, q_t) h_t(r_t)$ for some functions g_t, h_t. By the (information theory) chain rule, the Fisher information from s_t, $\mathcal{I}_{s_t}(q_t)$, and the Fisher information from r_t, $\mathcal{I}_{r_t}(q_t)$, satisfy $\mathcal{I}_{s_t} \geq \mathcal{I}_{r_t}$. Similarly, we know that $\mathcal{I}_{s_t \sim \mathcal{N}(0,\sigma_H^2)} > \mathcal{I}_{s_t \sim \mathcal{N}(0,\sigma_L^2)}$. Thus, it is impossible that $\limsup_{t\to\infty} \limsup_{\delta\to 1} ||R_t - R^{\hat{\omega}}||_\infty = 0$, whenever $\hat{\omega} = H$ but $\omega = L$. □

Intuition for Theorem B1. The consumer-reviewer game is a well-studied game between a long-run player and many myopic players, where standard existence results can be applied. The more general game always has an equilibrium where firms do not bribe and reviewers always report truthfully, so existence is established via construction.

Proof of Theorem B1. For the first part, notice the consumer-reviewer game is a dynamic game of incomplete information of the form given in Fudenberg and Levine (1992), where a long-run player plays against a sequence of myopic players who observe the public history H_t (in our model these agents actually live through that history). These players may observe the entire history of play (except for bribes) since the public signal q_t eliminates the private information received by the reviewer about q_t. Because the bribing strategy is given exogenously by $\tilde{F}_t(H_t)$, which is a function only of this public history, the payoffs are determined by nature given H_t.

For the second part, we establish that a trivial equilibrium (i.e., no bribing equilibrium) is always a full equilibrium. We thus prove existence by construction. Suppose a firm bribes, but receives a truthful review from either a truthful or strategic reviewer, who both always write truthful reviews in equilibrium. Then the firm decreases her payoff from bribing, yet, receives no benefit from the review because the recommendation written is exactly the same as it would have been without the bribe. Thus, the firm has no incentive to deviate. On the other hand, the reviewer in this equilibrium obtains the maximal influence she can by reporting truthfully (this is a consequence of Lemma B7). Because she receives no bribes regardless of H_t, her entire payoff is determined by her influence, which is maximized when she plays R_t according to $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t[\delta]-R^{\omega}||_{\infty}=0$. Thus, the reviewer has no incentive to deviate from reporting truthfully. $\square$

Intuition for Theorem B2. Lemma B3 establishes that the reviewer must mimic a truthful reviewer. Moreover, Lemma B7 shows that a reviewer can only mimic down to a worse quality reviewer. By Lemma B4, consumers learn this effective type of the reviewer and use it to inform their consumption decision in a precise way. By Lemma B5, the reviewer's strategy must eventually be time-invariant, and by Lemma B6, there is a unique maximal bribe given by the distribution presented in Section 3.1 that still satisfies the constraints for Lemma B3 to hold.

Proof of Theorem B2. Note that by Lemma B3, in every robust equilibrium, the reviewer mimics a type $\hat{\omega}$. Thus, by the strong law of large numbers, we know that $\mathbb{P}[q_t|r_t,H_t]$ converges almost surely to the expression given in Lemma B4. By Lemma B4, each consumer's expected quality is given by $\mathbb{E}[q_t|r_t]=\frac{r_t}{1+\hat{\sigma}^2}$, so consumers purchase iff $\phi(i)\leq\frac{r_t}{1+\hat{\sigma}^2}$. Simultaneously, firms observe the "shuffled" version of H_t (see Appendix B.1), and by the law of large numbers converge to know the effective type $\hat{\sigma}$ (with probability 1) of the reviewer; moreover, the firm knows that if $\hat{\sigma}=L$ the reviewer is with probability p actual type H, whereas if $\hat{\sigma}=H$ the reviewer is with probability 1 actual type H (by Lemma B7).

Finally, we need only show that there is a unique $\hat{\omega}$ the high-type reviewer mimics in the maximal equilibrium, under generic conditions (by Lemma B7, the type-L reviewer must mimic type L). Recall that by Lemma B3, for the type H reviewer, there are two options: (i) mimic type H (i.e., $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t[\delta]-R^H||_{\infty}=0$) or (ii) mimic type L (i.e., $\lim_{t\to\infty}\lim_{\delta\to 1}||R_t[\delta]-R^H||_{\infty}=0$). There are two possibilities in the former case: (a) the strategic reviewer plays according to H (in the ∞-norm sense) in equilibrium, in which case firms do not bribe eventually almost surely, or (b) the reviewer is truthful and type H, in which case firms eventually learn the reviewer is truthful almost surely and again stop bribing. Therefore, under (i), the reviewer who plays as true type H cannot solicit bribes as $t\to\infty$. For Case (ii), there is a maximal bribe b^* attainable via Lemma B6. As $\delta\to 1$, the payoff of the reviewer is $\beta I_{\infty}(H)$ if she plays as H and $\beta I_{\infty}(L)+b^*$,[22] if she plays as L (as a consequence of Kolmogorov's strong law, e.g., see Sen and Singer (1994) Theorem 2.3.10, and Proposition 1). Note that the payoff of the firm is independent of β, so does not affect the maximally supported bribe b^* (Lemma B6). Thus, the type-H reviewer is indifferent between these two if and only if $\beta I_{\infty}(H)=\beta I_{\infty}(L)+b^*$. However, for any sufficiently small $\epsilon>0$, setting $\beta\leftarrow\beta\pm\epsilon$ breaks the indifference. Thus, for generic β, there is a unique choice for $\hat{\omega}$ for the high-type reviewer. $\square$

[22] Both $I_{\infty}(H)$ and $I_{\infty}(L)$ are both defined in the proof of Proposition 1 in Appendix A.

B.5 Worked Examples

We present five worked examples that highlight distinct points made throughout the text. In all of these examples (with the exception of Appendix B.5.4), we look for when the *pooling equilibrium* exists (we know the *separating equilibrium* always does). The technique is to compute the largest bribe the firm is willing to offer and the lowest bribe the reviewer is willing to accept. This determines whether an environment is bribe-proof or not. In Appendix B.5.4, we establish a statistical feature of our fully dynamic model that firms cannot bribe multiple reviewers simultaneously.

In Appendix B.5.1, we explore the connection between the "friends" vs "non-friends" model of Section 2 and how a similar example exists within the model of Section 3 with bribes instead. In Appendix B.5.2, we work an example that demonstrates Theorem 1 about the non-monotonicity of bribes in θ. In Appendix B.5.3, we elaborate on the discussion in Section 4.2 by providing a formal example showing that bribing firms themselves might prefer a truthful world. Finally, in Appendix B.5.5, we show that while additional competition in bribe-proof environments always improves consumer welfare (more independent signals of quality), it may not when bribes are possible (i.e., bribes might become more lucrative and corrupt information transmission with more competition).

B.5.1 Worked Example: Connection between Section 2 (Reduced-Form Example) and Section 3 (Model)

Suppose there are two reviewer types given by $\sigma_H = 1$ and $\sigma_L = 2$ as in the example in Section 2, and that these types are equally likely. If the reviewer is high-skill, she can either commit to biasing her reviews or not. If the reviewer plays effective precision $\hat{\sigma}$, then her influence is (eventually) given by $I_\infty(\hat{\sigma})$.

Also as in Section 2, assume that 50% of firms are strategic ($\theta = 1/2$) and their outside option γ is zero. Consumers are distributed according to the following piecewise function:

$$\phi(i) = \begin{cases} -\infty, \text{ if } i \in [0, 1/3) \\ 1/2, \text{if } i \in [1/3, 2/3) \\ \infty, \text{ if } i \in (2/3, 1] \end{cases}$$

which means that 1/3 of the consumers never buy, 1/3 of the consumers buy only if $\mathbb{E}[q_t|r_t] \geq 1/2$, and 1/3 of the consumers always buy. We can compute the influence directly by noting the cutoff for the middle consumer group is $r_t \geq (1 + \hat{\sigma}^2)/2$. This yields $I_\infty(1) \approx 0.0266$ and $I_\infty(2) \approx 0.0146$. For the high-type reviewer to mimic low-type, her average bribe payment must exceed $\beta(I_\infty(1) - I_\infty(2)) \approx 0.012\beta$.

Consider $q_t, \varepsilon_t \sim \mathcal{N}(0, 1)$ and $\varepsilon'_t \sim \mathcal{N}(0, 3)$, all independent. If a firm bribes, then its review is distributed as $r^B_t = q_t + \varepsilon_t + |\varepsilon'_t|$; if it does not bribe, it is distributed as $r^{NB}_t = q_t + \varepsilon_t - |\varepsilon'_t|$. To compute the benefit of a bribe, we simply calculate the difference in probabilities (from bribing) that their review will exceed $(1 + \hat{\sigma}^2)/2 = 5/2$ when $\hat{\sigma} = L$. This benefit is equal to 0.08.

Thus, the firm gains (on average) 8% more of the consumer base by bribing. Since reviewer types are equally-likely, the true benefit from a bribe is only 4% (because like consumers, firms also cannot tell whether the reviewer is type H and biasing reviews, or is actually type L and reporting honestly). Similarly, the "on-average" bribe to a reviewer will be half of this (because $1 - \theta = 1/2$). Therefore, if $0.02 > 0.012\beta$ (i.e., $\beta < 5/3$) the strategic firm is willing to bribe any amount less than 0.04 and the reviewer is willing to accept anything greater than 0.024β

from strategic firms (which has a non-empty intersection). On the other hand, if $\beta > 5/3$, the environment is bribe-proof.

Finally, we look at consumer utility. A review $r_t \geq (1+\hat{\sigma}^2)/2$ has an expected quality q_t of:

$$\mathbb{E}[q_t | r_t \geq (1+\hat{\sigma}^2)/2] = \frac{1}{2\pi\hat{\sigma}\Phi^{-1}\left(-\frac{(1+\hat{\sigma}^2)}{2\hat{\sigma}}\right)} \int_{-\infty}^{\infty} \int_{(1+\hat{\sigma}^2)/2}^{\infty} q \exp(-q^2/2) \exp\left(-\frac{(r-q)^2}{2\hat{\sigma}^2}\right) \, dr \, dq$$
$$= \frac{1}{\Phi^{-1}\left(-\frac{(1+\hat{\sigma}^2)}{2\hat{\sigma}}\right)} \cdot \frac{1}{\exp\left(\frac{\hat{\sigma}^2+1}{8}\right)\sqrt{2(\hat{\sigma}^2+1)\pi}}$$

When $\hat{\sigma} = 1$, the above is equal to $(2e^{1/4}\sqrt{\pi}\Phi^{-1}(-1))^{-1} \approx 1.34$; when $\hat{\sigma} = 2$, it is equivalent to $(e^{5/8}\sqrt{10\pi}\Phi^{-1}(-5/4))^{-1} \approx 0.903$. As consumer utility is monotonically increasing in $\mathbb{E}[q_t | r_t \geq (1+\hat{\sigma}^2)/2]$ in this example, we see that when the reviewer mimics low-skill, consumer utility decreases.

B.5.2 Worked Example: Intermediate θ Admits Bribes

Let us revisit the setting of Appendix B.5.1. We first show that bribes are possible when θ lies in the interior of $[0,1]$. This is because the reviewer can offer favorable biases to a subset of firms at a price that both the reviewer and firms find agreeable. We then show, as suggested by Theorem 2, that as θ gets closer to 0 or 1, this type of biasing is either not profitable or not sustainable.

Unlike Appendix B.5.1, we will assume firms have an outside option of $\gamma = 0.37$. One can compute the fair consumption: $\bar{X}^* = \frac{1}{3} + \frac{1}{3} \cdot \mathbb{P}[r_L \geq 5/2] = \frac{1}{3} + \frac{1}{3} \cdot \Phi(-\sqrt{5}/2) \approx 0.378$, where r_L is the review given by a low-skill (honest) reviewer and Φ is the CDF of the standard normal. Therefore, the fair consumption net the outside option for the firm is approximately 0.008. This implies the bribing propensity is equal to $\psi = 0.008/(I_\infty(\sigma_H) - I_\infty(\sigma_L))$.

If $\theta = 1/2$, then when the high-skill reviewer mimics low-skill, she can accept bribes from the 50% of the firms who are willing to bribe, but loses $I_\infty(\sigma_H) - I_\infty(\sigma_L)$ in influence; simultaneously, the firm gains $\frac{1}{3} + \frac{1}{3}\mathbb{P}[r_t^B \geq 5/2] - \max\{\gamma, \frac{1}{3} + \frac{1}{3}\mathbb{P}[r_t^{NB} \geq 5/2]\}$ (before bribes). Notice that:

$$\mathbb{P}[r_t^{NB} \geq 5/2] = \mathbb{P}[(q_t + \varepsilon_t) - |\varepsilon_t'| \geq 5/2]$$
$$= \frac{1}{\sqrt{6}\pi} \int_{5/2}^{\infty} \int_0^{x-5/2} \exp\left(-\frac{x^2}{4}\right) \exp\left(-\frac{y^2}{6}\right) \, dy \, dx$$
$$= 0.0093$$

Since $\gamma > \frac{1}{3} + \frac{1}{3}\mathbb{P}[r_t^{NB} \geq 5/2]$, the firm would prefer to stay out instead of enter and not bribe. The benefit the firm gets from bribing is thus $\frac{1}{3} + \frac{1}{3}\mathbb{P}[r_t^B \geq 5/2] - \gamma$, with:

$$\mathbb{P}[r_t^B \geq 5/2] = \mathbb{P}[(q_t + \varepsilon_t) + |\varepsilon_t'| \geq 5/2]$$
$$= \frac{1}{\sqrt{6}\pi} \int_{-\infty}^{\infty} \int_{5/2-x}^{\infty} \exp\left(-\frac{x^2}{4}\right) \exp\left(-\frac{y^2}{6}\right) \, dy \, dx$$
$$= 0.264$$

Therefore, the payoff to the firm from bribing, assuming the reviewer is high-skill, is $1/3 + 1/3 \cdot (0.264) - 0.37 = 0.0513$. Since both types are equally-likely, the true benefit is half, which means roughly the maximal bribe is $b^* = 0.026$. Therefore, the largest "average" bribe the reviewer can solicit is half of this (since only half the firms bribe), which is equal to 0.013.

Is this setting bribe-proof? Even when β is slightly larger than the bribing propensity ψ, the loss of influence for the reviewer is $\beta \cdot (I_\infty(\sigma_H) - I_\infty(\sigma_L)) = 0.008$, whereas the bribes she can receive are as large as 0.013. Therefore, there exists a bribing equilibrium when $\theta = 1/2$.

In the case of θ close to 1, note that the largest bribe any reviewer could solicit is $1/3$ by $(1-\theta)$ fraction of the firms. Thus, for $\beta = \psi$, as long as $\theta > 0.98$ (i.e., fewer than 2% of firms possibly bribe), these bribes are not sufficient to compensate for the damage to the reviewer's influence. On the other hand, when θ is close to 0, by entering and bribing the firm gets slightly more than fair consumption, $\bar{X}^*$, as opposed to not entering and receiving the outside option. Therefore, the firm would be willing to bribe slightly above $.008$, but this does not compensate the reviewer sufficiently for her loss in influence when $\beta > \psi$. Thus, both of these extreme environments are bribe-proof, as predicted by Theorem 2, even though the case of $\theta = 1/2$ is not.

B.5.3 Worked Example: Bribing Firms' Welfare

Revisit the setting of Appendix B.5.1. If the reviewer has an effective type H, then the "intermediate" consumers (those with $\phi(i) = 1/2$) purchase whenever $r_t \geq 1$; if the reviewer has an effective type L, these consumers purchase whenever $r_t \geq 5/2$. Thus, the firm's welfare in a world where bribes do not exist is given by:

$$\begin{aligned}\frac{1}{3} + \frac{1}{6}\left(\mathbb{P}_{r_t \sim \mathcal{N}(0,2)}[r_t \geq 1] + \mathbb{P}_{r_t \sim \mathcal{N}(0,5)}[r_t \geq 5/2]\right) &= \frac{1}{3} + \frac{1}{6}(\Phi(-1/\sqrt{2}) + \Phi(-\sqrt{5}/2)) \\ &\approx 0.3953\end{aligned}$$

When a firm bribes an amount b, with probability 1/2 the reviewer is an actual low type, and thus just reports $r_t = q_t$. Thus, the bribe gives utility $\frac{1}{3} + \frac{1}{3}\mathbb{P}_{r_t \sim \mathcal{N}(0,5)}[r_t \geq 5/2] - b \approx 0.3773 - b$. On the other hand, with probability 1/2 the reviewer is an actual high type who will mimic an effective low type; the reviewer will report $r_t^B = q_t + \varepsilon_t + |\varepsilon_t'|$ where $q_t, \varepsilon_t \sim \mathcal{N}(0,1)$ and $\varepsilon' \sim \mathcal{N}(0,3)$. Thus, we compute the probability of capturing the intermediate consumers when bribing a high-type reviewer:

$$\begin{aligned}\mathbb{P}[r_t^B \geq 5/2] &= \mathbb{P}[q_t + \varepsilon_t \equiv \kappa \geq 5/2] + \mathbb{P}[\kappa \leq 5/2 \ \cap \ |\varepsilon_t| \equiv \zeta \geq 5/2 - \kappa] \\ &= \Phi\left(-\frac{5}{2\sqrt{2}}\right) + 2 \cdot \frac{1}{2\sqrt{6}\pi}\int_{-\infty}^{5/2}\int_{5/2-\kappa}^{\infty} \exp\left(-\frac{\kappa^2}{4}\right)\exp\left(-\frac{\zeta^2}{6}\right) d\zeta \ d\kappa \\ &\approx 0.255\end{aligned}$$

(Note this is almost the exact same numeric value from Appendix B.5.1 when we calculated $\mathbb{P}[r_t^B \geq 5/2] - \mathbb{P}[r_t^{NB} \geq 5/2]$. This was because the chance of getting a review greater than 5/2 while being biased downward was only a couple percentage points.) Thus, when the reviewer is actually the high type, the expected welfare of the bribing firm is $\frac{1}{3} + \frac{1}{3} \cdot 0.255 - b = 0.4183 - b$. This implies that the total expected welfare is $\frac{1}{2} \cdot (0.4183 - b) + \frac{1}{2} \cdot (0.3773 - b) = 0.3978 - b$.

Recall from Appendix B.5.1 that the largest bribe for the firm is 0.02. In this equilibrium, the firm has total welfare $0.3778 < 0.3953$. So, perhaps surprisingly, even the firms who bribe in the maximal bribing equilibrium would prefer to not bribe (given all others do the same) and receive a fair review. But because other firms bribe, a non-bribing firm will be given an especially bad review, which forces the firm's hand into bribing. If collusion was possible, firms would prefer to agree upfront that there will be no bribing (if such an agreement could be made enforceable).

B.5.4 Worked Example: Firms Cannot Bribe Multiple Reviewers

Here, we present the main ideas that show with long-run interaction between firms, reviewers, and consumers, the firms are incapable of bribing multiple reviewers (and receiving multiple biased reviews in the same period). The key step in the proof is to recognize that if firms bribe two reviewers simultaneously, who then simultaneously bias the review, this *will be detectable* by the consumers. This is because the reviewers will tend to "agree" on the reviews which cannot be explained by the realized qualities alone. The final step is to recognize that if the reviewers aim for zero conditional (on q_t) correlation across each other, and because bribing firms require positive bias (to pay for the bribe), at least one of the reviewers cannot match the marginal distribution of reviews of a truthful reviewer. This reduces to a setting where firms can only bribe (and receive biased reviews) from a single reviewer, with the other reviewers providing fair reviews about the product quality.

For simplicity, suppose that we have two reviewers, R1 and R2, who are both type H with high probability and β is close to 0. As in Appendix B.5.1, we assume $\sigma_H = 1$, $\sigma_L = 2$, and half of the firms are capable of bribing ($\theta = 1/2$). Because reviewers do not care about their influence, it is clear that with just a single reviewer, there is a bribing equilibrium. Can the firm bribe both R1 and R2 and receive biased reviews from both?

This possibility is explored in Figure 4. As observed in the first two figures, both reviewers match the review distribution of a truthful type-L reviewer, despite actually being type H and biasing the reviews of bribing firms as per the strategy outlined in Section 3 and formally shown as the pooling equilibrium with the largest bribe b^* possible in Lemma A1. Thus, when either R1 or R2 is the only reviewer, consumers cannot detect that the reviewer is misrepresenting the quality signals she receives. However, the third panel in Figure 4 allows consumers to detect this and eventually stop listening when both reviewers are present. Reviewers will not write conditionally independent (on q_t) reviews because they will both bias bribing firms upward and honest firms downward. The simulated conditional correlation (on q_t) for two type-L truthful reviewers is given by the orange line; however, the *actual* simulated conditional correlation for two type-H reviewers mimicking type L is given by the blue line. Thus, eventually the consumers will be able to use this correlation to detect whether both of the reviewers are misrepresenting themselves (in a consistent way) or just both less accurate reviewers.

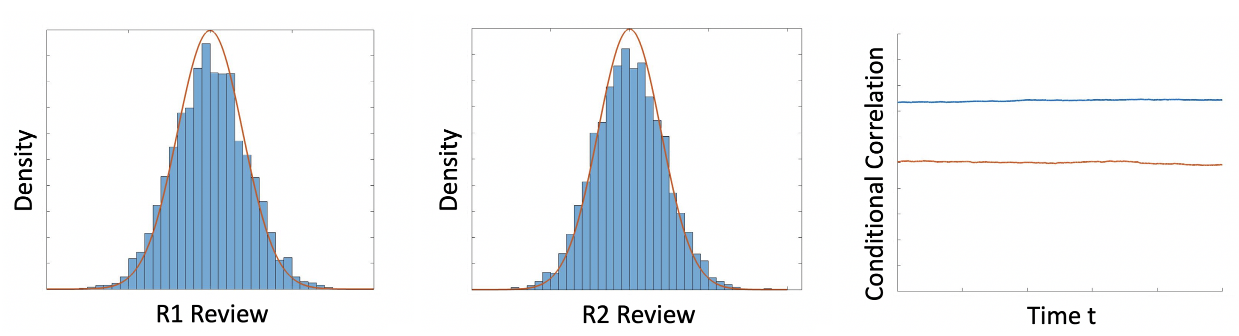

Figure 4. First panel: Histogram of R1's reviews who is type H but biases strategic firms, compared against theoretical distribution for type L truthful reviewer; Second panel: Similar plot for R2's reviews; Third panel: Conditional correlations (on q_t) between R1 and R2's reviews (top line) against theoretical correlation for truthful type-L reviewers (bottom line).

For R1 and R2 to go undetected (while manipulating reviews), they must match the zero conditional correlation between $r_{1,t}$ and $r_{2,t}$ over a long time horizon. Thus, if bribing firms want positive bias from *both* R1 and R2, which leads to positive (conditional) correlation between reviews, they must offset this with negative (conditional) correlation when truthful firms appear. In particular, one reviewer (without loss, say R2) must bias upward the truthful firms when the other biases them downward to create this persistent disagreement. The outcome of this reviewer strategy is pictured in the three panels of Figure 5. By biasing truthful firms upward, R2 creates persistent disagreement from R1 with truthful firms that counteracts the persistent agreement with R1 for bribing (strategic) firms. This can be seen in the third panel, where the observed and expected (conditional) correlations of the reviews from R1 and R2 are both close to zero. However, this strategy requires R2 to write reviews according to the distribution pictured in the second panel, as she upward biases both strategic and truthful firms. Eventually, it will become clear to consumers that R2 is consistently biased upward, and is thus manipulating her reviews. This again breaks this strategy as one that can be sustained in equilibrium for consumers.

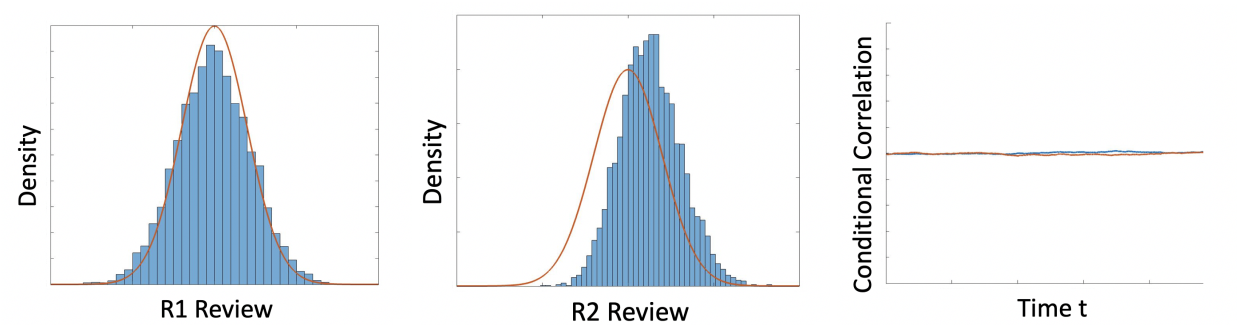

Figure 5. First panel: R1 successfully portrays herself as a truthful type L reviewer; Second panel: R2 is consistently upward biased in an attempt to make herself independent of R1's reviews; Third panel: Consumers cannot detect that R1 and R2 are both simultaneously accepting bribes in return for biased reviews.

The consequences of this are that bribing firms cannot simultaneously bribe several reviewers with the hopes of obtaining multiple biased reviews. Instead, eventually, each firm can only get one reviewer of the bunch to bias her review in each period in exchange for a bribe, and therefore only has an incentive to bribe one reviewer. Thus, the presence of competition squashes the ability of the firm to generalize its bribing scheme to all (or even more than one) reviewers.

B.5.5 Worked Example: Reviewer Competition Lowers Welfare

We assume first there that there is a single reviewer R1, who with high probability is type H (i.e., "super reviewer") with $\sigma_H = 1$, and with low probability is type L with $\sigma_L = 2$. We will assume the same distribution of $\phi(i)$ given in Appendix B.5.1 and continue to assume that half of the firms are honest ($\theta = 1/2$).

Monopolist. As we saw in Appendix B.5.1, the largest bribe the firm is willing to offer is 0.08. Thus, if $\beta = 4$, the loss in influence for the reviewer from playing effective type L is $4(0.0266 - 0.0146) = 0.048$, but since $\theta = 1/2$ (only half the firms bribe), the average bribe is 0.04, which does not compensate the reviewer for her loss in influence. **Thus, the environment with a monopolist reviewer R1 is bribe-proof.**

Duopoly. However, suppose we now add a reviewer R2 who is a type L (i.e., a "casual reviewer"). Note that when $\theta = 1$ (instead of $\theta = 1/2$), all reviewers report truthfully and the consumers get two independent signals with variances $\sigma^2_{\omega_1}$ and $\sigma^2_{\omega_2}$. Both these signals together provide the consumer with a tighter estimate of q_t than either signal by itself, so consumer welfare improves.

But, when $\theta = 1/2$, then, when R1 plays with precision $\hat{\sigma}_1$, the consumers with outside option $\phi(i) = 1/2$ purchase if and only if:[23]

$$\frac{r_{1,t}/\hat{\sigma}_1^2 + r_{2,t}/4}{1 + 1/\hat{\sigma}_1^2 + 1/4} = \frac{4r_{1,t} + \hat{\sigma}_1^2 r_{2,t}}{4\hat{\sigma}_1^2 + 4 + \hat{\sigma}_1^2} \geq \frac{1}{2}$$

This influence calculation is given by:

$$\begin{aligned}
I_\infty(\hat{\sigma}_1) =& \frac{1}{9}(\mathbb{P}[r_{2,t} \leq 5/2 \cap 4r_{1,t} + \hat{\sigma}_1^2 r_{2,t} \geq (4\hat{\sigma}^2 + 4 + \hat{\sigma}^2)/2] + \mathbb{P}[r_{2,t} \geq 5/2 \cap 4r_{1,t} + \hat{\sigma}^2 r_{2,t} \leq (4\hat{\sigma}^2 + 4 + \hat{\sigma}^2)/2]) \\
=& \frac{1}{9}(\mathbb{P}[r_{2,t} \leq 5/2 \cap r_{1,t} \geq (4\hat{\sigma}^2 + 4 + \hat{\sigma}^2 - 2\hat{\sigma}^2 r_{2,t})/8] + \mathbb{P}[r_{2,t} \geq 5/2 \cap r_{1,t} \leq (4\hat{\sigma}^2 + 4 + \hat{\sigma}^2 - 2\hat{\sigma}^2 r_{2,t})/8)]) \\
=& \frac{1}{9}\frac{1}{2(2\pi)^{3/2}\hat{\sigma}_1} \int_{-\infty}^{\infty} \int_{-\infty}^{5/2} \int_{(4\hat{\sigma}^2+4+\hat{\sigma}^2-2\hat{\sigma}^2 r_2)/8}^{\infty} \exp\left(-\frac{q^2}{2}\right) \exp\left(-\frac{(r_1-q)^2}{2\hat{\sigma}_1^2}\right) \exp\left(-\frac{(r_2-q)^2}{8}\right)\ dr_1\ dr_2\ dq \\
&+ \frac{1}{9}\frac{1}{2(2\pi)^{3/2}\hat{\sigma}_1} \int_{-\infty}^{\infty} \int_{5/2}^{\infty} \int_{-\infty}^{(4\hat{\sigma}^2+4+\hat{\sigma}^2-2\hat{\sigma}^2 r_2)/8} \exp\left(-\frac{q^2}{2}\right) \exp\left(-\frac{(r_1-q)^2}{2\hat{\sigma}_1^2}\right) \exp\left(-\frac{(r_2-q)^2}{8}\right)\ dr_1\ dr_2\ dq
\end{aligned}$$

which yields $I_\infty(1) \approx 0.0189 + .0053 \approx 0.0242$ and $I_\infty(2) \approx 0.0112 + 0.0043 = 0.0155$. Recall that the influence of R1 with precision $\hat{\sigma}_1$ when she is the sole reviewer was computed in Appendix B.5.1 to be $I_\infty(1) \approx 0.0266$ and $I_\infty(2) \approx 0.0146$. Thus, while the super-reviewer R1 loses influence as compared to the monopolistic setting, the casual reviewer R1 *gains* influence as compared to when she is a monopolist. The reason is that when R1 is high-skilled, the additional reviewer "steals some of the show," whereas when she is low-skilled, an additional reviewer bolsters the credibility of R1's review in the event that she corroborates it.[24]

To see how much firms are willing to bribe R1, recall the review written by R1 for a bribing firm is $r^B_{1,t} = q_t + \varepsilon_t + |\varepsilon'_t|$ and for a non-bribing firm is $r^{NB}_{1,t} = q_t + \varepsilon_t - |\varepsilon'_t|$ for $q_t, \varepsilon_t \sim \mathcal{N}(0,1)$ and

[23] This comes from the generalized inverse-variance formula of the signal extraction problem we considered with a single reviewer, given in Theorem 1. Consumers estimate $\mathbb{E}[q_t|\mathbf{r}_t, \hat{\boldsymbol{\omega}}]$ as:

$$\frac{\sum_{j=1}^k r_{j,t}/\hat{\sigma}_j^2}{1 + \sum_{j=1}^k 1/\hat{\sigma}_j^2}$$

[24] This is because type L would agree with type H occasionally. For example, imagine a product with extreme quality such that the signals of the L and H reviewers are similar. In that case they write similar reviews and the L reviewer, R2, shares some of the influence that used to only belong to R1. On the other hand, seeing two similar reviews from two type L reviewers increases the belief of the consumers that these reviewers know what they are talking about.

$\varepsilon_t' \sim \mathcal{N}(0,3)$. Thus, we must compute $\mathbb{P}[r_{1,t}^B + r_{2,t} \geq 3] - \mathbb{P}[r_{1,t}^{NB} + r_{2,t} \geq 3]$:

$$
\begin{aligned}
&\mathbb{P}[r_{1,t}^B + r_{2,t} \geq 3] - \mathbb{P}[r_{1,t}^{NB} + r_{2,t} \geq 3] \\
=&\mathbb{P}[r_{1,t}^B \geq 3 - r_{2,t}] - \mathbb{P}[r_{1,t}^{NB} \geq 3 - r_{2,t}] \\
=&\mathbb{P}[|\varepsilon_t'| \geq 3 - r_{2,t} - q_t - \varepsilon_t] + \mathbb{P}[|\varepsilon_t'| \leq r_{2,t} + q_t + \varepsilon_t - 3] \\
=&\frac{1}{\sqrt{3}(2\pi)^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{(3-r_2-q-\varepsilon)_+}^{\infty} \exp\left(-\frac{q^2}{2}\right) \exp\left(-\frac{\varepsilon^2}{2}\right) \exp\left(-\frac{(r_2-q)^2}{8}\right) \exp\left(-\frac{(\varepsilon')^2}{6}\right) \; d\varepsilon' \; dr_2 \; d\varepsilon \; dq \\
+&\frac{1}{\sqrt{3}(2\pi)^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{0}^{(q+\varepsilon+r_2-3)_+} \exp\left(-\frac{q^2}{2}\right) \exp\left(-\frac{\varepsilon^2}{2}\right) \exp\left(-\frac{(r_2-q)^2}{8}\right) \exp\left(-\frac{(\varepsilon')^2}{6}\right) \; d\varepsilon' \; dr_2 \; d\varepsilon \; dq \\
=&\frac{1}{\sqrt{3}(2\pi)^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{3-q-\varepsilon} \int_{3-r_2-q-\varepsilon}^{\infty} \exp\left(-\frac{q^2}{2}\right) \exp\left(-\frac{\varepsilon^2}{2}\right) \exp\left(-\frac{(r_2-q)^2}{8}\right) \exp\left(-\frac{(\varepsilon')^2}{6}\right) \; d\varepsilon' \; dr_2 \; d\varepsilon \; dq \\
+&\frac{1}{\sqrt{3}(2\pi)^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{3-q-\varepsilon}^{\infty} \int_{0}^{q+\varepsilon+r_2-3} \exp\left(-\frac{q^2}{2}\right) \exp\left(-\frac{\varepsilon^2}{2}\right) \exp\left(-\frac{(r_2-q)^2}{8}\right) \exp\left(-\frac{(\varepsilon')^2}{6}\right) \; d\varepsilon' \; dr_2 \; d\varepsilon \; dq \\
\approx&0.1449 + 0.0829 = 0.2278
\end{aligned}
$$

which implies that the maximal bribe the firm would be willing to offer is $0.2278/3 = 0.0759$. This is almost the bribe the firm is willing to offer when R1 is just a monopolist.

Note that when $\beta = 4$, the reviewer loses influence $4(0.0242 - 0.0155) = 0.0348$, yet the average bribe possible is as large as 0.0379 (half of 0.0759 given that $\theta = 1/2$, as before). Therefore, **the environment is no longer bribe-proof**: there is an equilibrium where the type H reviewer mimics type L in the presence of a second reviewer who is type L.

What is the impact on consumer welfare? When R1 is a monopolist, the consumer receives a signal $s_{1,t}$ with variance 1. When R1 and R2 compete, the consumer receives two independent signals $s_{1,t}$ and $s_{2,t}$ both of variance 4, which is equivalent to one signal (the average of both as per Footnote 23) with variance 2. Thus, the consumer is worse-off with the additional competition because it provides an avenue for bribery that makes good reviewers more willing to compromise their influence.