Strategic Reviews*

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 away from its first-best levels, and consequently leads to detrimental effects on consumer utility. This paper builds a dynamic two-sided reputation model where a reviewer can inflate her reviews in exchange for bribes. The problem the reviewer faces is the following: if she accepts bribes and misrepresents her reviews, then she builds her reputation as an inaccurate reviewer and eventually makes consumers less likely to follow her recommendations, which 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 the absence of bribes can sometimes lead to the *lowest* possible consumer utility, and that the introduction of bribes in these environments can restore some of that utility via implementing second-best information transmission.

1 Introduction

Firms have long recognized the impact that product reviews have on consumer purchasing behavior, and as a result have attempted to inflate these reviews in order to boost demand. These attempts have ranged from the farcical, like when Sony fabricated the movie critic 'David Manning' to shower its movies with praise[1] to more methodical quid-pro-quo arrangements, where side payments —"bribes"— are provided by firms in exchange for inflated ratings from reviewers. 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

*We thank Nick Arnosti, Ali Makhdoumi, and particpants at the Networks, Matching, and Platforms workshop for their helpful comments.

† Ross School of Business, University of Michigan.

‡ Laboratory for Information and Decision Systems, Massachusetts Institute of Technology.

[1]Sony ended up having to pay a settlement to those consumers who saw the movies based on the reviews of Sony's imaginary critic (Guardian (2004)).

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)).

In this paper, we focus on influential reviewers like the example above, but the model we build is equally suited to settings like the inspection and certification markets, where one side of the market makes consumption decisions based on the information it obtains from an expert who reviews or certifies the products and services offered by the other side of the market. Obviously, the presence of bribes distorts the information contained in reviews away from its first-best levels, and as a result can make consumers purchase products of low quality and/or avoid products of high quality, leading to an overall detrimental effect on consumer utility. This prompted the FTC to start taking measures against compensated reviews.[2] In this paper, we characterize the environments that are vulnerable to this kind of distortion, and in doing so uncover interesting properties about the effects of regulations. We also show that the presence of bribes is not always undesirable, and can in fact lead to an improvement in consumer utility under some circumstances.

We start by developing a two-sided reputation model that captures the above environments. A reviewer with a *private* ability type is situated between consumers and firms. The reviewer's type determines her accuracy in evaluating the quality of the products. Consumers learn this type over time as they compare product reviews with the ex-post realized qualities of the products. The more accurate the reviewer is, the more the consumers listen to her and the more influence she wields over their decisions. The reviewer may derive utility from being influential (for example because she gets more ad revenue on her dissemination channels of choice) and therefore would like to uphold her reputation with consumers as an accurate reviewer. At the same time, the reviewer wishes to maintain her reputation (and influence) with the consumers so that biased reviews actually translate into more consumption and greater profits for the firm, thereby making the firm willing to bribe.

The above setup resembles classic reputation problems (e.g. Kreps and Wilson (1982); Sobel (1985)) in that the reviewer wants to convince consumers that she is committed to playing as her true type. This however is complicated by the possibility that firms might offer bribes to the reviewer in exchange for biased reviews. The reviewer therefore finds herself simultaneously involved in two games. She can try to increase her payoff by maximizing her influence with consumers, but she can also increase her payoff by accepting bribes from firms. These two

[2] In addition to the example cited earlier, the FTC also started fining companies that are involved in review manipulation, see FTC (2019).

objectives are incompatible, since maximizing influence requires her to build a reputation for accurate reviews, which deters firms from offering her a bribe. On the other hand, accepting bribes requires the reviewer to misrepresent her reviews, which can diminish or completely remove her influence, and if that happens then firms are no longer interested in bribing the reviewer since she is no longer influential. How should the reviewer play this game? Can she maintain influence *and* accept bribes? And what is the impact of this behavior on consumer utility?

Contribution and overview of results We answer the above questions by analyzing a *dynamic, long-run repeated-game* model. Our model is unique in that it captures the dynamics of these markets while explicitly modeling *all* strategic actors (consumers, reviewers, and firms) in an incomplete information setting. This leads us to study signaling and cheap-talk games, the equilibria of which are notoriously difficult to analyze and refine (see Chen et al. (2008)), and this analysis is further complicated by the presence of two simultaneous and interlinked games (reviewer vs. consumers and reviewer vs. firms). Nevertheless, under rudimentary assumptions, we are able provide a technical characterization (existence and uniqueness) in Theorems 1 and 2 and a characterization of the optimal equilibrium strategies for all players in Theorems 3 through 5. We use these results to describe which environments are "bribe-proof". An environment is bribe-proof if the reviewer always reports her true signal. This can happen as a result of firms choosing not to offer bribes (because they believe that bribing will not be profitable for them), or as a result of the reviewer not accepting bribes even when they are offered (because of the damage to her reputation). We find that in the unique equilibrium, a reviewer either declines bribes and reports truthfully (in which case consumer utility is maximized) or commits to mimicking a lower type (i.e. less-skilled) reviewer forever (thereby distorting the information that consumers get away from first-best). This characterization allows us to derive the following results, together with their policy implications:

Bribes and Consumer Utility: Our model includes an auditing technology that operates independently of all agents. This means that a firm that chooses to bribe can be caught and punished. The strength of the auditing technology determines the proportion of firms that are willing to offer bribes (but who may still choose not to offer bribes in equilibrium) so that the higher the audit rate, the lower the proportion of those firms. While it is natural to think that decreasing the proportion of firms that offer bribes should improve consumer utility, Theorem 6 and Proposition 4 show that this is not necessarily the case. In particular, an increase in the audit rate that dissuades more firms from offering bribes may actually transition the system from being bribe-proof to becoming vulnerable to bribes and information distortion.[3] This non-monotonicity implies that such regulations should not be undertaken unless they *substantially* decrease the proportion of firms that are willing to offer bribes, instead of merely reducing it.

Our model also allows us to derive interesting comparative statics that describe how consumer utility changes as a function of the ex-ante information that the market has about the skill of the reviewer. When reviewers report truthfully, consumer utility is always increasing in the quality of this information. However, Proposition 5 and Proposition 6 together with Example 3 illuminate an important counterintuitive finding: increasing the quality of information (through an increase in the precision of the reviewer or an increase in the likelihood that the reviewer is highly skilled, or both) may not necessarily lead to better outcomes for consumers when some firms are willing to bribe. As we show, more skilled reviewers have more leeway in inflating their reviews, and this makes the prospect of bribing them more appealing for 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.

Finally, while bribes are generally harmful to consumer utility, their presence can sometimes increase the efficiency of the system. When writing reviews is costly (e.g., time consuming), reviewers can become extremely noisy or stop writing reviews altogether. Proposition 7 shows that by offering bribes, the private market provides a mechanism that funds the process of review writing. Crucially, the reviews still have to be relatively informative, because otherwise consumers will stop listening to the reviewer's recommendations and she will lose her value to the bribing firms. Thus, while the information in these reviews is necessarily distorted away from its first-best levels, bribes can still provide the correct incentives to implement second-best information transmission in the marketplace.

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, for example, 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

[3] Recall that an environment does not necessarily admit bribes simply because some firms are willing to offer them, given the current audit rate. If these bribes are not profitable for the firm, or not sufficient to compensate the reviewer, there may be no bribes in equilibrium even when some firms would be *willing* to bribe.

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. (2019)).

By contrast, our paper tries to capture these environments when consumers and firms interact through a reviewer, who, unlike the work mentioned up to this point, is a strategic self-interested player herself, 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 (2019) studies competition between influencers who choose how much 'sponsored' vs. 'organic' content to provide. Our work differs from these papers on two important dimensions. First, the expert (reviewer) in our model has an *unknown* type that must be learned by *both sides* of the market, and this learning happens through communicating about products of unknown quality. The static models in the above papers preclude the possibility of learning in this incomplete information setup.[4] Second, by incorporating this temporal aspect into our model, we are able to show how the expert can utilize this type uncertainty along with the dynamics of the environment to consistently shape the buyers' and sellers' beliefs, leading to different outcomes from the static models.

Finally, 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)), but unlike this literature, the reviewer in our setup plays two simultaneous reputation games on both sides of the market. This brings our paper closer 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 a seller in order to attract buyers in a two-period model. Our work shares some similarities with this paper in that the agent must trade off high precision (and more influence) with low precision (and increased likelihood of over-representing the true quality), but there are also several important distinctions. First,

[4]For example, Fainmesser and Galeotti (2019) assume that the expert's choice of how much sponsored content to post is observed by consumers, with the reasoning being that if the game were to be repeated, the expert would be detected by consumers if she deviates from that choice over time. By contrast, the presence of a hidden type and reputation in our model jointly allow the expert to consistently misrepresent her type without being detected.

while reputation is a long-run concept that is naturally studied in an infinite horizon setting (see Mailath et al. (2006)), it is necessarily temporary in a two-period model, which leads to qualitatively different results. Infinite-horizon models imply that the agent is disciplined more strictly and cannot misrepresent herself extensively for short-term gains, and hence the findings of Bouvard and Levy (2017) do not generalize to these interactions once the horizon extends beyond two periods. There are also differences that stem from that fact that while firms elect to be certified, they cannot choose not be reviewed, as per the Consumer Review Fairness Act of 2016 (FTC (2016a)). Finally, we explore different comparative statics and welfare implications, e.g. changing the strength of the auditing technology, as opposed to the effects from additional competition.

In the next section, we provide an example of some of the main ideas in the paper. Our formal model and equilibrium preliminaries are presented in Section 3. Section 4 provides existence and uniqueness results, while Section 5 describes the equilibrium strategies played by all three player types. Section 6 brings together all the previous technical results to examine the effect of bribes and regulation on consumer utility, and Section 7 concludes the paper.

2 Reduced-Form Example

We provide intuition for our results through a stylized example. To facilitate exposition, we abstract away several details of our full model. The points at the end of this section help connect the ideas presented in the example with the corresponding results in the paper.

A reviewer repeatedly interacts with a mass of consumers, where she writes reviews that consumers use to make 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. Can the reviewer successfully and repeatedly steer consumers towards products sold by her friends without having the *Bayesian* consumers adjust their consumption decisions to account for this?

We provide some details to answer this question. Firm t arrives at time t and sells a product whose quality q_t is distributed as $\mathcal{N}(0,1)$. After sampling the product, the reviewer 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. The reviewer's private type/skill ω is either **High** ($\sigma_H = 1$) or **Low** ($\sigma_L = 2$), 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. Naturally, 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 told 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 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 and that she might want to favor these firms, 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 we assume, for the purpose of this example, is an outcome that the reviewer 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**, which will be reflected in the distribution of her reviews being different from the distributions associated with these types. Consumers will then find out the reviewer is not telling the truth and punish her (by, for instance, ignoring her recommendations in the future).

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 5.2). 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 $r_t \sim \mathcal{N}(0,2)$, and 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**, but she can do it 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, as we mentioned earlier, is the distribution of the reviews written by a truthtelling reviewer of type **L**. In particular, if the reviewer behaves in this way then consumers cannot distinguish whether she is indeed type **L**, or is a type **H** pretending to be type **L** (and biasing her friends in the process).

Now, notice that even though the reviews over all firms are normally distributed, the reviews for friends are distributed as a half-normal, and so the average review for friends of quality q_t is:

$$\mathbb{E}[r'_t|q_t] = \mathbb{E}[s_t + \varepsilon'_t|q_t] = \mathbb{E}[s_t + |X_t|\ |q_t] = q_t + \sqrt{\frac{6}{\pi}}$$

whereas 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, we know that, having learned that the reviewer is playing as type **L**, she believes the expected quality given a review r'_t is equal to $\mathbb{E}[q_t|r'_t, \hat{\omega} = \mathbf{L}] = \frac{r'_t}{1+\sigma_L^2} = r'_t/5$, which means that the average quality belief for friends is:

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

and similarly, consumers (on average) believe the quality of non-friends to be $-\sqrt{\frac{6}{25\pi}}$. Because consumers cannot differentiate between 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 previous example shows that the reviewer can inject noise into her reviews in a particular way that ends up biasing consumers towards the firms that she prefers. Note that while this dilutes the information in the review, it does not completely remove its value to consumers, and so they continue perusing the reviews despite the possibility that they are biased. Importantly, and as mentioned above, consumers cannot tell whether the reviewer is a real type **L** or is a type **H** pretending to be type **L**: the former cannot bias the consumer towards certain products, but the latter can.

This example leaves open several questions that we address in the paper. We state these questions below and reference the results that include their answers at the end of each question.

1. As we saw in the example, the reviewer is misrepresenting her true signal to help her friends. Can she always do this without consumers learning (and adjusting) to this? (Theorems 3 and 4).

2. For simplicity, we assumed in the above example that some firms are friends of the reviewer. In our full model however, 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 does not follow through with a biased review. Firms, anticipating this, might not want to offer a bribe to begin with. Does the equilibrium of this game admit situations where the reviewer still gets (some) firms to offer bribes? (Theorems 2 and 5).

3. Generally, and as can be observed in the above example, the presence of bribes hurts consumer utility as consumers become biased towards certain products and against others (Proposition 2). In our example, half of the firms did not offer bribes. How does consumer utility change if the proportion of firms willing to offer bribes is different from half? (Theorem 6 and Proposition 4).

4. In this example, the different reviewer types were given by $\sigma_L = 2$ and $\sigma_H = 1$ and were equally likely. How do consumer utility and the size of bribes change as we decrease the skills of the reviewer (e.g., $\sigma_L = 4$ and $\sigma_H = 2$) or change the probability of her type (e.g., the reviewer is high-skill with probability 80%)? (Propositions 5 and 6).

As mentioned in the introduction, our model allows the possibility that the reviewer derives utility from being influential, independent of whether she accepts bribes. By accepting bribes and pretending to be a worse type, the reviewer loses some of that influence with consumers. Theorem 6 and Example 2 show how the tradeoff between influence and accepting bribes shape the decision of the reviewer and consequently, consumer utility. In the example in this section, if the reviewer cares a lot about her influence, she may choose not to favor her friends, and as a consequence, would not receive bribes. Our model also allows the reviewer to have no inherent value for influence, and Proposition 7 examines whether bribes can actually improve consumer utility in that case. In particular, the reviewer's friends may provide funding to the reviewer, who might otherwise choose not to write any reviews at all.

3 Model

We consider the following model. Firms are agents that produce and sell a good to a unit mass of consumers on $[0, 1]$. Firms arrive sequentially in discrete time, and "market" their product to consumers through a reviewer, who samples the firm's product and writes a review. Consumers then choose whether to listen to the reviewer and whether to purchase the good. The reviewer and consumer are infinitely-lived.

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 $\omega = H$ 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 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 $q_t \sim \mathcal{N}(0, 1)$.

(b) The reviewer does not directly observe the quality q_t of the good, but instead receives an unbiased, noisy signal s_t. In particular, $s_t = q_t + \varepsilon_t$ where ε_t is i.i.d. across time, and $\varepsilon_t \sim \mathcal{N}(0, \sigma_\omega^2)$.

(c) The reviewer posts a review $r_t \in \mathbb{R} \cup \{\varnothing\}$ of the product. Note that the reviewer can abstain from posting a review, but if she does post a review, she incurs a cost C. This cost may account for her time examining and using the product, producing media to disseminate her review over the approporiate channel(s), etc. Before posting the review, the firm may offer a bribe b_t to the reviewer, representing a "side payment" so that the reviewer speaks favorably about the product.

(d) Consumers have a belief π_t about the type of the reviewer. Every consumer $i \in [0, 1]$ observes the review r_t and then, based on r_t and π_t, chooses whether to purchase the good $x_{i,t} \in \{0, 1\}$ at unit price.

(e) Consumers who elect $x_{i,t} = 1$ are *active* consumers, and receive feedback about q_t through their observable payoff. Each active consumer receives an experience given by $e_{i,t} = q_t + \eta_{i,t}$, where $\eta_{i,t}$ are i.i.d., distributed symmetrically around zero, and have finite variance.

(f) After the period is over, the true quality q_t of the product is revealed through a public signal observed by everyone. This can be thought of as a critical or consumer consensus that emerges after the product has been out in the marketplace for a while.[5]

3.2 Payoffs

We now describe the payoff structure for all three agent types:

Consumers The continuum of consumers on $[0, 1]$ have heterogeneous preferences for quality. Each consumer i has an outside option ϕ_i she can obtain (for the same unit price) instead of purchasing the good. Throughout the rest of the paper, we will impose the following structure on ϕ_i:

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

This assumption guarantees that for any quality q_t, some positive fraction of consumers always purchases the good and some positive fraction of consumers abstain. Consumers are myopic and maximize their current-period utility given posted reviews r_t and beliefs π_t; that is, consumer i chooses her consumption $x_{i,t} \in \{0, 1\}$ according to:

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

It is easy to see that each consumer will employ a cutoff strategy $x^*_{i,t}(r_t, \pi_t) = 1$ iff $\mathbb{E}[q_t|r_t, \pi_t] \geq \phi_i$. Therefore, we can define $X^*_t(r_t, \pi_t) = \mu\left(\{i : x^*_{i,t}(r_t) = 1\}\right)$, where μ is the Lebesgue measure, as the *total* consumption of the product.[6] Note that X^*_t depends on both the review r_t and also the history of reviews and observations from the consumer in the past, via the belief π_t.

Firms Firms receive a payoff proportional to the number of active consumers $0 < X^*_t < 1$ (i.e., those who purchase the good) less any bribes they pay to the reviewer, b_t. We may normalize the bribe payment to be in product units. Thus, we can write the total payoff of the firm simply as $U_t = X^*_t(r_t, \pi_t) - b_t$. Implicit in this form are that firms are myopic (as would be the case when arriving firms are different in each period). Firms cannot observe bribes from previous periods, but instead must infer from past play (via public history, as we discuss shortly) whether a reviewer may be receptive to bribery.

Firms can be one of two types: truthful or strategic. Truthful firms do not offer bribes, while strategic firms can choose to offer a bribe or not. The proportion of truthful firms depends

[5]For example, the consensus that the reception on the iPhone 4 was especially poor (see Helft and Bilton (2010)) or that the keyboard on the 2015-2018 Apple MacBook Pros fail often (see Stern (2019)).

[6]Note under Assumption 1 the set $\{i : x^*_{i,t}(r_t, \pi_t) = 1\}$ is always an interval, so in particular it is always measurable.

on an *auditing technology* that operates independently of all agents. This technology might investigate the firm (with probability α) and such an audit perfectly reveals whether the firm bribed the reviewer ($b_t > 0$) or not ($b_t = 0$). The higher the audit rate, the higher the proportion of truthful firms.[7] We denote by $\theta(\alpha)$ the proportion of truthful firms, which is continuous, strictly increasing in α and satisfies $\theta(0) = 0$ and $\theta(1) = 1$. Similar to consumers, we assume that firms have an outside option $\gamma \geq 0$, and that they only enter the market place if they expect their payoff to exceed that option. [8]

Reviewer In addition to her skill type, the reviewer has another private attribute that describes whether she is truthful or strategic. If she is truthful (with small probability $\epsilon > 0$), then she always reports her signal ($r_t = s_t$),[9] whereas if she is strategic, she may report any r_t desired.[10] A strategic reviewer is a far-sighted agent who *may* care about her reputation and influence, with discount factor $\delta \in (0, 1)$. We can define the *influence*, I_t, of the reviewer to be equal to the average (squared) deviation in consumption following her review:

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

In other words, the reviewer's influence measures how much the marginal consumption reacts to her review on average, compared to the baseline of $\mathbb{E}[q_t] = 0$. We assume that the reviewer cares about her influence with propensity $\beta \geq 0$.[11] The case $\beta > 0$ accounts for the possibility that the reviewer may derive direct revenue and/or utility from her influence, for example through more ads on her dissemination channel or through some altruistic or egotistical utility from helping consumers and influencing their decisions. Reviewers are patient and maximize their (average)

[7] We assume that firms are heterogeneous in how they respond to getting caught in an audit. To keep our model parsimonious, we do not explicitly model this response mechanism. For example, firms might have different levels of risk aversion that is perhaps connected to their reputation in the marketplace, and depending on the audit rate and the penalty of getting caught it might not be profitable for all of them to offer a bribe. Similarly, firms can have different cost structures that might again make it profitable for some of them to offer bribes but not others.

[8] The outside option —which accounts for the fact that firms may have opportunity costs (e.g., other ventures) when deciding to enter the market— helps eliminate certain unnatural equilibria that arise where the reviewer can aggressively "punish" the firm for not bribing, thereby soliciting massive bribes just to avoid punishment.

[9] The existence of a truthful type is a standard assumption in the reputation literature (e.g. Fudenberg and Levine (1989)) to eliminate "babbling" equilibria where the reviewer never posts informative reviews, and the consumer forever assumes her prior belief $\mathbb{E}[q_t|r_t] = \mathbb{E}[q_t] = 0$ (see Best and Quigley (2016)).

[10] Note that because the reviewer's payoff does not directly depend on r_t, but only through how r_t affects the actions of the firm and consumer, this is a game of repeated cheap talk (with reputational concerns). This has been studied in models such as Ottaviani and Sørensen (2006), but unlike this model, the reputation of the reviewer is *endogenous* and studied in a fully dynamic setting as opposed to a two-period model with an exogenous continuation payoff.

[11] For technical reasons, throughout we will always assume $\beta > 0$, but will simply write $\beta = 0$ to mean $\beta \to 0$ (i.e., the influence propensity vanishes).

discounted payoff:

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

where $c_t = C$ if the reviewer posts a review and 0 otherwise. In the special case where the reviewer does not care or intrinsically value her influence (i.e., $\beta = 0$), she simply maximizes the discounted sum of bribes that she receives over the entire duration of the game, less her running costs of writing the reviews. Of course, this might still require the reviewer to maintain some influence with the consumers, otherwise the firms would not be interested in offering her a bribe.

3.3 Belief Evolution and Equilibrium Preliminaries

The reviewer has a private history h_t consisting of both private quality signals $\{s_1, s_2, \dots, s_{t-1}\}$ and bribes $\{b_1, b_2, \dots, b_{t-1}\}$. There is also a public history H_t which consists of the reviews $\{r_1, r_2, \dots, r_{t-1}\}$ and the true qualities $\{q_1, \dots, q_{t-1}\}$ in every prior period. Everyone observes the public history H_t at the beginning of period t, and the reviewer also observes h_t.

Firm t and the consumers hold beliefs in period t about the type of the reviewer. This belief is entirely determined by the public history H_t, given by some $\pi_t \in \Delta(\{\sigma_L, \sigma_H\})$ (whereas, naturally, the reviewer knows her own type). We consider the following equilibrium concepts as the reviewer becomes very patient ($\delta \to 1$):

1. *Consumer-reviewer equilibrium*: Consider a (fixed) sequence of bribes $b_t(H_t)$ which may be conditional on the public history. Let $\tilde{b}_t(H_t)$ be the sequence of bribes that are forced to zero at every time t where the firm is honest. Then a *consumer-reviewer equilibrium* is a perfect Bayesian equilibrium of the dynamic game between the reviewer and the consumer with bribes given by $\tilde{b}_t(H_t)$, where $b_t(H_t)$ is common knowledge.

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

Robust Equilibrium. The above solution concept admits a multiplicity of equilibria. Here, we introduce a refinement that allow us to obtain uniqueness. The consumer-reviewer game is a cheap-talk game, which is a subset of signaling games. Unlike other signaling games though, cheap talk games are notoriously more difficult to refine, with a sizable literature that proposes different selection arguments (see Chen et al. (2008); Sobel (2009) for a summary and discussion). For our purposes, we focus on the *risk-dominant* equilibrium (Harsanyi et al. (1988)). The reason behind this is simple and intuitive: one of the possible equilibria is what we call a *babbling-trigger* equilibrium (see Best and Quigley (2016)), whereby consumers stop listening to the reviewer forever if they believe she is a strategic type who is accepting bribes to inflate her reviews. The fact that babbling-trigger guarantees the reviewer a perpetual payoff of zero and the fact that the reviewer is uncertain whether consumers will select this equilibrium jointly provide an incentive for the reviewer to "play it safe" and assume that consumers will indeed play babbling-trigger. This gives rise to the risk-dominant equilibrium. Second, to break indifferences, we suppose the auditing technology punishes firms in a way that increases with the size of the bribe offered. However, we assume the fraction of the penalty which is proportionate to the bribe is much smaller than the penalty at-large. Therefore, when indifferent between two or more distinct bribes, firms choose to break ties in favor of the lower bribe. We say an equilibrium is *robust* if (i) it is a full equilibrium, and (ii) it satisfies the aforementioned two conditions.

4 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 use this as a foundation for our characterization of equilibrium strategies in Section 5 and for the comparative statics and second-best characterization in Section 6.

By standard existence results, one can see:

Theorem 1. *Given any discount factor δ:*

(a) For any fixed bribe scheme $b_t(H_t)$, a consumer-reviewer equilibrium exists;

(b) A full equilibrium always exists.

While Theorem 1 guarantees our solution concept is well-posed for any δ, we will focus on the case of a patient reviewer, i.e. $\delta \to 1$. If the reviewer is less patient, she may be able to manipulate consumers by lying for some time, eventually driving consumers away, and then forever after abstaining from writing reviews. Because this is unsurprising, we investigate whether any form of review manipulation can persist forever. We capture this by assuming the reviewer values payoffs long in to the future, and therefore obtains significant value from her reputation as a good reviewer. For these same reasons, and because consumers and firms need to learn the reviewer's skill over time, we look only at strategies in the limit $t \to \infty$.

Let $\{b_t(\delta_k), r_t(\delta_k), X_t(\delta_k)\}_{k=1}^n$ be a sequence of (full) equilibrium strategies for a sequence of parameters δ, with $\delta_k \to 1$.[12] We define a limit equilibrium when $\delta \to 1$:

Definition 1. We say the equilibrium strategies (b_t, r_t, X_t) are a *limit equilibrium* if for any sequence $\delta_k \to 1$, there exists a sequence of equilibrium strategies $\{b_t(\delta_k), r_t(\delta_k), X_t(\delta_k)\}_{k=1}^n$ that converges (in $\sup$-norm) to (b_t, r_t, X_t) as $k \to \infty$.

Second, we look at the equilibrium strategies of a limit equilibrium for large t. To do this, we consider fixed (time-invariant) strategies (b^*, r^*, X^*) and look to see if the limit equilibrium approaches any strategy of this form as $t \to \infty$. Formally:

Definition 2. We say the limit equilibrium *approaches* a strategy (b^*, r^*, X^*) if $\limsup_{t\to\infty} ||(b_t^*, r_t^*, X_t^*) - (b^*, r^*, X^*)||_\infty = 0$.

In such cases, we may refer to (b^*, r^*, X^*) as the limit equilibrium, with a slight abuse of terminology, but where the distinction should be obvious. Even in this setting, we have a multiplicity of equilibria because of indifferences that exist for the reviewer. These indifferences exist because of a *lack of commitment* by the reviewer: if the firm bribes the reviewer, she is not committed to giving any particular review.

This multiplicity can be seen most transparently by noting that there is always a trivial equilibrium. We say a full equilibrium is *trivial* if incoming firms always offer bribes $b_t = 0$, and reviewers always report truthfully $r_t = s_t$. In the one-shot game, the reviewer can always choose to accept the bribe but then report her belief of the product quality truthfully anyway, and if the firm is aware the reviewer might do this, then it offers no bribe to begin with. As repeating the static equilibrium is always a repeated equilibrium, we see that this trivial equilibrium can exist in any dynamic setting. On the other hand, a full equilibrium is *bribing* if there exists bribes (almost surely) at some point along the history of play. Under certain conditions, there may be no other equilibrium (other than the trivial one) with bribing, 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, we concentrate on bribes which are maximally-supported between the firm and the reviewer, as we define next:

Definition 3. We say a limit equilibrium, (b_1^*, r_1^*, X_1^*), is *maximal,* if for any other limit equilibrium (b_2^*, r_2^*, X_2^*), given fixed bribes b_1^*, (r_2^*, X_2^*) is still a consumer-reviewer equilibrium, and b_1^* is the largest bribe supported in any full limit equilibrium with (r_1^*, X_1^*).

[12]For ease of notation, we have collapsed the firm's decision about entering the market, and then whether (and how much) to bribe, into a single choice variable b_t. This can be done by introducing an additional action, $b_t =$ exit, which (deterministically) awards the firm its outside option γ.

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 is bribe-proof in order to know whether the setting is bribe-proof under any equilibrium. With this, we obtain the following uniqueness result:

Theorem 2. *There always exists a maximal robust equilibrium and it is generically unique.*

When the environment is bribe-proof, the unique maximal robust equilibrium coincides with the trivial equilibrium. The proof of Theorem 2 is the backbone for characterizing the behavior of the different players in the next section.

5 Equilibrium: Characterization

We now provide a characterization of the unique robust equilibrium strategies. We start with the reviewer, and provide the general strategy she uses to manipulate (or not) her signal through her review. Next, we characterize the consumer's Bayesian consumption decision, given the strategy of the reviewer in the consumer-review equilibrium. Finally, we comment on the conditions that allow firms to bribe in the full equilibrium, and use this to analyze consumer welfare.

5.1 Reviewer

First, we define a *mimic-down* strategy. A reviewer of type ω may mimic any $\hat{\omega}$ with $\sigma_{\hat{\omega}} \geq \sigma_{\omega}$. In other words, the reviewer can misrepresent her precision, but cannot artificially improve her own skill – she may only "worsen" it. This means that a high-type reviewer may appear as high or low-type, whereas the low-type must represent herself correctly. Thus the reviewer follows the following strategy:

(a) If $\hat{\omega} = \omega$, then the reviewer reports truthfully in each period, $r_t = s_t$.

(b) If $\hat{\omega} \neq \omega$, then the reviewer produces an orthogonal (to s_t) component of noise ε'_t at every period, where $\varepsilon'_t \sim \mathcal{N}(0, \sigma^2_{\hat{\omega}} - \sigma^2_{\omega})$. She then reports the review $r_t = s_t + \varepsilon'_t$.

Then we obtain the following characterization of reports for the reviewer:

Theorem 3. *In the (limit) equilibrium, the reviewer plays the same* mimic-down *strategy at every point in time.*

The reviewer therefore commits to playing a particular type throughout the entire horizon. Note that a reviewer following a mimic-down strategy has reviews $r_t \sim \mathcal{N}(0, 1 + \sigma^2_{\hat{\omega}})$. We refer to ω as the true type of the reviewer, whereas $\hat{\omega}$ is the *effective* type she chooses to mimic. As we show, consumers and incoming firms cannot identify whether the reviewer is of type $\hat{\omega}$, or a more skilled type ω mimicking the type $\hat{\omega}$. For instance, an effective low-type may be a true low-type, or the high type who has chosen to mimic the low-type by injecting noise into her review.

5.2 Consumers

For any given effective type $\hat{\omega}$, we can define the *inverse-variance* expression for quality given review r_t as:

$$IV(r_t, \hat{\omega}) = \frac{r_t/\sigma^2_{\hat{\omega}}}{1 + 1/\sigma^2_{\hat{\omega}}} = \frac{r_t}{1 + \sigma^2_{\hat{\omega}}}$$

Note here that the 1 in the denominator is a consequence of assuming $q_t \sim \mathcal{N}(0, 1)$. This is a classic signal extraction problem (e.g. Cochran (1954)); take $Z = X + Y$ where $X \sim \mathcal{N}(0, 1)$ and $Y \sim \mathcal{N}(0, \sigma^2_Y)$. Then $\mathbb{E}[X|Z = z] = \frac{z/\sigma^2_Y}{1+1/\sigma^2_Y}$ (i.e., the prior of $\mathcal{N}(0, 1)$ regresses the estimate of X closer to zero). The consumer thus "tempers" the review on both extremes: she believes very good reviews overhype the quality, whereas very negative reviews are overly harsh. Given this we can characterize the optimal consumption decision of every consumer:

Theorem 4. *Every consumer (eventually) assigns probability 1 to the correct effective type $\hat{\omega}$ played by the reviewer. Moreover, $\mathbb{E}[q_t|r_t, \pi_t]$ converges to $IV(r_t, \hat{\omega})$, and consumption $X^*_t(r_t) \to \phi^{-1}(IV(r_t, \hat{\omega}))$.*

Reviewers, as in Theorem 3, commit to an effective type which is eventually learned by all consumers over time. When the reviewer has a more skilled effective type, consumers place a higher weight on her reviews (and therefore she has more influence) compared to if she has a lower effective skill. On the other hand, the higher the effective type (e.g. a high type reviewer acting like her true type), the less leeway the reviewer has in biasing her reviews. While the reviewer can obtain maximal influence by truth-telling (i.e., choosing $\hat{\omega} = \omega$), but then cannot differentiate her review on any basis other than her quality signal. The reviewer thus faces the following tradeoff: if she chooses to mimic a lower skill type, she loses influence from the precision of her reviews, but has more flexibility to bias her review above or below her signal at her own discretion.

Theorem 4 shows why the reviewer in Section 2 was able to consistently manipulate the (Bayesian) consumers into purchasing particular products. Because consumers can only learn the effective type, they cannot distinguish between a reviewer who is biasing them towards specific products (in which case they would stop listening to her) or a reviewer who is just naturally less precise and favors no firms over others.

5.3 Firms

In the stylized example in Section 2, we assumed that the reviewer has preferences over certain firms, and we demonstrated how she can benefit the firms she favors through positive reviews. We now return to our model and recall that each firm that enters is either a truthful type (with probability $\theta(\alpha)$) or strategic type that may offer bribes to the reviewer in exchange for preferential reviews (with probability $1-\theta(\alpha)$), which is controlled by the intensity of the auditing technology α. We will often write just θ, suppressing the dependence on α.

To prevent situations where firms consistently prefer their outside options to entering the marketplace, assuming they can receive an honest review without having to bribe, we make the following assumption:

Assumption 2. The firm's outside option γ is upper-bounded by her *fair consumption*, denoted by $\bar{X}^*$, which is the expected value of consumption under an honest review from the worst-skill reviewer (i.e., $\mathbb{E}_{r_t \sim \mathcal{N}(0,1+\sigma_L^2)}[\phi^{-1}(r_t/(1+\sigma_L^2))]$).

One can interpret fair consumption as the worst (ex-ante) consumption the firm should expect given an honest review, before entering the market. We note this assumption also guarantees that non-bribing firms who anticipate other firms will not bribe at auditing technology α, would prefer to enter themselves than accept the outside option.[13]

For notational convenience, define $\mathcal{N}_\theta(\mu, \sigma^2)$ as the conditional normal distribution of the bottom θ percentile, and let $1-\mathcal{N}_\theta(\mu, \sigma^2)$ be the conditional normal distribution of the top $(1-\theta)$ percentile. We obtain the following result:

Theorem 5. *Suppose the reviewer is type ω mimicking type $\hat{\omega}$ in equilibrium. Then, on the equilibrium path at time t:*

(a) If the firm is truthful, she enters the market and offers no bribe ($b_t = 0$), and the reviewer writes the review $r_t = s_t + \varepsilon'_t$, where $\varepsilon'_t \sim \mathcal{N}_\theta\left(0, \sigma_{\hat{\omega}}^2 - \sigma_\omega^2\right)$.

[13] In reality, if non-bribing firms are aware that some firms may still bribe at audit rate α, 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 find it profitable to both enter and not bribe, regardless of whether these firms expect a biased-down review.

(b) *If the firm is strategic, she enters the market and offers a bribe* ($b_t = b^*$)*, and the reviewer writes the review* $r_t = s_t + \varepsilon'_t$*, where* $\varepsilon'_t \sim 1 - \mathcal{N}_\theta\left(0, \sigma^2_{\hat{\omega}} - \sigma^2_\omega\right)$.

If the reviewer decides to mimic a less precise type, she adds additional noise to her review so she can inject (positive) bias for those firms who bribe her. Truthful firms do not bribe and are punished by being placed in the bottom θ percentile of this additional adjustment, *relative to the reviewer's true quality signal.* On the other hand, strategic firms do bribe and are rewarded by being placed in the top θ percentile of this distribution. Therefore, similar to the example in Section 2, there is systematic bias from the consumers towards the products of bribing firms. Note that while bribing firms get preferential reviews, the reviewer's biases do not completely drown out the true measure of quality in the review, but rather dilutes it. This can be illustrated in the next example.

Example 1. 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 choose to commit to review biasing or not. If the reviewer plays effective precision $\hat{\sigma}$, then her influence is (eventually) given by $I_\infty(\hat{\sigma})$ (see Section 5.4).

Assume that only 50% of firms are willing to bribe at the current auditing rate α (i.e., $\theta = 1/2$) and their outside option is zero (i.e., $\gamma = 0$). 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 $r_t \geq (1 + \hat{\sigma}^2)/2$; therefore:

$$I_\infty(\hat{\sigma}) = \frac{1}{9} \frac{1}{2\hat{\sigma}\pi} \int_{-\infty}^{\infty} \int_{(1+\hat{\sigma}^2)/2}^{\infty} \exp(-q^2/2) \exp\left(-\frac{(r-q)^2}{2\hat{\sigma}^2}\right) \, dr \, dq$$

which 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 $r_t^B = q_t + \varepsilon_t + |\varepsilon'_t|$; if it does not bribe, it is distributed as $r_t^{NB} = q_t + \varepsilon_t - |\varepsilon'_t|$. 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 $\hat{\sigma} = L$:

$$
\begin{aligned}
\frac{1}{3}\cdot\left(\mathbb{P}[r_t^B \geq 5/2] - \mathbb{P}[r_t^{NB} \geq 5/2]\right) &= \frac{1}{3}\mathbb{P}\left[(q_t+\varepsilon_t) - |\varepsilon_t'| \leq 5/2 \leq (q_t+\varepsilon_t) + |\varepsilon_t'|\right] \\
&= \frac{1}{3}\mathbb{P}\left[|(q_t+\varepsilon_t) - 5/2| \leq |\varepsilon_t'|\right] \\
&= \frac{1}{3}\Bigg(1 - \frac{1}{2\sqrt{6}\pi}\int_{-\infty}^{0}\int_{\kappa}^{-\kappa} \exp\left(-\frac{(\kappa+5/2)^2}{4}\right)\exp\left(-\frac{\varepsilon^2}{6}\right)\, d\varepsilon\, d\kappa \\
&\quad - \frac{1}{2\sqrt{6}\pi}\int_{0}^{\infty}\int_{-\kappa}^{\kappa} \exp\left(-\frac{(\kappa+5/2)^2}{4}\right)\exp\left(-\frac{\varepsilon^2}{6}\right)\Bigg) \\
&\approx \frac{1}{3}\left(1 - 0.75 - 0.01\right) \approx 0.08
\end{aligned}
$$

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 mimicking down or is actually a low type). Similarly, the "on-average" bribe to a reviewer will be half of b_t, where b_t is the bribe amount of a strategic firm. Therefore, if $0.02 > 0.012\beta$ (i.e., $\beta < 5/3$) the strategic firm bribes an amount $.01 + .006\beta$. On the other hand, if $\beta > 5/3$, the environment is bribe-proof. $\square$

Finally, note the consumer's utility is always worse when there is a bribe equilibrium in Example 1. We show this formally in Proposition 2, but informally, this occurs because bribes decrease the precision of the estimator $q_t|r_t$, which leads to more ex-post unsatisfied consumers (higher ϕ_i than q_t) and more inactive consumers (higher q_t than ϕ_i) than optimal. In Example 1, a review $r_t \geq (1+\hat{\sigma}^2)/2$ has an expected quality q_t of:

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

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

5.4 Influence and Welfare

Based on Example 1, we see that as the high-skill type mimics down, both her influence over the consumer's decision and the welfare of the consumer decrease. This holds more generally, regardless of the consumer distribution.

Proposition 1. *For any ϕ_i, the influence $I_t(\hat{\sigma})$ for a reviewer who mimics type $\hat{\omega}$ converges almost surely as $t \to \infty$ to:*

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

where $I_\infty(\hat{\sigma})$ is strictly decreasing in the precision $\hat{\sigma}$, and $\mathbf{1}_{\hat{\sigma}}$ is the belief π_t that assigns probability 1 to precision type $\hat{\sigma}$.

Note that because the reviewer is extremely patient ($\delta \to 1$), the value of her influence is entirely pinned-down by $I_\infty(\hat{\sigma})$. Reviewers who mimic greater precision have a higher impact on the decisions of consumers. This is true, even though less precise reviews more aggressively endorse or slander products relative to those with higher precise reviews. Despite this, a mild endorsement from a high-skilled reviewer carries more weight than a more extreme one from a reviewer with a worse reputation for accurate reviews. Consumers learn to take such reviews with a grain of salt. For instance, in Example 1, the low-skill reviewer writes much higher variance reviews, but is about half as influential on purchasing decisions.

Definition 4. Consumer utility at time t is given by the average consumer utility of *active consumers* conditional on reviews r_t, $CU_t = \int_{i \in \mathcal{A}_t|r_t} (q_t - \phi_i)\, di$.[14] The average consumer utility at time t is the time-average of consumer utility up until time t, $\bar{CU}_t = \frac{1}{t}\sum_{\tau=1}^{t} CU_\tau$.

In other words, consumer utility measures the total surplus of the active consumers, given the information encoded in the review. The perfectly efficient set of active consumers is precisely $\mathcal{A}_t = \phi^{-1}([-\infty, q_t])$, that is, those consumers whose outside option is less than the quality of the good. However, the set of active consumers conditional on review r_t, $\mathcal{A}_t|r_t$, will possibly contain agents whose outside option exceeds q_t or not contain agents whose outside option falls below q_t. The extent to which $\mathcal{A}_t|r_t$ matches the efficient set of active consumers measures total welfare.

Proposition 2. *Average consumer utility converges almost surely as $t \to \infty$ to some $\bar{CU}_\infty(\hat{\sigma})$ which depends only on the effective type of the reviewer. Moreover, $\bar{CU}_\infty(\hat{\sigma})$ is strictly decreasing in $\hat{\sigma}$.*

The expected quality of the product, *conditional on the review endorsing the product*, is lower for less skilled reviewers. This relates directly to Proposition 1; even when the review threshold

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

for purchasing the product is greater, the expected quality is lower. To illustrate again with Example 1, middle consumers purchase the product only if the low-skill review exceeds 2.5, whereas they will purchase if the review written by the high-skill reviewer exceeds only 1. Even so, the expected quality of a product purchased based on the review of the skilled reviewer is higher (1.34 for H and only 0.903 for L), even though on average the less-skilled review will be stronger.

We make two additional comments about Proposition 2. The first is that while *average* utility for each consumer increases as the reviewer becomes more precise, this does not guarantee that for a given product, a more precise review leads to higher expected utility. For instance, if $\phi_i = 5$ for some consumer i (i.e., the consumer only consumes the best products) and $q_t < \phi_i$, the consumer would prefer not to consume if she knew the quality with certainty. If the reviewer has very little precision, the consumer would make the correct decision since $\mathbb{E}[q_t|r_t] \approx \mathbb{E}[q_t] = 0$ pretty much regardless of the review. On the other hand, if the reviewer is more precise, there is a greater probability that r_t will exceed the threshold needed for consumer i to purchase (and regret it). This does not contradict Proposition 2, however, because *on-average* (over many products), the consumer would always prefer a more precise review. Second, we point out that long-run average consumer utility depends only on $\hat{\sigma}$, and nothing else (e.g., frequency or size of the bribes). Therefore, to gauge consumer welfare effects it is sufficient to characterize the equilibrium choice of $\hat{\sigma}$ as a function of other parameters.

6 How do Bribes Impact Consumer Utility?

We now examine the impact of bribes on consumer utility. As a benchmark, we start by providing a simple characterization of the setting where the auditing technology is perfect and always catches bribes, so that all firms are truthful. Second, we derive conditions under which, even if firms could bribe, reviewers prefer to exert full influence by reporting truthfully. This provides the *first-best* amount of information transmission, and by Proposition 2, maximizes consumer utility. The comparative statics and associated policy recommendations arising from this characterization are particularly interesting. Finally, we show that the absence of bribes can sometimes lead to the *lowest possible* consumer utility, and that the introduction of bribes in these environments can restore some of that utility via implementing second-best information transmission.

6.1 Perfect Auditing Technology

Suppose that $\alpha = 1$, so all firms are truthful (i.e., $\theta = 1$). When this is the case, the reviewer derives utility solely from her influence. Thus, the reviewer participates in writing reviews only if her expected long-term utility from influence exceeds the cost of writing them. As such, all that matters is the relative difference between the reviewer's propensity for influence, β, and the cost of writing, C. When the reviewer does write a review, she maximizes her influence by reporting truthfully, which leads to maximum information transmission. Otherwise, no reviews are written and consumers can do no better than use their prior for their purchasing decisions. This is summarized in the next result.

Proposition 3. *For every cost C, there exists $\beta^*(C)$ such that:*

(a) If $\beta > \beta^*(C)$, the reviewer reports truthfully $(r_t = s_t)$ and consumer utility is maximized;

(b) If $\beta < \beta^*(C)$, the reviewer abstains $(r_t = \varnothing)$ and consumer utility is minimized.

Moreover, $\beta^*(C)$ is (strictly) increasing in the cost C.

Informally, the first-best outcome is always obtained with perfect auditing technology, assuming reviewers care enough about their influence and the cost of writing a review is not too high.

6.2 Bribe-Free Environments

Having defined the first-best in the previous section, we now derive conditions under which this first-best is achieved in the presence of imperfect auditing (i.e., $\alpha < 1$). Recall from Section 5.3 that *fair consumption* is the ex-ante expected consumption for a firm when the reviewer is the worst skill possible but guarantees an honest review. We then make the following definition:

Definition 5. The *consumption-influence ratio*, ψ, is the ratio of the fair consumption to the change in influence possible from mimicking, i.e., $\psi = (\bar{X}^* - \gamma)/(I_\infty(\sigma_H) - I_\infty(\sigma_L))$.

The consumption-influence ratio measures the relative value of fair consumption (discounted by the firm's outside option γ) to the influence a reviewer must relinquish in order to accept bribes. When this ratio is large (resp. small), either firms are willing to pay large (resp. small) bribes for just a fair (as opposed to preferential) treatment, or reviewers are willing to easily (resp. rarely) compromise influence for bribes. One of the main results of our paper is the following theorem, which shows that the property of being bribe-proof is not monotonic in the extent of the auditing technology.

Theorem 6. *Suppose the influence propensity satisfies $\beta > \psi$. Then there exists $0 < \underline{\alpha} < \bar{\alpha} < 1$ such that if $\alpha < \underline{\alpha}$ or $\alpha > \overline{\alpha}$, the setting is bribe-proof.*

We show next, by example, that the environment may not be bribeproof with an intermediate amount of auditing technology, $\underline{\alpha} < \alpha < \bar{\alpha}$, even when $\beta > \psi$:

Example 2 (Intermediate Auditing and Non-Monotonicity)**.** Let us revisit the setting of Example 1. We will first show that even when $\beta > \psi$, bribes can be possible when the auditing technology is intermediate (i.e., θ lies well in the interior of $[0, 1]$). This is because the reviewer can offer favorable biases to a subset of firms at a price that they both find agreeable. We then show that Theorem 6 in that when θ gets closer to 0 or 1, this type of biasing is either not profitable or not doable.

Unlike in Example 1, we will assume firms have an outside option of $\gamma = 0.37$. Then the fair consumption is given by $\bar{X}^* = \frac{1}{3} + \frac{1}{3} \cdot \mathbb{P}[r_L \geq 5/2] = \frac{1}{3} + \frac{1}{3} \cdot \Phi(-\sqrt{5}/2) \approx 0.378$, 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 consumption-influence ratio is equal to $\psi = 0.008/(I_\infty(\sigma_H) - I_\infty(\sigma_L))$.

When the high-skill reviewer mimics low-skill, she can accept bribes from 50% of the firms, 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:

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

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:

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

Therefore, the payoff to the firm from bribing, assuming the reviewer is high-skill, is $1/3 + 1/3(0.264) - 0.37 = 0.0513$. 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 bribeproof? Let us consider some $\beta = \psi + \nu$, for small $\nu > 0$. The loss of influence for the reviewer is $\beta \cdot (I_\infty(\sigma_H) - I_\infty(\sigma_L)) = 0.008 + \nu \cdot (I_\infty(\sigma_H) - I_\infty(\sigma_L))$, whereas the bribes she can receive are as large as 0.013. Therefore, there exists a bribing equilibrium when $\theta = 1/2$, even though $\beta > \psi$.

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 this β, as long as the auditing technology is good enough such that $\theta > 0.98$ (i.e., fewer than 2% of firms possibly bribe), these bribes are not sufficient to compensate for the damage in reviewer 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 maximal bribe would be slightly above $b^* = .008$, but this does not compensate the reviewer sufficiently for her loss in influence when $\beta > \psi$. Thus, both of these environments are bribeproof, as predicted by Theorem 6. $\square$

Theorem 6 (along with Example 2) describes a phase transition whereby the environment is bribe-proof when there is little auditing technology. Once the auditing technology increases beyond a certain threshold the environment may be susceptible to bribes and continues to be so in an intermediate region, and then transitions again to being bribe-proof as the auditing technology stops sufficiently many firms from bribing. The reasoning behind Theorem 6, as outlined in Example 2, is the following. Recall that the adjustment noise that the reviewer adds to the reviews of bribing firms is drawn from a normal distribution that is parametrized by the proportion of truthful firms θ. As $\alpha \to 0$, θ becomes small enough that it makes the bias received by a bribing firm not worth much. Consequently, the firm only offers a small bribe that is not enticing enough for the reviewer to give up some of her influence, and therefore she reports truthfully and the environment is bribe-proof.

At the other end, when $\alpha > \bar{\alpha}$, the fraction of bribing firms is low enough that it becomes unprofitable for the reviewer to trade off her influence with the bribes she collects. This makes her keep her influence by reporting truthfully, which again leads to a bribe-proof environment. It is in the intermediate region $[\underline{\alpha}, \overline{\alpha}]$ that the reviewer and the strategic firms are both better off in the presence of bribes: the reviewer can credibly convince consumers that she is a lower effective type and strategic firms find that the bias they receive in that case is worth more sizable bribes. This in turn make it profitable for the reviewer to trade influence for these larger bribes.

Theorem 6 has immediate policy implications. This is because interventions that aim to reduce the number of bribing firms can move the whole system from the bribe-proof environment $[0, \underline{\alpha})$ to the regime $[\underline{\alpha}, \overline{\alpha}]$ that admits bribes and reduces welfare, essentially accomplishing the opposite goal of what the intervention is designed for. Similarly, even when the proportion of truthful firms is high enough and the environment is bribe-proof, a small slip that takes that proportion below $\bar{\theta}$ is all that is required for the environment to revert back to being vulnerable to bribes. These non-monotone welfare conclusions, when varying the intensity of the auditing technology, generalize immediately to more arbitrary consumer distributions by the following result:

Proposition 4. *There exists β^*, $\bar{p}$, cutoffs $0 < \alpha_1 < \alpha_2 < \alpha_3 < \alpha_4 < 1$, and an outside option for the firm, γ, such that $\alpha \in (0, \alpha_1) \cup (\alpha_4, 1)$ is bribeproof, but consumer utility is strictly below first-best levels (i.e., there are bribes) when $\alpha \in (\alpha_2, \alpha_3)$ for all $p > \bar{p}$.*

The next result examines whether the environment is bribe-proof as a function of the difference between the possible reviewer skills. We see that the behavior is again not monotone: if reviewers have comparable skills, then there is little to be gained by the firm if the reviewer mimics down, since by doing so she resembles someone who does not look that different from her to begin with, and that makes bribing unattractive. On the other hand, if the skill gap between reviewers is too high, then there might be no bribe that will make it profitable for the reviewer to relinquish this much influence when she pretends to be a much worse type. This occurs for two reasons. First, the reviewer sacrifices the payoff she gets from exerting higher influence by accepting these bribes. Second, and more importantly, when the skill gap is large, firms are unwilling to pay bribes to a reviewer who has little influence on the decisions of the consumers. These together imply that consumer welfare is non-monotone in the difference in reviewer skills.

Proposition 5. *Fix some $\beta > 0$ and let $\Delta = \sigma_L - \sigma_H$ be the difference in the precisions of the most extreme types. Then first-best levels of consumer welfare are attained as $\lim_{\Delta \to 0}$ or $\lim_{\Delta \to \infty}$.*

The existence of bribing examples (i.e., Example 1) should convince the reader that when Δ is in an intermediate range, consumer welfare may be below first-best levels. Proposition 5 highlights that the property of being bribe-proof is more a matter of relative skill levels than absolute skill levels: whether bribes exist in the system depends on if the reviewer can successfully imitate a less informative reviewer, without losing her influence over consumer decisions altogether.

Finally, our last comparative static shows that the role that prior beliefs about the reviewer's skill have on welfare can be quite interesting. In particular, when the prior beliefs that consumers and firms have about the reviewer skill move in a direction that makes it more likely that the reviewer is low skill, then information transmission from the reviewer to consumers *increases* and welfare increases as well (holding fixed the true type of the reviewer). Recall that firms, like consumers, cannot distinguish between a reviewer who really is low-type and a reviewer who is the high-type but behaves as an effective low-type. The latter can help inflate the reviews of bribing firms but the former cannot. When the prior about the reviewer being low type increases, firms cease to offer bribes as the chances that the reviewer cannot help them increase, and this leads to first-best information transmission.

Proposition 6 (Prior Beliefs)**.** *Suppose the environment is bribe-proof and consider decreasing the prior probability of the high-type reviewer p (and thus increase the prior the reviewer is low-type). Then the environment is still bribeproof.*

We now turn our attention to a more subtle point about consumer welfare and prior beliefs. While shifting prior beliefs closer to the low-skill types increases welfare for a given reviewer type, we also increase the likelihood that the reviewer herself will be born with less skill, which decreases welfare. These two effects compete to give an ambiguous net effect on expected consumer welfare (where the expectation is over initial reviewer types): a higher likelihood of poor reviewers means fewer bribes and information distortion, but also just lesser information for the reviewer to transmit. We show in the following example, perhaps surprisingly, that reducing the expected skills of the reviewer can lead to an *improvement* in (average) consumer welfare because of its reduction in bribes:

Example 3 (Reviewer Skill and Welfare)**.** Let us revisit the setting of Example 1, except where we vary the probability p that the reviewer is high skill. The value of p is a proxy for the expected skill of the reviewer or transparency in the product's quality upon consumption. The maximal bribe offered by the firm is $b^* = p/25$, which implies the environment is bribeproof if $\beta \geq 10/3 \cdot p$, or $p \leq 3/10 \cdot \beta$. Consumer utility is proportional to $1.34p + 0.90(1 - p)$ when the environment is bribeproof, and equal to 0.90 when the environment is not. In other words, consumer utility is equivalent to $0.90 + 0.44 \cdot p \cdot \mathbf{1}_{p \leq 3/10 \cdot \beta}$. For different values of β, we plot ex-ante expected consumer utility in Figure 1 as a function of p. Note in particular that it is non-monotone in p, and is maximized at $p = \min\{3/10 \cdot \beta, 1\}$. The reason for this non-monotonicity lies in the fact that increasing the (expected) skill of the reviewer improves information but also increases incentives for bribes. Holding the latter constant, the former leads to an increasing in consumer utility, per Proposition 2. However, the latter effect can dominate whereby the introduction of bribes into the system when skills increase leads to a reduction in information transmission. $\square$

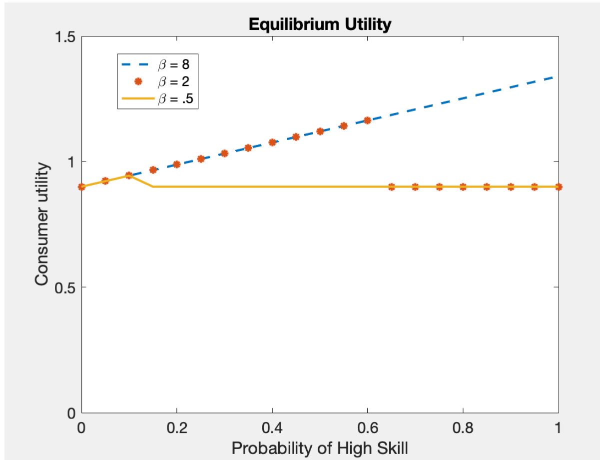

Figure 1. Equilibrium Utility for Varying Ex-Ante Skill.

Example 3 illuminates an important counterintuitive finding: increasing the quality of information and efficacy of reviews may not necessarily lead to better outcomes for consumers. Because firms can also leverage the improvement in information for their own gain, it is possible reviewers and firms extract greater surplus through bribes at the cost of hurting consumers.

6.3 Second-Best Outcomes and Bribe Funding

We now show that the presence of bribes can sometimes have a positive impact on consumer utility. We assume throughout that reviewers do not value their own influence ($\beta = 0$) but instead derive benefits entirely from bribes, and yet reviews are costly to write ($C > 0$). In this case, it is straightforward to see that first-best information transmission is not possible, simply because any (strategic) reviewer of high-skill will mimic-down to $\hat{\omega} = L$: the high-skill reviewer can only obtain payoffs from bribes, which only occurs if the reviewer can bias bribing firms. Similarly, since the cost of writing a review is positive, she may abstain altogether unless there is compensation from bribes. This is the worst case for consumer utility, which leads into the following definition:

Definition 6 (Second-Best)**.** An environment has *second-best* information transmission if both reviewer types elect to play effective type $\hat{\omega} = L$ instead of abstaining.

We say this is second-best because of all the environments where $\beta = 0$, this outcome maximizes consumer utility. We present our main result in this setting for second-best outcomes. Interestingly, the worst outcome for consumer utility occurs when auditing technology is most stringent. With lesser auditing however, utility increases and we eventually get the (second-) best outcome. This is summarized in the following result.

Proposition 7. *There exists some prior $\bar{p} \in (0,1)$, $C \in (\underline{C}, \bar{C})$, and $0 < \underline{\alpha} < \overline{\alpha} < 1$ such that for all $p > \bar{p}$ and $C \in (\underline{C}, \bar{C})$: (i) if $\alpha > \overline{\alpha}$, the reviewer abstains and consumer utility is minimized, (ii) if $\alpha < \underline{\alpha}$, the second-best outcome is obtained.*

Recall from Proposition 3 that when $C > 0$ and $\beta = 0$, the reviewer abstains from writing a review and consumer utility is minimized. However, the introduction of bribing firms can entice reviewers to participate. The private market funds the review writing (as a form of advertisement), and while the incentives for writing these reviews are perverse and lead to review manipulation, they nevertheless guarantee an improvement in utility compared to the alternative of having no reviews altogether.

Finally, we consider an example with more reviewer types that generalizes the findings in Proposition 7. In particular, we once again see that second-best outcomes may occur due to bribe funding: the reviewer may mimic only one type down or even truthtell, which is funded entirely by the firms.

Example 4 (Two-Sided Reputation)**.** Consider $\Omega = \{L, M, H\}$ where the high-skill type is highly likely but $\beta = 0$, so the influence of the reviewer has no inherent payoff to her. Let us assume the same ϕ distribution as in Example 1, $\theta = 1/2$ and $\sigma_H = 1$, but now we vary $1 \leq \sigma_M \leq \sigma_L \leq 2$. Similar to before, the high-skill reviewer will mimic down to either low or middle, whichever can solicit the larger bribe from the firm. Thus, we simply calculate the maximal bribe when mimicking some type $\hat{\sigma}$ from Example 1:

$$\begin{aligned}
b^*(\hat{\sigma}) &= \frac{1}{3}\mathbb{P}[|(q_t + \varepsilon_t) - (1+\hat{\sigma}_L^2)/2|] \\
&= \frac{1}{3}\Bigg(1 - \frac{1}{2\pi\sqrt{2(\hat{\sigma}^2-1)}}\int_{-\infty}^{0}\int_{\kappa}^{-\kappa} \exp\left(-\frac{(\kappa + (1+\hat{\sigma}^2)/2)^2}{4}\right)\exp\left(-\frac{\varepsilon^2}{2(\hat{\sigma}^2-1)}\right)\, d\varepsilon\, d\kappa \\
&\quad - \frac{1}{2\pi\sqrt{2(\hat{\sigma}^2-1)}}\int_{0}^{\infty}\int_{-\kappa}^{\kappa} \exp\left(-\frac{(\kappa + (1+\hat{\sigma}^2)/2)^2}{4}\right)\exp\left(-\frac{\varepsilon^2}{2(\hat{\sigma}-1)}\right)\, d\varepsilon\, d\kappa\Bigg)
\end{aligned}$$

In Figure 2, we see b^* as a function of the effective type $\hat{\sigma}$. This provides a tight characterization of the information transmission in equilibrium: the high-type mimics $\arg\max_{\hat{\sigma}\in\{\sigma_M,\sigma_L\}}$. For example, in terms of $\hat{\sigma}^*$, the solution to $\arg\max_{\hat{\sigma}} b^*(\hat{\sigma})$, we have:

(i) If $\hat{\sigma}_M < \hat{\sigma}_L < \hat{\sigma}^*$, then the high-type reviewer and the middle-type reviewer always mimic the low type, or both abstain.

(ii) If $\hat{\sigma}^* < \hat{\sigma}_M < \hat{\sigma}_L$, then the high-type reviewer mimics the middle type and the middle type reports truthfully (i.e., second-best), or both abstain.

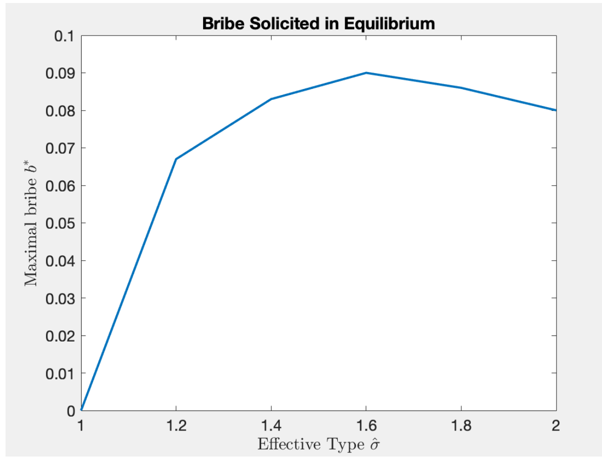

Figure 2. Maximal Bribes.

Whether the reviewer abstains in (i)/(ii), or mimics as described, depends on if b^* exceeds C. In this case, we get bribe funding, and in the special case of (ii), this bribe funding obtains the second-best outcome, even when it is possible to mimic lower precision types. Note also the similarity to our conclusions in Proposition 6 and Example 3: as skills atrophy, information decreases, but the incentives for bribes also change, often in a way that can benefit consumer utility. Here, this comes in the form of maintaining a stronger reputation with the consumers, thereby improving the efficacy of the bias, and soliciting larger bribes. This highlights an important feature: the presence of (larger) bribes in the system does not always translate into a decrease in consumer utility.

The key tradeoff for the reviewer, as seen in Figure 2, is the following. When mimicking a low type, bribing firms receive more bias in that, considering $r_t = (q_t + \varepsilon_t) + |\varepsilon_t'|$, a larger amount of their review variation comes from the (good) bias $|\varepsilon_t'|$, which makes bribing more attractive. On the other hand, we know from Theorem 4 that when the reviewer mimics worse types, consumers more heavily discount her review, thereby making this additional bias less advantageous. While the reviewer doesn't care about inherent influence, she *does care about influence* insofar as it may help her get larger bribes. $\square$

Example 4 provides two related insights. The first is that bribe funding can spawn review writing, and the degree of review manipulation depends on the prior distribution over reviewer skills and the consumer base in nuanced (and non-monotone) ways. Second, the reviewer may choose to possess higher influence, even when she does not care about influence directly. Rather, the desire for influence can arise endogenously from the two-sided interaction and the reputation balance she needs to maintain between the firms and the consumers. In the extreme case, the reviewer may even truthtell by accepting bribes while only "pretending" to bias, as to capture the largest possible influence with the consumers.

7 Final Remarks

The recent proliferation of review manipulation shows that the problem is multi-faceted and can take on multiple forms. In this paper, we focused on the case where an influential reviewer or expert can misrepresent her reviews in order to favor specific firms. We provided conditions under which the reviewer can repeatedly do this without getting detected, and analyzed the impact of such practice on consumer utility. Interestingly, we show that an increase in the number of non-bribing firms might lead to the market becoming more vulnerable to bribes and review distortions, and is therefore an aspect that perhaps should be taken into account when regulators aim to crack down on bribing firms. Another possibility that could dilute the effect of bribes is recognizing that at the center of the reviewer's decision problem is a tradeoff where she tries to balance her loss of influence (that comes from providing less accurate reviews) with accepting side payments from firms, and implicit in this is the observation that an increase in how reviewers value influence can offset the payoffs obtained from bribes and lead to bribe-proof environments where reviewers always reports truthfully. The flip side is that when reviewers have little or no inherent value for influence, their reviews may become extremely noisy or they may stop writing reviews altogether, leading to a decline in consumer utility. In this case, bribes provide a market mechanism that supports reviewers and funds the process of review writing, and while the reviews will unavoidably be biased in this case, they still transmit useful information to consumers and lead to an improvement in welfare.

As mentioned, review manipulation can manifest in other ways. For example, there are large-scale manipulation attempts on platforms like Amazon, where product ratings are aggregated from many reviews, and where the seller can source out these reviews to multiple agents who would provide inflated product ratings in exchange for payments. This problem can possibly be addressed by designing aggregation systems that take into account factors that suggest whether a review is authentic or not, in similar ways to how some websites like fakespot or reviewmeta flag some reviews as suspicious.

Finally, we remark that our work is applicable beyond the setup of online reviews. The testing, inspection, and certification market 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 share several attributes with the model developed in this paper, where again an expert or an intermediary evaluates members of one side of the market, and that evaluation provides information that goes into the decisions of market participants and impacts the overall efficiency of the marketplace. Despite the importance of these interactions and their impact on all parties involved, their dynamic nature have not been analyzed or perfectly understood, and we hope that the model in this paper serves as a first step in this direction.

A Proofs

Proof of Theorem 1. For part (a), 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 short-run (myopic) players. Short-run players may observe the the entire history of play (except for the bribes) since the public signal q_t eliminates the private information received by the reviewer about q_t. Because bribes are given by $b_t(H_t)$, which is a function only of this public history, the payoffs are determined by nature given H_t.

For part (b), consider firm t who chooses $b_t(H_t)$ conditional on the public history H_t. Note the information set of the firm is uniquely determined by H_t, so can only choose (mixed) strategies dependent on this. The firm takes as given the public belief of the reviewer's type π_t, and acts myopically (choosing b_t) by maximizing $\mathbb{E}_{\pi_t}[U_t]$, where the action of the reviewer and consumers is given by the reviewer-consumer equilibrium identified in part (a). This is a static game of incomplete information between the firm and the reviewer, who receives both a current payoff and the continuation payoff from the reviewer-consumer equilibrium. Existence is implied immediately by the existence of a Bayesian Nash equilibrium (see Section 6.4 in Fudenberg and Tirole (1991)). □

Proof of Theorem 2. Since $\beta > 0$, by Fudenberg and Levine (1992) there always exists an equilibrium where the reviewer receives at least $\beta \cdot I_\infty(\omega) > 0$ (i.e., her Stackelberg payoff under Assumption 2). Consider the the public belief $\tilde{\pi}_t$ that the reviewer is a truthful type (of any precision). It is clear that $\tilde{\pi}_t$ is a bounded martingale, so converges almost surely as $t \to \infty$. Let us denote by $\mathbf{r}_t^*(\omega)$ the (distribution) of reviews by an honest reviewer of type ω at time t (which is given by $\mathcal{N}(0, 1+\sigma_\omega^2)$ and $\text{cov}(\mathbf{r}_t^*(\omega), q_t) = 1$ as $t \to \infty$). We know that either the belief $\tilde{\pi}_t$ converges to zero or the reviewer mimics the truthful type in the sense that, $\limsup_{t\to\infty} \tilde{\pi}_t \cdot ||\mathbf{r}_t - \mathbf{r}_t^*(\omega)||_\infty = 0$ for some $\omega \in \{H, L\}$ (see Cripps et al. (2004)). Note for fixed $\delta < 1$, there always exists a cutoff $\underline{\tilde{\pi}}_t$ such that for $\tilde{\pi}_t < \underline{\tilde{\pi}}_t$, there is an equilibrium where the reviewer abstains ($r_{t'} = \varnothing$) for all $t \geq t'$ (and the consumer ignores any review written off the equilibrium path). We call this equilibrium *babbling-trigger*. In a robust equilibrium (where the equilibrium is assumed to be risk-dominant), the consumer must necessarily babble-trigger, as this equilibrium obtains a payoff approaching 0 for a patient reviewer, which is the lowest payoff possible. For the reviewer to obtain a payoff of at least $\beta \cdot I_\infty(\omega)$ in equilibrium, she must not be babble-triggered for at least:

$$T^* = \frac{\log\left(1 - \frac{\beta \cdot I_\infty(\omega)}{2}\right)}{\log(\delta)} - 1$$

periods. Since $T^* \to \infty$ as $\delta \to 1$, we know for any fixed t, it cannot hold that $\liminf_{\delta\to 1} \tilde{\pi}_t \to 0$ in equilibrium, otherwise we have arrived at a contradiction of Fudenberg and Levine (1992) that the reviewer obtains at least $\beta \cdot I_\infty(\omega)$. Therefore, as $\delta \to 1$, it must be the case that $\limsup_{t\to\infty} ||\mathbf{r}_t - \mathbf{r}_t^*(\hat{\omega})||_\infty = 0$ for some type $\hat{\omega}$ from the set of commitment types in equilibrium.

Since $\mathbf{r}_t^*(\hat{\omega})$ is identifiable for every precision type $\hat{\omega}$, it is clear the public belief π_t of the reviewer's mimicking type $\hat{\omega}$ will converge to the true mimicked type almost surely by LLN. Since as $\delta \to 1$ (and $t \to \infty$), we know $||\mathbf{r}_t - \mathbf{r}_t^*(\hat{\omega})||_\infty \to 0$, consumers can compute $\mathbb{E}[q_t|r_t]$:

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

And since consumers are myopic, they employ a cutoff strategy whereby if $\phi_i \le \mathbb{E}[q_t|r_t] = \frac{r_t}{1+\sigma_{\hat{\omega}}^2}$, which is unique except on a set of measure zero.

Let S_t be a strategy for the reviewer (at time t) which maps bribes, signals, and the public history, (b_t, s_t, H_t), into reviews, r_t. For a bribing firm t, we let $\tilde{S}_t(H_t)$ be its mixed strategy over bribes, as a function of the public history H_t. Note the strategy employed by the reviewer is in fact:

$$S_t^* = \theta \cdot S_t(0, s_t, H_t) + (1-\theta) \cdot \int_{b_t \in \tilde{S}_t(H_t)} S_t(b_t, s_t, H_t)\, db_t$$

By our previous observation, it must be the true that $\lim_{t\to\infty} ||S_t^* - \mathbf{r}_t^*(\hat{\omega})||_\infty \to 0$ for some $\hat{\omega}$. Note because the consumer can also observe H_t, it must be true that for all H_t, $\lim_{t\to\infty} ||S_t^*(H_t) - \mathbf{r}_t^*(\hat{\omega})||_\infty \to 0$. But, recall that $\mathbf{r}_t^*(\hat{\omega})$ depends only on s_t and not b_t or H_t. Therefore, in the limit $t \to \infty$, S_t^* cannot depend on H_t in the sense that $||S_t^*(s_t, H_t) - S_t^*(s_t, H_t')||_\infty \to 0$.

Consider any equilibrium where $S_t(0, s_t, H_t)$ admits different $\mathbb{E}_{S_t(0,s_t,H_t)}[\phi^{-1}(r_t/(1+\sigma_{\hat{\omega}}^2))]$ for different histories, for infinitely many t (i.e., a sequence $\tau_1, \tau_2, \ldots, \tau_k, \ldots$). Without loss of generality, suppose that

$$\mathbb{E}_{S_t(0,s_t,H_t)}\left[\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)\right] < \mathbb{E}_{S_t(0,s_t,H_t')}\left[\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)\right]$$

at some time τ_k. Instead, suppose the reviewer sets $S_t(0, s_t, H_t') \leftarrow S_t(0, s_t, H_t)$ and $S_t(b_t, s_t, H_t') \leftarrow S_t(b_t, s_t, H_t)$ at all times τ_k. It is clear that this can still be supported in equilibrium because

$$\mathbb{E}_{S_t^*}\left[\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)\right] = \mathbb{E}_{(S^*)_t'}\left[\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)\right] = \mathbb{E}_{r_t^*}\left[\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)\right]$$

Moreover, it must be the case that for bribing firms:

$$\mathbb{E}_{b_t}\left[\mathbb{E}_{S_t(b_t,s_t,H_t)}\left[\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)\right]\right] > \mathbb{E}_{b_t}\left[\mathbb{E}_{S_t(b_t,s_t,H_t')}\left[\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)\right]\right]$$

which implies that there exists an equilibrium where $b_\tau' > b_\tau$ for all such τ, so such an equilibrium is not maximal. This implies that $S_t(0, s_t, H_t)$ cannot depend on H_t as $t \to \infty$, which immediately shows that $\int_{b_t \in \tilde{S}(H_t)} S_t(b_t, s_t, H_t)\, db_t$ may not depend on H_t either.

In a robust equilibrium, every firm acting at time t who enters chooses $\min_{b_t}(b_t \in \arg\max_{b'_t} \mathbb{E}_{H_t,\pi_t}[U_t])$ because she breaks indifferences toward lower bribes because of the infinitesimal proportional cost. Therefore the firm always plays a pure-strategy, and $\int_{b_t \in \tilde{S}(H_t)} S_t(b_t, s_t, H_t)\, db_t$ immediately reduces to $S_t(b_t(H_t), s_t, H_t)$. In the limit $t \to \infty$, because $\mathbf{r}_t^*(\hat{\omega})$ cannot depend on b_t either, it must necessarily be the case that S_t does not depend on H_t; that is, $||S_t(b_t(H_t), s_t, H_t) - S_t(b_t, s_t)||_\infty \to 0$, for some $S_t(b_t, s_t)$. 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 (given by S_t), and all strategic firms are identical, the outcome of review manipulation is the same for all strategic firms. We can thus write $S_t(s_t)$, and note that in a maximal equilibrium, b_t is given by the supremum over all bribes supported by some equilibrium $S_t(s_t)$ as $t \to \infty$, denoted by b^* (not such a maximal bribe cannot depend on t, because the reviewer's strategy does not depend on H_t, including t). The unique equilibrium outcome for the firms is given by strategic firms bribing b^* and truthful firms bribing 0. If $\lim_{t\to\infty} ||S_t^*(s_t) - \mathbf{r}_t(\hat{\omega})||_\infty \to 0$ for some $\hat{\omega}$, then $S_t^*(s_t)$ must converge in distribution to a normal distribution $\mathcal{N}(0, \sigma_{\hat{\omega}}^2 + 1)$ with cov$(S_t^*(s_t), q_t) = 1$. Since this distribution is entirely determined by its covariance matrix, $S_t^*(s_t)$ converges (in distribution) to a unique distribution.

Recall that we can write $S^*(s_t) = \theta \cdot S(0, s_t) + (1-\theta) \cdot S(b^*, s_t)$, and by the arguments in the above paragraph, we know that $S_t^*(s_t)$ is constrained (in the limit) to have a normal distribution $\mathcal{N}(0, \sigma_{\hat{\omega}}^2 + 1)$ with cov$(S^*(s_t), q_t) = 1$. Consider some arbitrary $S^*(s_t)$ which satisfies these distributional constraints. Note that $S^*(s_t) - s_t$ is normally distributed and uncorrelated with q_t, so in particular, are independent of q_t. For any strategy $S^*(s_t) - s_t$ to be independent of q_t as $t \to \infty$, since q_t is not observed ex-ante by the reviewer, it must be the case that $S^*(s_t) - s_t$ is independent of s_t. Therefore, $S^*(s_t) - s_t$ must have the distribution $\mathcal{N}(0, \sigma_{\hat{\omega}}^2 - \sigma_\omega^2)$ and be orthogonal to s_t. Therefore, we can write $S^*(s_t) = s_t + \varepsilon'_t$, where ε'_t is orthogonal to s_t and has this distribution. Rewritten:

$$S^*(s_t) = s_t + \theta \cdot (S(0, s_t) - s_t) + (1-\theta) \cdot (S(b^*, s_t) - s_t)$$

It is easy to see that $S(0, s_t) - s_t$ must be independent of s_t in any maximal bribe, for the same reason S_t must be independent of H_t. Since $S^*(s_t) - s_t$ is independent of s_t, this implies that $S(b^*, s_t) - s_t$ is independent of s_t (again, because all distributions are Gaussian). Consider some $\tilde{S}^*$ with $\tilde{S}(0, s_t) = s_t + \varepsilon'_t$ with $\varepsilon'_t \sim \mathcal{N}_\theta(0, \sigma_{\hat{\omega}}^2 - \sigma_\omega^2)$ and $\tilde{S}(b^*, s_t) = s_t + \varepsilon'_t$ with $\varepsilon'_t \sim 1 - \mathcal{N}_\theta(0, \sigma_{\hat{\omega}}^2 - \sigma_\omega^2)$ as given in Theorem 5. It is easy to check that $\tilde{S}^*(s_t)$ satisfies the distributional constraints. We claim that $\tilde{S}(b^*, s_t) - s_t$ first-order stochastically dominates all other distributions subject to the distributional constraints: (i) $S(b^*, s_t) - s_t$ is orthogonal to s_t; (ii) $S^*(s_t) - s_t \sim \mathcal{N}(0, \sigma_{\hat{\omega}}^2 - \sigma_\omega^2)$; and therefore, generates the maximal bribe b^*. Consider some $\hat{S}(b^*, s_t)$ satisfying these conditions. Note that for any $a < (\sigma_{\hat{\omega}}^2 - \sigma_\omega^2) \cdot \Phi^{-1}(1-\theta)$, then $\mathbb{P}[\tilde{S}(b^*, s_t) - s_t \geq a] = 1$, and so $\mathbb{P}[\tilde{S}(b^*, s_t) - s_t \geq$

$a] \geq \mathbb{P}[\hat{S}(b^*, s_t) - s_t \geq a]$ trivially. For any $a \geq \sqrt{\sigma_{\hat{\omega}}^2 - \sigma_{\omega}^2} \cdot \Phi^{-1}(1-\theta)$:

$$\begin{aligned}
\mathbb{P}[\tilde{S}(b^*, s_t) - s_t \geq a] &= \frac{1 - \Phi\left(\frac{a}{\sqrt{\sigma_{\hat{\omega}}^2 - \sigma_{\omega}^2}}\right)}{1-\theta} \\
&= \frac{1}{1-\theta}\mathbb{P}[\hat{S}^*(s_t) - s_t \geq a] \\
&= \frac{1}{1-\theta}\left(\theta \cdot \mathbb{P}[\hat{S}^*(s_t) - s_t \geq a | b = 0] + (1-\theta) \cdot \mathbb{P}[\hat{S}^*(s_t) - s_t \geq a | b = b^*]\right) \\
&= \frac{\theta}{1-\theta}\mathbb{P}[\hat{S}^*(s_t) - s_t \geq a | b = 0] + \mathbb{P}[\hat{S}^*(s_t) - s_t \geq a | b = b^*] \\
&\geq \mathbb{P}[\hat{S}^*(s_t) - s_t \geq a | b = b^*] \\
&= \mathbb{P}[\hat{S}^*(b^*, s_t) - s_t \geq a]
\end{aligned}$$

which shows that $\tilde{S}(b^*, s_t) - s_t$ FOSD $\hat{S}(b^*, s_t) - s_t$, so the expected review under $\tilde{S}(b^*, s_t)$ is always strictly higher than under $\hat{S}(b^*, s_t)$, and so is expected consumption.

It is also to easy to see that in every equilibrium the firm must enter under Assumption 2. Otherwise, the reviewer is forced to abstain for $(1-\theta) > 0$ fraction of firms, which implies she gets discounted average payoff of 0 as $\delta \to 1$ (either from babbling-trigger or from always abstaining), which is a contradiction, because the reviewer must obtain at least her Stackelberg payoff.

Finally, we need only show that there is a unique $\hat{\omega}$ the high-type reviewer mimics in the maximal equilibrium, under generic conditions (since the low-type reviewer's strategy is fixed, by Theorem 3). Recall the reviewer's (realized) payoff is given by $V(\hat{\omega}) = (1-\delta)\sum_{t=1}^{\infty} \delta^t(\beta \cdot I_t(\hat{\omega}) + b_t(\hat{\omega}))$. Suppose that for $\hat{\omega}_1 < \hat{\omega}_2$, we have $V(\hat{\omega}_1) = V(\hat{\omega}_2)$ as $\delta \to 1$. Because $(1-\delta)\mathbb{E}\left[\sum_{t=1}^{\infty} \delta^t I_t(\hat{\omega}_1)\right] > (1-\delta)\mathbb{E}\left[\sum_{t=1}^{\infty} \delta^t I_t(\hat{\omega}_2)\right]$ it must necessarily be the case that $(1-\delta)\mathbb{E}\left[\sum_{t=1}^{\infty} \delta^t b_t^1\right] < (1-\delta)\mathbb{E}\left[\sum_{t=1}^{\infty} \delta^t b_t^2\right]$, where b_t^1 and b_t^2 are the bribes given under $\hat{\omega}_1, \hat{\omega}_2$, respectively. Thus, for any sufficiently small $\epsilon > 0$, setting $\beta \leftarrow \beta \pm \epsilon$ breaks the indifference and does not affect the payoffs of the firm or consumer, so b_t is still supported in equilibrium. Thus, for generic β, there is a unique choice for $\hat{\omega}$ for the high-type reviewer. □

Proof of Theorem 3. By Theorem 2, as $t \to \infty$ we know the reviewer mimics a commitment type $\hat{\omega}$. Thus, it just remains to show that the reviewer must mimic-down. Consider $\sigma_{\hat{\omega}} < \sigma_{\omega}$. Let $f_t : s_t \mapsto r_t$ be any stochastic function (possibly time-varying) mapping the reviewer's signals into reviews. 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. Thus, the Fisher information from s_t, $\mathcal{I}_{s_t}(q_t)$, and the Fisher information from r_t, $\mathcal{I}_{r_t}(q_t)$, satisfies $\mathcal{I}_{s_t} \geq \mathcal{I}_{r_t}$ by the chain rule. It is also easy to see that $\mathcal{I}_{\hat{s}_t} > \mathcal{I}_{s_t}$ if $\sigma_{\hat{\omega}} < \sigma_{\omega}$. Thus, it is impossible that $\lim_{t\to\infty} ||\mathbf{r}_t(\omega) - \mathbf{r}_t^*(\hat{\omega})||_\infty \to 0$, which is a contradiction. □

Proof of Theorem 4 and 5. These follow immediately from the arguments in Theorem 2. □

Proof of Proposition 1. By Theorem 4, we know that $\pi_t \overset{a.s.}{\to} \mathbf{1}_{\hat{\sigma}}$, and therefore $\mathbb{E}[q_\tau | r_\tau, \pi_\tau] \overset{a.s.}{\to} \mathbb{E}[q_\tau | r_\tau, \mathbf{1}_{\hat{\sigma}}]$, so by Theorem 3, we know that these are eventually i.i.d. random variables. By

Kolmogorov's strong law (see Sen and Singer (1994) Theorem 2.3.10), we know that:

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

where the final equality can be seen by simply integrating out q_τ. Recall we have that:

$$I_\infty(\hat{\omega})=\frac{1}{\sqrt{2\pi(1+\hat{\sigma}^2)}}\int_{-\infty}^{\infty}\left(\phi^{-1}\left(\frac{r_t}{1+\sigma_{\hat{\omega}}^2}\right)-\phi^{-1}(0)\right)^2\cdot\exp\left(-\frac{r_t^2}{2(1+\sigma_{\hat{\omega}}^2)}\right)\,dr_t$$

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

$$I_\infty(\hat{\omega})=\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\left(\phi^{-1}\left(\frac{\alpha}{\sqrt{1+\sigma_{\hat{\omega}}^2}}\right)-\phi^{-1}(0)\right)^2\cdot\exp(-\alpha^2/2)\;d\alpha$$

Note that $\phi^{-1}(\cdot)$ is an increasing function, and $\left|\alpha/\sqrt{1+\sigma_{\hat{\omega}}^2}\right|$ is decreasing in $\sigma_{\hat{\omega}}$, so the integrand is decreasing in $\sigma_{\hat{\omega}}$ pointwise for all α. Thus $I_\infty(\hat{\omega})$ is decreasing in $\sigma_{\hat{\omega}}$. $\square$

Proof of Proposition 2. As in Proposition 1, we can write:

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

and via Theorem 3 and 4, we know these are eventually i.i.d. random variables, which implies by Kolmogorov's strong law:

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

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

$$CU_i=\frac{1}{2\hat{\sigma}\pi}\int_{-\infty}^{\infty}\int_{(1+\hat{\sigma}^2)\phi_i}^{\infty}(q-\phi_i)\exp(-q^2/2)\exp\left(-\frac{(r-q)^2}{2\hat{\sigma}^2}\right)\;dr\;dq$$

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

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

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

$$CU_i = \frac{\sqrt{1+\hat{\sigma}^2}}{\sqrt{2\pi}(1+\hat{\sigma}^2)^{3/2}} \int_{\sqrt{1+\hat{\sigma}^2}\phi_i}^{\infty} \left(\sqrt{1+\hat{\sigma}^2}\alpha - (1+\hat{\sigma}^2)\phi_i\right) \exp\left(-\alpha^2/2\right)\, d\alpha$$
$$= \frac{1}{\sqrt{2\pi(1+\hat{\sigma}^2)}} \int_{\sqrt{1+\hat{\sigma}^2}\phi_i}^{\infty} \left(\alpha - \sqrt{1+\hat{\sigma}^2}\phi_i\right) \exp(-\alpha^2/2)\, d\alpha$$

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

$$CU_i = \frac{1}{\sqrt{2\pi(1+\hat{\sigma}^2)}} \int_0^{\infty} \kappa \exp\left(-\frac{(\kappa+\sqrt{1+\hat{\sigma}^2}\phi_i)^2}{2}\right)\, d\kappa$$
$$= \frac{1}{\sqrt{2\pi(\hat{\sigma}^2+1)}} \exp\left(-\frac{\phi_i^2(1+\hat{\sigma}^2)}{2}\right) - \frac{1}{2}\phi_i \text{erfc}\left(\frac{\phi_i\sqrt{1+\hat{\sigma}^2}}{\sqrt{2}}\right)$$

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

$$CU_i = \frac{1}{\sqrt{2\pi}\zeta} \exp\left(-\frac{\phi_i^2\zeta^2}{2}\right) - \frac{1}{2}\phi_i \text{erfc}\left(\frac{\phi_i\zeta}{\sqrt{2}}\right)$$

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, CU_i is decreasing in ζ for all i, which implies it is decreasing in $\hat{\sigma}$. Hence, the average consumer utility is also decreasing in $\hat{\sigma}$. $\square$

Proof of Proposition 3. When firms are truthful, the payoff to the reviewer is just given by her influence $\beta \cdot I_\infty(\hat{\sigma})$, less the cost of writing a review C, given she mimics type $\hat{\sigma}$. By Proposition 1, the reviewer's influence is maximized when $\hat{\sigma} = \sigma$ (given Theorem 3 which says the reviewer may only mimic down), so if the reviewer writes a review, she reports honestly $r_t = s_t$ to maximize her long-run (discounted) payoff. The reviewer abstains if on the other hand $\beta \cdot I_\infty(\sigma) < C$; thus, $\beta^*(C) = C/I_\infty(\sigma)$, which is increasing in C. Finally, by Proposition 2, and by Theorem 3 the reviewer can only mimic down, consumer utility is maximized when the reviewer is truthful and is minimized when the reviewer abstains, because this is identical to the $\hat{\sigma} \to \infty$ case. $\square$

Proof of Theorem 6. The firm's bribe is bounded above by $b^* = 1-\gamma$, so the expected bribe received by the reviewer is at most $(1-\theta)\cdot(1-\gamma)$. Because $\beta > 0$ and $I_\infty(\sigma_H) > I_\infty(\sigma_L)$, 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 to mimic her true type than mimic down and receive 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 $(\bar{X}^* - \gamma)$, as per Theorem 5 she receives review $r_t \sim q_t + \epsilon_t + \varepsilon_t'$ where $\varepsilon_t' \sim \lim_{\theta\to 0} 1 - \mathcal{N}_\theta(0, \sigma_L^2 - \sigma_H^2) = \mathcal{N}(0, \sigma_L^2 - \sigma_H^2)$, 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 she does not bribe she receives $r_t \sim q_t + \epsilon_t + \varepsilon_t'$ where $\varepsilon_t' \sim \lim_{\theta\to 0} \mathcal{N}_\theta(0, \sigma_L^2 - \sigma_H^2) = Dirac(-\infty)$, 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 0. The difference in these gives the maximal bribe of $b^* = \bar{X}^* - \gamma$, whereas the reviewer loses an (average) influence payoff of:

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

Because b^* is continuous in θ, and as $\theta \to 1$ the environment is bribeproof because the loss in influence payoff (strictly) exceeds the maximal bribe, there must exist $\bar{\theta}$ such that for all $\theta > \bar{\theta}$ this property still holds, so the environment remains bribeproof. By continuity, strict increasing, and the boundary conditions on $\alpha(\theta)$, applying the intermediate value theorem we can find lower and upper bounds on α such that the environments are bribeproof as well. □

Proof of Proposition 4. 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 $\zeta(\theta)/\theta$. Therefore, the maximal (average) bribe is given by $\bar{b}^*(\theta) = (1-\theta)\left(p\kappa(\theta) + \min\{-p\zeta(\theta)/\theta, (1-p)\bar{X}^* - \gamma\}\right)$ (see Proposition 6). 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

$\lim_{\theta\to 0}\Phi^{-1}(\theta) = -\infty$, we see $\lim_{\theta\to 0}\eta(\theta) = 0$. Therefore, fixing some small $\gamma > 0$, we obtain on an open interval $\theta \in (0, \underline{\theta})$ where $\partial\bar{b}^*(\theta)/\partial\theta > 0$ 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 $\bar{b}^*(\theta^*) > \lim_{\theta\to 0}\bar{b}^*(\theta)$. Choosing β^* such that:

$$\frac{\lim_{\theta\to 0}\bar{b}^*(\theta)}{I_\infty(\sigma_H) - I_\infty(\sigma_L)} < \beta^* < \frac{\bar{b}^*(\theta^*)}{I_\infty(\sigma_H) - I_\infty(\sigma_L)}$$

implies there is a region $(0, \theta_1)$ that is bribeproof, whereas $(\theta^* - \epsilon, \theta^* + \epsilon)$ experiences bribes for some small ϵ. Simultaneously, by the same logic as in Theorem 6, one can always find θ_4 such that with $\theta \in (\theta_4, 1)$ for this particular β^* that admits a bribeproof environment. Noting the continuity, strict increasing, and boundary conditions of $\alpha(\theta)$ as in Theorem 6 then proves the claim. $\square$

Proof of Proposition 5. It is easy to verify that $\bar{CU}_\infty(\hat{\sigma})$ from Proposition 2 is continuous in $\hat{\sigma}$, so as $\Delta \to 0$ we know that $\bar{CU}_\infty(\sigma_L) \to \bar{CU}_\infty(\sigma_H)$, and thus the first-best consumer utility is attained. When the reviewer mimics low-type, we know the payoff from a bribe is given by:

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

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:

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

which implies the maximal bribe satisfies $b^* \to 0$. Since $\beta > 0$, the high-type reviewer mimics her own type as $\sigma_L \to \infty$, which makes the environment bribeproof and therefore attains first-best consumer utility. $\square$

Proof of Proposition 6. Note the value of a bribe is given given by the difference between the firm entering and bribing and her next best option (either entering and not bribing, or staying out). Let us denote by $\underline{X}$ and $\overline{X}$ as the 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, she receives $p\overline{X} + (1-p)\bar{X}^*$, whereas if the firm enters and does not bribe, she receives $p\underline{X} + (1-p)\bar{X}^*$, and if she stays out she 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 bribeproof if $p\overline{X} + \min\{-p\underline{X}, (1-p)\bar{X}^* - \gamma\} < \beta(I_\infty(\sigma_H) - I_\infty(\sigma_L))$ and otherwise admits bribes (see Theorem 2). Since the right-hand side of the inequality is constant with respect to p but the left-hand side 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. $\square$

Proof of Proposition 7. As in Theorem 6 and Proposition 4, it is enough to find thresholds

$0 < \underline{\theta} < \overline{\theta} < 1$ instead. Once again, bribes are bounded above by 1, so for $(1-\theta) < \underline{C}$, or $\theta > 1 - \underline{C}$, the reviewer will opt to not write a review in any period, which minimizes consumer utility as noted in Proposition 3. Therefore, taking $\bar{\theta} = \max\{0, 1-\underline{C}\} < 1$, we establish the upper threshold for any given $\underline{C}$.

In any equilibrium where the high-skill reviewer does not abstain, we know she mimics low type because $\beta = 0$. As $\theta \to 0$, firms pay a maximal bribe given by $\min\{p\bar{X}^*, \bar{X}^* - \gamma\}$, which can be seen by using the expression in Proposition 6 and substituting $\overline{X} = \bar{X}^*$ and $\underline{X} = 0$. By Assumption 2 (i.e., $\bar{X}^* - \gamma > 0$) implies $b^* > 0$ for $p > \bar{p}$ with $\bar{p}$ sufficiently close to 1. Letting b^* be the bribe for $p = \bar{p}$, we see that as $\theta \to 0$ the reviewer receives at least $b^* > 0$, which implies there is an open interval $(\overline{\theta}, 1)$ where the reviewer can receive bribes bounded below by some $\underline{b} > 0$ for $\theta \in (\overline{\theta}, 1)$. Choosing $\overline{C} = \underline{b}/2$ and $\underline{C} = \underline{b}/4$ completes the proof, as the high-type and low-type reviewer both receive the same bribes and mimic the low type for all $\theta \in (\overline{\theta}, 1)$, which is second-best. $\square$

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.

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

Best, James and Daniel Quigley (2016), "Honestly dishonest: A solution to the commitment problem in bayesian persuasion." Technical report, Mimeo.

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.

Cripps, Martin W., George J. Mailath, and Larry Samuelson (2004), "Imperfect Monitoring and Impermanent Reputations." *Econometrica*, 72, 407–432, URL https://www.jstor.org/stable/3598908. 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.

Fainmesser, Itay P and Andrea Galeotti (2019), "The market for online influence." *Johns Hopkins Carey Business School Research Paper*.

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.

FTC (2019), "Ftc brings first case challenging fake paid reviews on an independent retail website."

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.

Fudenberg, Drew and Jean Tirole (1991), *Game Theory*. MIT Press.

Guardian (2004), "Fake critic still haunting sony." *The Guardian*, URL https://www.theguardian.com/film/2004/jan/30/news2.

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

Helft, Miguel and Nick Bilton (2010), "Consumer reports says iphone 4 has design flaw." *The New York Times*, URL https://www.nytimes.com/2010/07/13/technology/13apple.html.

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.

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.

Mostagir, Mohamed, Asu Ozdaglar, and James Siderius (2019), "When is society susceptible to manipulation?" *Available at SSRN*, URL https://ssrn.com/abstract=3474643.

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.

Stern, Joanna (2019), "Apple still hasn't fixed its macbook keyboard problem." *The Wall Street Journal*, URL https://www.wsj.com/graphics/apple-still-hasnt-fixed-its-macbook-keyboard-problem/.