

Harvesting Ratings

Johannes Johnen[†] and Robin Ng[‡]

[†]CORE/LIDAM, UCLouvain

[‡]University of Mannheim

February 4, 2024

Abstract

Evidence suggests lower prices lead to better ratings, but better ratings induce firms to charge higher prices in the future. We model that consumers are only willing to make the effort to rate a seller if this seller provides a *sufficient* value-for-money. Using this model, we explore how firms use prices to impact their own ratings. We show that firms harvest ratings: they offer lower prices in early periods to trigger consumers to leave a good rating in order to earn larger profits in the future. Because especially low-quality firms harvest ratings, harvesting makes ratings less-informative about quality. Based on this mechanism, (i) we argue that rating harvesting causes rating inflation; (ii) we show that a marketplace that facilitates ratings (e.g. through reminders, one-click ratings etc.) may get more ratings, but also less-informative ratings; (iii) a marketplace that screens the quality of sellers makes ratings less-informative if the screening is insufficient. Counter to the conventional wisdom that consumers benefit from ratings via the information they transmit, we show that consumers prefer somewhat, but never fully informative ratings. Nonetheless consumers prefer more-informative ratings than average sellers. We apply these results to characterise when a two-sided platform wants to facilitate ratings, and argue that efforts of major platforms to facilitate ratings did not just lead to less-informative ratings, but also shifted surplus from consumers to sellers.

JEL: D21, D83, L10

Keywords: Ratings and Reviews, Digital Economy, Reputation

This paper was previously titled ‘Ratings and Reciprocity’. We thank Yair Antler, Paul Belleflamme, Maren Hahnen, Paul Heidhues, Botond Kőszegi, Xavier Lambin, François Maniquet, Simon Martin, Martin Peitz, Andrew Rhodes, Robert Somogyi, Clément Staner, Nikhil Vellodi, Julian Wright, and many seminar participants for helpful comments. We thank the Fédération Wallonie-Bruxelles for funding the Action de recherche consercée grant 19/24-101 ‘PROSEco’. Robin Ng gratefully acknowledges support from the Belgian National Fund for Scientific Research (FNRS) Aspirant Grant (FC46885) and the Deutsche Forschungsgemeinschaft (DFG) CRC TR 224 (Project B05).

1 Introduction

Ratings feature prominently in the decisions we make each day. We often rely on the experience of others, through ratings, to make decisions. Examples include choosing a holiday resort, buying a car, or where to have lunch. But how far can we trust these ratings to inform us about the quality of a product?

Despite having an influential role on the way consumers make purchase decisions, evidence suggests that ratings can be a poor signal of objective product quality (De Langhe et al., 2016; Siering et al., 2018). A key reason is that especially prices influence ratings. For example, Li and Hitt (2010) show that a 1% increase in price lowers ratings by 0.7 on a 1-10 point scale, so the effect of price on ratings can be quite sizeable.¹ While this suggests that prices affect ratings, other evidence shows that firms with better ratings set higher prices (Luca & Reshef, 2021). Combined, this evidence suggests a dual role between prices and ratings. Firms use lower prices to obtain good ratings, as this allows them to set higher prices in the future.

The body of theoretical literature on reputation tends to assume that reputation is solely a consequence of quality (See Bar-Isaac and Tadelis (2008) for a survey). In doing so, these papers focus on only one direction of effects: how ratings can influence future prices. In this paper, we contribute to the theoretical literature by considering the dual role of prices in ratings. We ask: ‘how do firms use prices strategically to influence their own future ratings, and how does this impact how informative ratings are about quality?’

To develop a framework where prices influence ratings, we posit that consumers rate if the value-for-money is sufficiently large to compensate them for the cost they face from leaving a rating. We motivate that consumers rate this way through two plausible psychological mechanism. First, building on the well-established concept of reciprocity (Bolton & Ockenfels, 2000; Dufwenberg & Kirchsteiger, 2004; Rabin, 1993), this rating mechanism results if consumers perceive good deals as a kindness by the firm and reciprocate with a good rating. Second, this captures that consumers rate if they feel good about a deal. With either interpretation, this mechanism captures that given the true quality and price, the consumer is willing to make the effort to rate if a seller offered a sufficiently large value-for-money. This way, we endogenize the rating decision of consumers and enable firms to use prices to influence ratings.

Formally, we study a two period model with asymmetric information. Two actors participate

¹They look at overall product ratings on CNET.com and Dpreview.com. The effect has a similar magnitude when consumers are asked to rate value-for-money explicitly.

in each period—a long lived monopolist and short lived consumers. The firm is endowed with a product of high or low quality. The firm knows the quality of its product, and sets prices in each period to maximise lifetime profits. Initially, consumers only know the distribution of quality. Each consumer only participates in a single period. At the start of each period, consumers observe the historical ratings of the firm and its price in the current period. Crucially, consumers cannot distinguish if the product has a good rating because it has a high quality, or because it was sold cheaply in the past. Consumers only learn of the true quality of the product after consumption, and may then choose to leave a rating. Thus, consumers are uncertain about quality at the beginning of each period; but they may use ratings to transmit some information about quality to future consumers.

This model represents markets such as Amazon and Taobao, where ratings play the important role of developing trust between anonymous users. In these markets, new consumers are mostly unaware of the true quality of sellers they are transacting with. They rely on information left by prior consumers—through ratings—to form beliefs over the quality of a product. Further, these websites only list historical ratings and current prices, so consumers cannot distinguish if a rater left a good rating because she enjoyed a high quality, or because she paid a low price.²

Our key novel mechanism is the following trade-off of low-quality firms. First, in period 1 they can set a price which is sufficiently low to trigger consumers to leave a good rating. By doing so, the low-quality firm receives a good rating and earns larger profit in period 2. We call this strategy *‘ratings harvesting’*. Second, in period 1 the low-quality firm could charge the same (larger) price as the high-quality firm to make consumers believe this seller could be of a high quality. Because both firms charge the same price in period 1, we call this strategy *‘price mimicking’*.

In line with evidence that high-quality firms are more likely to get good ratings (Ananthakrishnan et al., 2019; Li et al., 2020), we focus on equilibria where high-quality firms set sufficiently low prices to obtain a good rating. But the rating of low-quality firms depends on their pricing strategy. If low-quality firms harvest ratings, both types of firms get a good rating, and ratings do not help consumers to distinguish firms. But if low-quality firms mimic prices, only high-quality firms obtain a good rating, allowing future consumers to better distinguish firms.

In equilibrium, low-quality firms may play a mixed strategy and harvest ratings and mimic prices with strictly positive probability. If leaving a rating is very costly for consumers,

²While some services may exist to attempt to track historical prices, their validity cannot be ascertained and these services are unable to reflect the transaction price associated with each rating.

low-quality firms mimic prices so that ratings perfectly signal quality. But if leaving a rating becomes sufficiently easy, low-quality firms can harvest ratings at a larger price so that low-quality firms harvest ratings and mimic prices with strictly positive probability. Intuitively, low-quality firms harvest ratings to get a good rating and free-ride on the reputation of high-quality firms. But this undermines the value of a good rating until, in equilibrium, low-quality firms are indifferent between rating harvesting and price mimicking.

Our key trade-off induces several interesting equilibrium features.

First, our equilibrium links the pricing strategy of firms with how informative ratings are about product quality: low-quality firms obtain a good ratings if and only if they harvest ratings. This is why the probability that low-quality firms mimic prices in equilibrium also measures how precisely ratings signal quality.

Second, rating harvesting causes rating inflation.³ If low quality firms harvest ratings, they get a better rating. But if low-quality firms get better ratings, ratings become less informative about quality. But our results also shed new light on this well-documented strategy to build reputation: low-quality firms do so to free-ride on the reputation of high-quality firms, which undermines how well ratings signal product quality.

Third, the equilibrium captures evidence on the dual nature of prices and ratings we described above. In our setting, high-quality firms and low-quality firms that harvest ratings get a good rating and charge larger prices over time. This captures evidence that firms build reputation in early periods by obtaining good ratings and then raise prices (Cabral & Hortaçsu, 2010; Cabral & Li, 2015; Li et al., 2020).

Major platforms regularly change their rating environment. We leverage our insight that rating harvesting makes ratings less informative to understand how changes in the ratings environment impact the informativeness of ratings. To do so, we explore how different features of a marketplace influence ratings, namely (i) how easy it is to rate and (ii) the extent of quality controls.

Many major platforms try to encourage and facilitate ratings by lowering the effort it takes for consumers to leave a rating. For example, Amazon transitioned to a one-click rating system, arguing that—in the spirit of the law of large numbers—more ratings “more accurately [...]

³Rating inflation here refers to the observation that rating scores are improving over time, and most of the improvements cannot be attributed to product quality. Thus, leading to ratings becoming a less effective signal of quality. This observation is made in Filippas et al. (2022), suggesting that other attributes such as the cost of leaving ratings, kindness to seller and other forms of retaliation contributes to rating inflation.

reflect the experience of all purchasers”.^{4,5} We show that this logic ignores how a lower effort to rate impacts prices. Making it easier to rate means that firms need only transfer a smaller rent to consumers to encourage a good rating. Thus, low-quality firms can harvest ratings with higher prices and will do so more often in equilibrium. Even though this leads to more ratings in equilibrium, ratings are also less-informative. Evidence suggests these effects can be quite large: Cabral and Li (2015) pay consumers to leave any rating to lower their opportunity cost to rate; they find that giving consumers \$1 to leave any rating leads to 22% less negative ratings.⁶ More generally, together with evidence that platforms facilitate ratings over time, this result suggests platforms engage in a race towards uninformative ratings and reinforce rating inflation.

This result also suggests that recent policy proposals to improve rating environments may not go far enough. Crawford et al. (2023) propose that sellers should not be able to pay raters conditional on the content of their ratings or reviews. We show that unconditional payments for ratings—even when not discriminating between worse or better ratings—leads to less-informative ratings.

We also explore the impact of rating harvesting on surplus. Even though rating harvesting makes ratings less informative, consumers may still benefit: low-quality firms who harvest ratings lower prices and provide some surplus to consumers. We call this surplus ‘harvesting rent’. Since consumers only receive this surplus when firms harvest ratings, they benefit despite less informative ratings. This has two key implications: (i) Consumers do not prefer fully-informative ratings. If low-quality firms do not harvest ratings, ratings are perfectly informative, but consumers do not receive any harvesting rent. (ii) Counter to the conventional wisdom that consumers mainly benefit from ratings via the information they transmit, we show that consumers may benefit from lower prices that firms charge to induce better ratings.

We show that easier ratings affect buyer and seller surplus differently. Sellers have polarized views on ratings: High-quality firms unambiguously prefer a higher effort to rate, because this prevents other firms from free-riding on their reputation. But a higher effort to rate reduces harvesting rents and average sellers prefer easier ratings. Consumers, however, prefer an intermediate effort level that leads to somewhat informative ratings. Thus, consumers prefer more-informative ratings than the average seller, but less-informative ratings than

⁴Quote from article Rey (2020) on vox.com.

⁵Prior to 2019, Amazon had required raters to leave a 20-word review along with their rating. Documented by Amazon reviews and forums (Amazon Customer, 2012; crebel, 2017).

⁶They do their field experiment in a field experiment on eBay. Effectively, \$1 discount amounts to a price reduction of 25%.

high-quality firms.

These results suggest that a platform can facilitate or discourage ratings to shift surplus between sellers and buyers. Using this insight, we show when a two-sided platform may facilitate ratings to encourage more, but less-informative ratings. This insight suggests that the above-mentioned efforts of major platforms to facilitate ratings did not just lead to less-informative ratings, but also shifted surplus from consumers to sellers.

Marketplaces do not just facilitate ratings, but also employ quality controls to weed out low-quality firms. For example, Amazon suspends sellers who do not meet a minimum standard.⁷ We show that improving the aggregate quality in the market can discourage ratings harvesting, leading to more-informative ratings and less rating inflation. But when aggregate quality is low, quality improvements can also foster rating inflation instead.

Finally, we study a range of extensions and robustness checks. We show that competition between sellers encourages ratings harvesting and leads to more rating inflation. Additionally, our results are robust when consumers can leave negative ratings. In this extension, we show that value-for-money based ratings can explain why we observe extreme (very positive or negative) ratings in practice. We also discuss how our results extend to rating systems with more different ratings such as 5-star ratings. In an extension beyond two periods, we capture the evidence that firms harvest ratings, but that lower-quality firms are less likely to maintain good ratings (Cabral & Hortaçsu, 2010; Jin & Kato, 2006). We also show that our results persist when we introduce a continuum of firms.

We introduce the basic model in Section 2, and discuss the equilibrium in Section 3. Section 4 shows how various features in the rating system influence how well ratings reflect quality. We then discuss implications on surplus in Section 5. We present extension and robustness checks in Section 6. Section 7 connects our results to the literature, and Section 8 concludes.

2 Basic Model

We set up a two-period model of incomplete information with a long lived monopoly firm and a unit mass of buyers in each period.

Firms. The firm can be of a high or low quality. A firm of type $j \in \{H, L\}$ has quality q^j , where $q^H > q^L$. The probability that the firm is of type q^H or q^L is common knowledge

⁷See information disseminated to sellers by Amazon (Rushdie, 2018) and a Bloomberg news article (Soper, 2019).

and given by $\gamma \in (0, 1)$ and $(1 - \gamma)$, respectively. The realized quality is private information to the firm and is constant between periods. In each period, the firm sets the price of its product to maximise its lifetime profit, $\sum_{t=1}^2 p_t^j \cdot d_t^j$, where p_t^j is the price of firm j in period t , and $d_t^j \in \{0, 1\}$ is the demand of firm j in period t . We assume that the cost of production is zero regardless of quality.⁸ After selling in period 1, sellers may receive a rating R_1 from consumers. If they do, this rating is made common knowledge to both the firm and consumers in subsequent periods.

Consumers. Consumers have homogeneous valuations. They participate in only one of the two periods and a new unit mass of consumers arrives in each period. When choosing to consume a product in period t , consumers observe the price on offer and past ratings, R_{t-1} . They do not observe the firm's quality or the price the rater paid.⁹ Consumers may choose to buy or not to buy. After they buy, they observe the firm's true quality and decide on leaving a rating.

To simplify illustration in the main model, we focus on a binary rating system where consumers can choose between leaving a rating or not. More precisely, ratings take the form $R_t \in \{0, 1\}$. The informational content of a rating will be determined in equilibrium, but we say that a rating is good if $R_t = 1$, and when consumers choose not to provide any rating, we say $R_t = 0$. We focus on positive ratings to simplify exposition and we show below that our main results are qualitatively robust when consumers are additionally able to leave bad ratings. We also discuss how our results extend to settings with even more rating messages, such as 5-star ratings. Without loss of generality, we say $R_0 = 0$, i.e. firms have no previous ratings.¹⁰

We distinguish between consumption utility and rating utility. This serves two purposes. First, we are able to capture the phenomenon that consumers do not factor the intention to rate into their purchase decision.¹¹ Second, this simplifies presentation of results. The consumption utility for consuming a product from firm j in period $t \in \{1, 2\}$ is given by $u_t = q^j - p_t^j$, where p_t^j represents the price that the firm sets in period t . If a consumer does

⁸We assume zero marginal cost to simplify the analysis. We only require that low-quality firms face a lower marginal cost than high-quality firms.

⁹Our results are robust as long as consumers cannot be sure if a given rating is the result of a high quality or a low price. Thus, our results are robust when consumers can observe past prices, as long as they do not observe the price that a rater paid when she gave her rating.

¹⁰This captures cases where the firm is new in the market. But also any history where two firms of different quality have the same rating at the beginning of period 1.

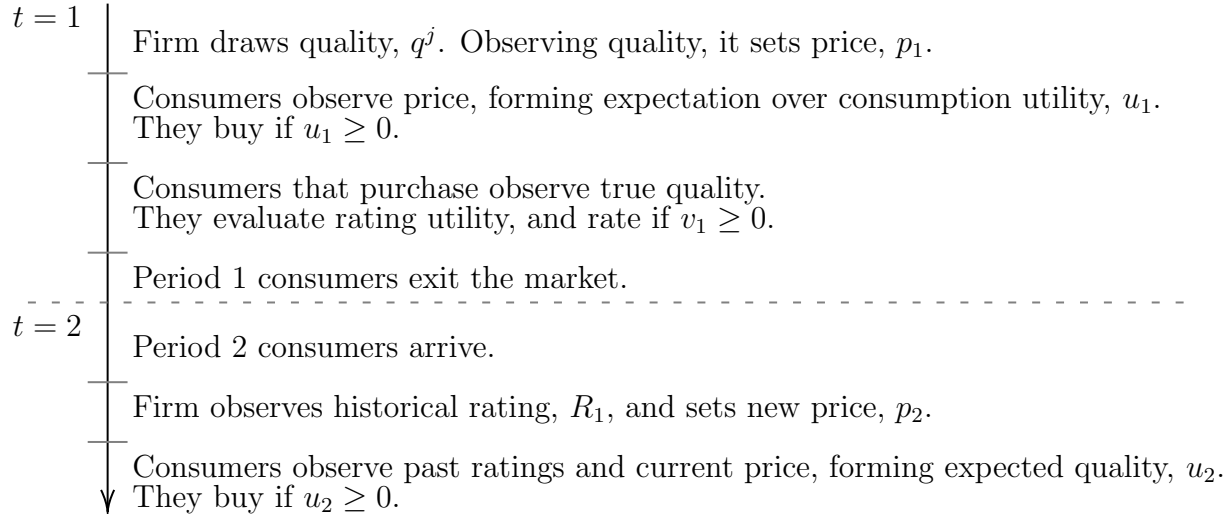
¹¹More precisely, Cabral and Li (2015), find that incentivizing consumers to rate does not change their willingness to pay for a product, suggesting that whether consumers rate or not does not affect their value of the product.

not buy, she receives her outside option which we normalize to zero.

The rating utility captures the mechanism that motivates consumer to rate. Formally, the rating utility of consumers in period t is $v_t = q^j - p_t^j - e$ if $R_t = 1$, and $v_t = 0$ if $R_t = 0$. Here, $e \geq 0$ reflects the time and effort consumers need to invest to provide a rating. Thus, consumers leave a rating if they earn a sufficiently large value-for-money, i.e. if $q^j - p_t^j \geq e$.

We provide two foundations for this mechanism. First, this rating utility is a simplification of models of intrinsic reciprocity (Dufwenberg & Kirchsteiger, 2004; Rabin, 1993).¹² Reciprocity is a long-studied and well established decision-making phenomenon and therefore provides a solid foundation for our rating utility. Second, consumers rate out of self-expression, i.e. if they feel good about a purchase. Firms can induce this good feeling by giving consumers a good value-for-money. Either interpretation captures the idea that consumers are happy to make the effort to rate a firm if they receive a sufficient value-for-money.

Timing of game. To summarize the timing of the game,



This model represents markets such as Amazon, Taobao, eBay, and Google reviews. Consumers rely on ratings to form or update their expectations of product quality. A key feature of our model is that consumers cannot distinguish whether a rater left a good rating because she enjoyed a high quality or paid a low price. Indeed, in these applications consumers

¹²All our results apply for the following generalized rating utility: $v_t = [\kappa q^j - p_t^j] \Delta - e$ if $R_t = 1$ and $v_t = 0$ if $R_t = 0$. $\kappa \in [0, 1]$ represents the proportion of surplus that consumers think is equitable for firms to receive; $\Delta > 0$ represents the warm glow a consumer enjoys from reciprocating the value-for-money the firm provides with a good rating. We obtain the rating utility in the main text for the special case where $\kappa = 1$ and $\Delta = 1$. All our results are qualitatively robust for $\kappa \in [0, 1]$ and $\Delta > 0$. Details can be found in the Web Appendix B.3.

usually do not observe the prices that a previous consumer paid when she left a rating.¹³ Some of these platforms have more detailed ratings systems such as 5-star ratings. But our simple rating system is quite close to the rating system by eBay that allows for positive and negative reviews; and we discuss below how our results extend to rating systems with negative ratings or 5-star ratings.

In the main model, we make some simplifying assumptions. We discuss in Section 6 how results are robust to various extensions, namely a game with more than 2 periods, a continuum of quality types of the firm, and rating systems with more messages.¹⁴

3 Equilibrium

We look for a perfect Bayesian equilibrium. We apply two equilibrium selection assumptions.¹⁵

Selection Assumption 1. *We focus on the equilibria where high-quality firms obtain a rating of $R_1 = 1$ with probability 1.*

Without this selection assumption, there may exist equilibria where the high-quality firm does not get a rating. In these equilibria, ratings signal low quality so that ratings do not help sellers build a reputation. Such situations, however, seem economically less interesting: evidence suggests that high-quality firms are more likely to obtain ratings (Ananthkrishnan et al., 2019; Li et al., 2020). The selection assumption implies that, in equilibrium, consumers who observe a rating expect a weakly higher product quality, and firms can indeed use ratings to build a reputation for quality.

Our second selection assumption limits beliefs off the equilibrium path.

Selection Assumption 2. *For all prices in period t such that low-quality firms obtain no rating, the expected quality in period t is independent of prices.*

Selection Assumption 2 implies that in a given period t , consumers have the same beliefs

¹³Also reviews very rarely mention prices. We discuss this more carefully below in the conclusion.

¹⁴Lewis and Zervas (2019) show that only the relative difference in stars affects the pricing decision of firms. This suggests that it is more important to consider the effect of a relative difference in ratings, which we already do with our simple binary-rating framework. Other papers find that negative ratings have a statistically insignificant impact on prices (Bajari & Hortacısu, 2003; Cabral & Hortacısu, 2010; Livingston, 2005; Resnick et al., 2006). Hence, we believe that this simplification is well justified. Although, as we show in our extension, allowing for more flexibility in the ratings scale does not qualitatively change our result.

¹⁵These selection assumptions are similar in spirit to the restrictions used in Rhodes and Wilson (2018), who study how advertisement can signal quality.

about quality for any price that induces the same rating. Without this selection assumption, there exist additional, qualitatively similar, equilibria at different price levels. In these equilibria, firms do not deviate to higher prices, because consumers who observe an off-equilibrium price have overly pessimistic beliefs, e.g. q^L , and no longer purchase. Instead, this selection assumption leaves an equilibrium where firms' pricing decisions are restricted by the threat of ratings.¹⁶ Thus, firms can extract the conditional expected surplus of consumers, but they may choose not to do that in order to manipulate ratings.

The following proposition characterizes equilibria in this game. The formal proof is in Appendix A.

Proposition 1. *All perfect Bayesian equilibria satisfy the following.*

1. *Ratings build reputation. In period 2, prices equal expected quality conditional on ratings, where $E[q_2|R_1 = 1] > E[q_2|R_1 = 0]$.*
2. *In period 1, high-quality firms charge $\bar{p} \equiv \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$ with probability 1 and receive a good rating.*
3. *In period 1, low-quality firms randomize their price.*
 - a. *They charge \bar{p} with probability δ^* and obtain no rating.*
 - b. *They charge $\underline{p} \equiv q^L - e$ ($< \bar{p}$) with probability $1 - \delta^*$ and obtain a good rating.*
 - c. *$\delta^* \in (\frac{1}{2}, 1)$ if and only if*

$$(1 - \gamma)(q^H - q^L) > e, \tag{1}$$

and $\delta^ = 1$ otherwise.*

The equilibrium is unique up to off-equilibrium-path beliefs and exists if $q^H - e \geq \frac{\gamma q^H + \delta^(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$.*¹⁷

To understand the key trade-off in this equilibrium, we first make two observations. Observations 1: firms who build reputation earn larger profits in period 2. To see this, we look at high-quality firms. By our Selection Assumption 1, we focus on equilibria where high-quality firms build reputation. This has two implications. First, consumers who observe a rating will expect a larger quality, so that ratings lead to larger prices in period 2. Thus, ratings are 'good' ratings in equilibrium. Second, to get a good rating in period 1, high-quality firms

¹⁶Alternatively, we could use the common D1 refinement (Cho & Sobel, 1990), but our selection assumption is actually weaker.

¹⁷The existence condition captures that e cannot be too large, or high-quality firms prefer not to build a reputation. Intuitively, if e is very large, firms need to charge a very low price to get a good rating.

charge a price \bar{p} . This price is low enough so that consumers who learn the firm is of high quality get sufficient value-for-money to leave a good rating.

Observation 2: in period 1, firms charge one of two prices in equilibrium. High-quality firms always set the high price \bar{p} to extract all expected surplus (conditional on observing \bar{p}). If consumers could distinguish high- and low-quality firms, this price would be equal to q^H . But because in period 1 consumers cannot distinguish firms, low-quality firms can mimic the high-quality firms and charge \bar{p} . This, in equilibrium, lowers the conditionally expected surplus below q^H . The low price \underline{p} is just low enough so that consumers rate low-quality firms positively, i.e. \underline{p} is such that $q^L - \underline{p} - e = 0$.

The following key trade-off drives our results: in period 1, low-quality firms choose between these two prices. Low-quality firms benefit from ratings by free-riding on the reputation of high-quality firms to earn higher prices in period 2 (point 1 in Proposition 1). But to get a rating, they have to charge a low price \underline{p} ($< \bar{p}$) in period 1 so that consumers are willing to rate them positively. This is why we call this strategy ‘ratings harvesting’. Alternatively, in period 1 low-quality firms can mimic the high price \bar{p} of high-quality firms; but then in period 2 they get no ratings and lower profits. Because firms who follow this strategy copy the price of high-quality firms, we call it ‘price mimicking’. Thus, low-quality firms trade-off ratings harvesting and price mimicking. The probability that low-quality firms mimic prices and do not receive a rating, δ^* , captures how firms resolve this trade-off in equilibrium.

Why may low-quality firms set $\delta^* \in (\frac{1}{2}, 1)$? Intuitively, low-quality firms harvest ratings to free-ride in the reputation of high-quality firms. This, however, undermines the expected quality associated with a good rating until—in equilibrium—low-quality firms are indifferent between ratings harvesting and price mimicking.¹⁸

The equilibrium has several interesting features.

First, the equilibrium features the dual role of prices and ratings we discussed in the Introduction: lower prices allow firms to build reputation and charge larger prices in the future. In our equilibrium all firms who receive a good rating (high-quality firms and low-quality firms who harvest ratings) raise their price in period 2.

Related to this feature, our equilibrium captures the growing evidence on how quality and prices affect ratings. More precisely, consumers seem to rate based on the value-for-money they obtain from a purchase. For example, studying marketplaces for digital cameras, Li and

¹⁸Formally, benefit of harvesting is the price difference in period 2, i.e. $E[q_2|R_1 = 1] - E[q_2|R_1 = 0] = \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} - q^L$. This decreases in δ^* . In equilibrium, $\delta^* > \frac{1}{2}$, because the benefit of harvesting needs to be sufficiently large.

Hitt (2010) highlight that a 1% price increase reduces ratings by 0.36 stars in 5-star ratings and 0.71 stars for 10-star ratings. On AirBnB, Gutt and Kundisch (2016) and Neumann et al. (2018) show that prices negatively impact ratings. On Yelp, Luca and Reshef (2021) provide evidence that a 1% price increase lowers average ratings by 3-5%. Abrate et al. (2021) suggests that a 1% increase in hotel prices lowers overall ratings by 1 star (out of 10). Because these articles control for product quality, they suggest that it is value-for-money that influences ratings, and not quality alone.

Second, rating harvesting induces rating inflation. Rating inflation describes the phenomenon that the vast majority of ratings on many websites are the best ratings, e.g. 5 out of 5 stars (Filippas & Horton, 2022; Filippas et al., 2022; Nosko & Tadelis, 2015; Zervas et al., 2021). A major concern about rating inflation is that when sellers receive mostly the best ratings, ratings are not very informative about product quality.

Rating harvesting induces rating inflation in the following way. Suppose (1) is violated so that e is just above $(1 - \gamma)(q^H - q^L)$. Then low quality firms do not harvest ratings and ratings are a perfect signal of quality. But as e decreases and (1) starts to hold, low-quality firms lower prices to offer a larger value-for-money and boost their ratings. This benefits these firms, but, in equilibrium, makes ratings less-informative about product quality. This way, if firms do more rating harvesting, the equilibrium features more rating inflation.

Third, the equilibrium links the pricing strategy of firms with how precisely ratings inform about product quality. In particular, as δ^* increases and low-quality firms mimic prices more, low-quality firms obtain a rating less often, and ratings are more likely to indicate a high-quality firm. In turn, if low-quality firms harvest ratings more, so that δ^* decreases, a rating is more likely to represent a low-quality firm that charged a low price in period 1. Thus, δ^* captures the firm's pricing strategy, but also how precisely ratings inform about quality.

In the next section, we exploit this feature and study how various design features of markets impact how informative ratings are. To do so, we restrict the remainder of our analysis to situations where (1) holds so that the low-quality firm plays a mixed-strategy. The condition is satisfied when the difference in quality is sufficiently large, i.e. when ratings are more relevant in the first place.

4 Designing ratings environments

In this section, we discuss how common features of a rating environment influence the informativeness of ratings.

Facilitating ratings

A key design tool that major platforms use to influence ratings is how easy it is to rate. For example, raters may have to complete a verification process, they may be asked to rate along multiple dimensions, or raters may receive monetary rebates and reminders to rate. These design features influence the time and cognitive effort it takes to evaluate a product, and therefore the cost of leaving a rating.

Intuitively, the law of large numbers would suggest that collecting more ratings would lead to more precise and more-informative ratings. From this perspective, reducing the cost of leaving a rating seems to be a good idea. We show that this intuition is misleading as it ignores how firms adjust their pricing strategy when ratings become easier.

In our setting, less-costly ratings lead more ratings, but also less-informative ratings. If it gets easier to rate (i.e. e decreases), a consumer's rating utility v_t increases. This means that firms can provide a smaller value-for-money and still induce consumers to give a good rating. So low-quality firms can harvest ratings with a larger price, which encourages them to harvest ratings with a larger probability. But then, in equilibrium, a good rating is now more likely to represent a low-quality firm, which makes ratings less-informative about quality. The following corollary summarized this result.

Corollary 1. *If (1) holds, then $\frac{\partial \delta^*}{\partial e} > 0$.*

The corollary provides another channel through which value-for-money based ratings induces rating inflation: easier ratings encourage especially low-quality firms to harvest ratings. In turn, this means that making ratings costly can make them more-informative.

This result connects well to evidence. Cabral and Li (2015) measure quality using shipping speed. They find that offering rebates for rating decreases the proportion of negative ratings, especially for low-quality products. Since rebates offset the cost of rating, their result highlights—as in our model—that lower costs of ratings make ratings less-informative about quality.

This insight sheds new light on the effort of major platforms to encourage ratings. The cost to leave a rating has decreased over time on many rating platforms. Yelp and Google encourage ratings with various perks, such as invitations to exclusive events and discount

codes.¹⁹ On google reviews, users receive constant reminders to leave ratings and can leave one-click ratings on their smartphone. Amazon has also moved towards fuss-free rating. Prior to 2020, Amazon required customers to write a review in order to leave a rating, subsequently removing this requirement and allowing for one-click ratings.²⁰ Publicly, and in line with the law of large numbers, Amazon stated that more feedback would lead to more accurate ratings.²¹ Our result suggests that these efforts of major platforms can backfire, because they ignore that easier rating affects prices. Although more customers may rate, easier rating encourages rating harvesting. This can lead to more ratings, but ultimately less-informative ratings.

Quality control

We have shown that (i) rating harvesting and (ii) easier rating (e) lead to less-informative ratings and therefore trigger rating inflation. In both cases, however, rating inflation is the result of less-informative ratings. But in practice, ratings may also improve because the average product quality improved over time.

The aggregate quality on a marketplace may change for a variety of reasons. First, low-quality firms may invest in better quality (Klein et al., 2016) or leave the market (Cabral & Hortaçsu, 2010; Nosko & Tadelis, 2015). Second, platforms may weed out low-quality firms (Casner, 2020; Nosko & Tadelis, 2015; Wang, 2021). Both channels can improve the aggregate quality on the market and ultimately lead to better average ratings. If firms have higher average quality, however, the incentives of low-quality firms to harvest ratings also changes.

We show that when the aggregate quality improves (i.e. γ increases), the remaining low-quality firms may harvest more or less ratings, depending on the aggregate quality level of the market. More precisely, for intermediate levels of aggregate quality, firms harvest ratings most intensely and ratings are least informative.

Intuitively, when most firms have similar quality (γ is either high, or low), consumers do not need ratings to predict firm quality, which is why ratings are less useful to build reputation; and when ratings build less reputation, low-quality firms harvest ratings less. In turn, when firms are heterogeneous (intermediate levels of γ), ratings are more useful to build reputation. But when ratings build reputation, low-quality firms will harvest ratings more to free-ride

¹⁹Yelp Elite Squad and Google Local Guides perks as described by Yelp (Yelp, 2022) and the Harvard Business Review article Donaker et al. (2019).

²⁰A timeline of Amazon’s rating system as documented by Forbes (Masters, 2021).

²¹As reported by TechCrunch (Perez, 2019).




Figure 1: Effect of quality control on informativeness of ratings, evaluated at $q^H = 0.5$, $q^L = 0.3$, $e = 0.05$.

on that reputation. As a result, $\delta^*(\gamma)$ is U-shaped (See Figure 1.).

Thus, for low aggregate quality, a quality increase encourages rating harvesting. But for large aggregate quality, a quality increase discourages rating harvesting. The following proposition summarizes this result.

Proposition 2. *There exists a unique $\bar{\gamma} \in (0, 1)$ such that $\frac{\partial \delta^*}{\partial \gamma} < 0 \iff \gamma < \bar{\gamma}$, and $\frac{\partial \delta^*}{\partial \gamma} > 0 \iff \gamma > \bar{\gamma}$.²²*

This result connects with the observation of quality-controls in practice. For example, Amazon actively enforces seller quality, suspending sellers who do not meet a minimum standard.²³ Also Uber has announced that it will remove both riders and drivers with consistently poor ratings,²⁴ and both Uber and their subsidiary Uber Eats suspend drivers who fall below a minimum rating.²⁵ Booking.com suspends properties for quality control purposes,²⁶ Airbnb bans hosts based on a combination of factors, including being in the bottom 1% of overall ratings and guest feedback.²⁷

Our results suggest how such quality screening impacts ratings. Especially for platforms with low aggregate quality, screening sellers can encourage rating harvesting and make ratings less informative. In turn, platforms with large aggregate quality who screen sellers can boost how informative their ratings are. Thus, platforms with low aggregate quality may only want to

²² $\bar{\gamma}$ is given by (11).

²³Amazon makes this decision through a combination of customer reviews, feedback and other measures (Rushdie, 2018; Soper, 2019).

²⁴Uber announces that it will begin deactivating riders with poor ratings (Dickey, 2019).

²⁵Leaked documents from Uber suggest drivers risk deactivation if they fall below 4.6 stars as reported by TechCrunch (Dickey, 2019). Uber Eats communication with couriers as posted on a forum for drivers (UberLyftDriver, 2017).

²⁶Partners report that Booking.com closes their apartments on Booking.com Partner Hub (Prodius, 2020).

²⁷Information from Airbnb help center information (Airbnb, 2022) and a third party Airbnb management service (Zodiak, 2021).

start screening sellers, if they can do so sufficiently well.

Our insights also point to how platforms can use ratings to screen seller quality. Even though good ratings may not be fully-informative about underlying quality, no ratings or—as we show in an extension—bad ratings are quite informative about low quality.²⁸ Thus, even in scenarios when good ratings are noisy signals of quality, platforms can indeed use bad ratings to weed out low-quality firms.

5 Surplus Analysis

So far, we studied how rating harvesting affects the informativeness of ratings. We now explore how rating harvesting affects surplus.

Conventional wisdom suggests that more informative ratings create better matches between firms and consumers and ultimately raise surplus. Our results highlight a novel pricing benefit for consumers: even though rating harvesting leads to less informative ratings, it also leads to lower prices for consumers. Evidence underlines the relevance of this mechanism for consumer surplus: Cai et al. (2014) show that consumers benefit from ratings even before anyone rates, because firms lower their prices.

We now explore the link between the informativeness of ratings and surplus more carefully. The key insight is that consumer-optimal rating systems are often somewhat, but never fully, informative.

To start, we investigate how rating harvesting affects consumer surplus, which in our setting equals consumption utility.²⁹ In order to induce consumers to leave a good rating, firms need to set a price below the consumers' ex-post willingness to pay. This is why sellers, even though they are monopolists, leave a rent to consumers. Thus, the expected consumer surplus is strictly positive and reflects this harvesting rent

$$CS = (1 - \gamma)(1 - \delta^*)e. \quad (2)$$

In period 2, and in period 1 when consumers face a high price of \bar{p} , firms set prices to extract all conditionally expected consumer surplus. But in period 1 when consumers observe a low price, which happens with probability $(1 - \gamma)(1 - \delta^*)$, they get the surplus $q^L - \underline{p} = e$.

²⁸In practice, consumers may not rate for various reasons, so that products without ratings are not necessarily low quality. This is why we check that the result also holds in the setting with negative ratings which are informative about low quality.

²⁹The reason is that in equilibrium, the rating utility is equal to zero.

Low-quality firms share this surplus so that consumers leave them a good rating, allowing low-quality firms to free-ride on the reputation of high-quality firms.

Consumers prefer somewhat, but never fully-informative ratings. To see this, note that rating harvesting benefits consumers through lower prices, which induces less-informative ratings. This, however, does not imply that consumers prefer uninformative ratings in equilibrium. If ratings were completely uninformative, low-quality firms would not harvest ratings, and consumers would earn no surplus.

We now discuss implications of this result more carefully for the impact of rating effort on consumer surplus.

Cost of ratings

A larger cost of ratings (e) has two opposing effects on consumer surplus. First, and following directly from Corollary 1, low-quality firms harvest ratings less often; this tends to reduce consumer surplus. Second, however, the low price, \underline{p} , decreases: when consumers can rate less easily, they are less inclined to leave a good rating. This is why low-quality firms who harvest ratings have to charge a lower \underline{p} to receive a good rating, raising consumer surplus.

We show that these two opposing effects pin down a positive level of rating effort that maximizes consumer surplus.

Proposition 3. *There exists a level of effort $e^{cs} > 0$ that maximizes consumer surplus. This is true if*

$$(1 - \gamma)^2 \gamma (q^H - q^L)^2 \geq (1 + \gamma)^2 e^2. \quad (3)$$

Otherwise, consumers prefer $e^{cs} = 0$.

The proposition characterizes when the conventional wisdom that more-informative ratings benefit consumers holds in our setting. When e is sufficiently small ($e < e^{cs}$), making it more costly to rate leads to more-informative ratings (Corollary 1), and increases consumer surplus. In this case, the price effect on \underline{p} dominates and more-informative ratings put pressure on prices for low-quality firms.

More surprisingly, when e is sufficiently large ($e \geq e^{cs}$), more-informative ratings harm consumers. Intuitively, when e is large, low-quality firms charge a very low \underline{p} to harvest ratings. This is why, as e increases further, low-quality firms get discouraged from harvesting ratings, which reduces consumer surplus.

Also when considering effort to rate, the proposition implies that consumers still prefer a somewhat-informative rating system. The reason is that low-quality firms are only willing to harvest ratings if they can free-ride on the good reputation of high-quality firms; but this requires ratings to be somewhat-informative.

By condition (3) consumers prefer somewhat-informative ratings if the difference between high- and low-quality firms is sufficiently large. This is similar in spirit to (1) and rather intuitive: if the condition is violated and the difference in quality is small, then the price difference in period 2 is also small. Low-quality firms have little incentive to harvest ratings and hardly ever do so. Because low-quality firms harvest ratings so rarely, consumers want firms to harvest ratings more often and prefer $e^{cs} = 0$. This result, however, seems economically less relevant, since it only applies when the quality differences are small so that ratings are less relevant in the first place.

We now explore the impact of costly ratings on seller surplus. We show that sellers have polarized views on ratings. While high-quality sellers unambiguously prefer a high effort to rate, sellers on average prefer uninformative ratings. Intuitively, seller surplus is largest when sellers leave the smallest harvesting rent to buyers. This is true when $e = 0$ and ratings are least-informative.

While sellers on average prefer a small e and less-informative ratings, high-quality firms prefer a large e and perfectly-informative ratings: informative ratings allow high-quality firms to distinguish themselves from low-quality firms who try to free-ride on their reputation. As e increases, harvesting requires a larger discount and free-riding becomes more costly, which leads to more-informative ratings and allows high-quality firms to extract more of the surplus they generate. The following corollary summarizes these results.

Corollary 2. *When (3) holds, average seller surplus is maximal at $e^s = 0$ ($< e^{cs}$). Moreover, profits of high-quality firms increase in e , and profits of low-quality firms decrease in e .*

The corollary implies that only high-quality firms unambiguously prefer perfectly informative ratings because this limits free-riding on their reputation. Neither buyers nor sellers, on average, prefer perfectly informative ratings. But buyers prefer more informative ratings than the average seller. The reason is that somewhat informative ratings push sellers to harvest ratings, which puts pressure on prices.

The result suggests that marketplaces that facilitate ratings do not just affect how informative ratings are, they also shift surplus between buyers and sellers. Thus, marketplaces can encourage or discourage ratings to shift surplus between buyers and sellers.

In practice, as we outlined above, many platforms such as Google Reviews and Amazon have facilitated ratings in recent years. Our results suggest that these design choices lead to more ratings harvesting. This has further implications: First, firms obtain a good rating for a lower harvesting rent. Thus, facilitating ratings effectively transfers surplus from consumers to firms.³⁰ Second, by Corollary 1 ratings become less informative.

Quality controls

We now discuss the effect of aggregate product quality on consumer surplus. Our first insight is that improvements in aggregate quality can make consumers worse off.

The intuition has two steps. First, we know from Proposition 2 that low-quality firms harvest ratings most intensely for intermediate levels of γ . Because harvesting induces the harvesting rent, this result implies that consumer surplus is strictly concave in γ . However, since only low-quality firms harvest ratings and leave harvesting rents to consumers, consumer surplus is maximized for an aggregate quality below $\bar{\gamma}$.

The next proposition summarizes these results.

Proposition 4. *Equilibrium consumer surplus is strictly concave in γ . There exists an aggregate quality level, denoted by γ^{cs} , that maximises consumer surplus, where $\bar{\gamma} > \gamma^{cs}$ and $\gamma^{cs} > 0$.*

Proposition 4 implies that, even when larger aggregate quality leads to more-informative ratings, consumers can be worse off. This is the case when $\gamma \notin [\gamma^{cs}, \bar{\gamma}]$. For $\gamma < \gamma^{cs}$, better quality encourages low-quality firms to harvest ratings, which benefits consumers but undermines the informativeness of ratings. For $\gamma > \bar{\gamma}$, an increase in aggregate quality makes ratings more-informative; but as low-quality firms participate in less ratings harvesting, consumers surplus diminishes. Thus, to evaluate a rating system, observing that ratings reflect quality more closely is not enough to conclude that consumers benefit.

Sellers unambiguously benefit if their average quality increases. High-quality firms are able to set higher prices in both periods. Low-quality firms benefit either from being able to set a higher price in the first period when they mimic prices, or from setting a higher price in the second period when they harvest ratings.

Lemma 1. *The profits of high- and low-quality firms increases in γ . Seller surplus is maximised at $\gamma^s \rightarrow 1$*

³⁰This argument assumes that these platforms do not set $e > e^{CS}$. We argue that such large efforts are not optimal for a platform, as lower e would benefit both buyers and sellers and raise overall activity on a platform. We discuss a profit maximising platform's optimal design in Web Appendix B.1.1.

Because all firms benefit from a higher average quality, this also implies that the remaining firms prefer a higher level of quality controls than consumers.

6 Extensions and Robustness

Designing a profit maximising ratings system. In the main text, we take the properties of a rating system as given. But in this extension, we explore how a two-sided platform designs its rating system. The platform needs to offer value to both consumers and firms to get them to use the platform and generate transactions. Our insight from Corollary 2 implies that the platform can use e to shift surplus between buyers and sellers. We show that platforms may choose a very low e and induce a rather uninformative rating system when it favors sellers; but it may also design informative rating systems when it cares more about attracting consumers. Because the platform optimal rating system features an effort below e^{CS} , this result speaks towards concerns that the design of rating systems are insufficiently informative for consumers (Competition and Markets Authority (UK), 2017). We also discuss how a platform’s strategic choices may change over time, i.e. they may start with an informative rating system to attract buyers, but then facilitate ratings and accept less informative ratings to extract more profits from sellers. This reflects how Amazon’s rating system changed over time: from introducing new verification methods to allowing ratings to be provided without reviews (Amazon, 2021). For details, see Web Appendix B.1.1.

Competitive environment. We study how competing firms use ratings and introduce a competitive fringe who sells a product of known quality, at a price equal to its marginal cost c . We show that when c decreases and competition gets more fierce, low-quality firms participate in more ratings harvesting, and ratings become less-informative. Despite less-informative ratings, competition exerts pressure on all firms, which respond by lowering prices, benefiting consumers. Intuitively, competition improves the outside option of consumers, so sellers need to leave more rents to consumers. But low-quality sellers who harvest ratings already leave a harvesting rent to consumers, so they are less affected by competition and sellers harvest more ratings. For details, see Web Appendix B.1.2.

Negative ratings. We relax our simplifying assumption and allow for negative ratings—in addition to positive ratings and no ratings. In equilibrium, negative ratings arise from retaliation. When firms leave a sufficiently small value-for-money, consumers retaliate with a bad rating. In equilibrium, only low-quality firms receive negative ratings, which is why negative ratings transmit the same information as no ratings in the main model, and our

results are qualitatively robust.³¹ But—different from the main model—ratings are more extreme. For details, see Web Appendix B.1.3.

More generally, our results also extend to rating systems with even more messages such as 5-star rating systems. Intuitively, in our framework, the value of ratings is determined endogenously in equilibrium. Thus, also with more complex rating systems, there exist equilibria where one rating has the same informational content as our good rating, other ratings have the same informational content as our bad rating, and the others are uninformative or not used in equilibrium. This equilibrium is plausible because it reflects the common finding that ratings are strongly bimodal and raters leave either five stars or zero stars (Dellarocas & Wood, 2008; Filippas & Horton, 2022; Filippas et al., 2022; Hu et al., 2009; Nosko & Tadelis, 2015).³²

Continuum of Firms. We also show that our results extend beyond firms with two types and study a continuum of firms with different quality types. We find a pure-strategy equilibrium where (i) an interval of highest quality firms gets a good rating; (ii) an interval of lowest quality firms gets a bad rating, and (iii) a middle interval of firms harvests ratings and gets a good rating as well. In this equilibrium, middle-quality firms harvest ratings and make ratings less-informative. Additionally, as ratings become easier (e decreases), the middle interval of firms that harvest ratings grows, leading to less-informative ratings, just like in the main model. For details, see Appendix Web Appendix B.1.4.

Longer Horizon Model. We show that our results are robust to more than two periods by looking at a three period model. We find equilibria that are similar to those described in our base model, and show that low-quality firms harvest ratings and mimic prices with strictly positive probability in every non-terminal period. We show that, in line with evidence by Cabral and Hortaçsu (2010) and Jin and Kato (2006), low-quality firms are less likely to sustain a good rating. For details, see Web Appendix B.1.5.

7 Related Literature

Our work connects to different literatures. Our key novelty which we have not seen anywhere, is that we study how firms price their products to freeride on the reputation of others. Based

³¹Fradkin and Holtz (2022) show that paying guests of hosts without ratings to rate leads to more negative reviews. In line with this, consumers who do not rate are more likely to have had a low-quality product, so encouraging them to rate when negative ratings are possible leads to more negative ratings.

³²For example, Nosko and Tadelis (2015) find that more than 99% of ratings on ebay are positive.

on this mechanism, we derive novel predictions for how features of the rating environment affect surplus, and how well ratings inform about product quality. These predictions help shed new light on the impact of recent changes on online marketplaces, such as facilitating ratings, and quality controls.

We connect to the wider theoretical literature on **trust and information transmission in the digital economy**. Platforms may recommend products (Hagiü & Jullien, 2011; Peitz & Sobolev, 2022), shroud additional fees and features of third-party sellers (Johnen & Somogyi, 2022), and marketplaces may have fake reviews (He et al., 2022). We contribute to this literature by studying information transmission via ratings and study how firms can use prices to affect their own ratings.

Many studies have shown that prices have a significant influence on ratings. The evidence highlights two key patterns. First, for a given quality, lower prices induce better ratings (Cai et al., 2014; Carnehl et al., 2021; Li & Hitt, 2010; Luca & Reshef, 2021; Neumann et al., 2018). Second, firms with better ratings charge higher prices in future periods (Cabral & Hortaçsu, 2010; Cabral & Li, 2015; Cai et al., 2014; Carnehl et al., 2021; Ert & Fleischer, 2019; Gutt & Herrmann, 2015; Jin & Kato, 2006; Jolivet et al., 2016; Lewis & Zervas, 2019; Li et al., 2020; Livingston, 2005; Luca & Reshef, 2021; McDonald & Slawson, 2002; Neumann et al., 2018; Proserpio et al., 2018). We explain the pattern of setting low prices to build reputation, and setting higher prices to reap the rewards of reputation in a single framework, capturing this with our mechanism of ‘ratings harvesting’.

We contribute to the **theoretical literature on reputation** (e.g. Cabral (2000), Jullien and Park (2014), Kovbasyuk and Spagnolo (2021), Martin and Shelegia (2021), and Tadelis (1999); see also Bar-Isaac and Tadelis (2008) for a survey), and **word-of-mouth** (Chakraborty et al. (2022)). In existing work usually (i) buyers do not endogenously choose if and how to rate, and (ii) ratings mostly reflect quality, and prices do not affect how consumers rate. While some papers relax some of these assumptions (e.g. Chakraborty et al. (2022) and Martin and Shelegia (2021) relax (i), Sobolev et al. (2021) and Stenzel et al. (2020) relax (ii)), no article seems to feature that buyers choose strategically if and how to rate, and sellers price to freeride on the ratings of others. We capture both of these features, which allows us to capture empirical patterns of prices and ratings we discuss earlier in this section. We are also the first to study how rating design via the effort to rate and average quality affects how informative ratings are in equilibrium.

Martin and Shelegia (2021) study a signalling model where consumers leave a good rating if the product was better than expected. High-quality sellers may mimic low-quality sellers

to overdeliver and get a good rating. In contrast, we study how low-quality firms use prices to improve their ratings, and we provide novel insights about how the design of rating environments leads to more informative ratings.

Some recent papers study the impact of **value-for-money** on ratings. Maybe the first theoretical work on how prices affect ratings is Stenzel et al. (2020). They focus on prices in long-run equilibria where ratings transmit precise information about quality. Instead, we focus on a setting where firms can harvest ratings to freeride on the reputation of other sellers, biasing ratings also on the path of play. Also in Sobolev et al. (2021) ratings may be a noisy signal of quality on the path of play. But their mechanism is very different: they start with the premise that more sales can lead to more or less precise ratings, e.g. because the additional raters might know the products better or worse than existing raters. Instead, we study how firms price their products to freeride on the reputation of others. In addition to both papers, we provide novel insights about how the design of rating environments leads to more informative ratings.

Some researchers argue that consumers should get **paid to rate**. One argument is that sellers should be allowed to pay for feedback: because high-quality firms are more inclined to pay for feedback, feedback is a credible signal for quality (Halliday & Lafky, 2019; Kihlstrom & Riordan, 1984; Milgrom & Roberts, 1986; Nelson, 1974). Others argue that feedback is like a public good that is underprovided (Avery et al., 1999; Bolton et al., 2004; Chen et al., 2010). In contrast, we show that encouraging feedback via ratings (e.g. when marketplaces pay consumers to rate, introduce simpler one-click ratings, or reminding consumers to rate) encourages low-quality firms to harvest ratings, leading possibly to more ratings, but also less-informative ratings. This result is in line with evidence by Cabral and Li (2015) that we discussed above.

We also connect to the ongoing **debate on rating inflation**. Rating inflation describes the observation that rating scores improve over time, and most of the improvements cannot be attributed to product quality (Filippas & Horton, 2022; Filippas et al., 2022; Nosko & Tadelis, 2015; Zervas et al., 2021). We propose multiple channels through which value-for-money leads to more rating harvesting and therefore rating inflation: (i) platforms encouraging consumers to rate, (ii) increased quality controls by platforms, and (iii) increased competition between sellers.

We provide a theoretical explanation for why identical products may get **different ratings across platforms** (Chevalier & Mayzlin, 2006). Some evidence suggests that this is due to user self-selection onto marketplaces (Granados et al., 2012; Raval, 2020). We provide a

complementary explanation and show that differences in features of the rating system (e.g. cost of ratings, aggregate quality, competition between sellers) can lead to different ratings for identical products. Our results align with experimental evidence on how the design of rating systems can influence ratings (Schneider et al., 2021).

We connect to the literature on **consumer information about differentiated** products. A wide range of articles highlight how firms prefer well-informed consumers to amplify product differentiation and relax competition (Anderson & Renault, 2006; Armstrong & Zhou, 2022; Hefti et al., 2022; Johnen & Leung, 2022). Indeed, Armstrong and Zhou (2022) show that firms prefer more-informative information structures than consumers. We show that this intuition may not translate to the context of ratings: When firms harvest ratings, they may prefer less-informative ratings than consumers.

8 Conclusion

Ratings are an essential element of the online economy, building trust between strangers. But ratings are only able to build trust if they are informative about the underlying products and services. In this paper, we study how firms use prices to influence their own ratings. We explore a qualitatively novel trade-off between rating harvesting and price mimicking, which connects well to evidence on the dynamic interaction between prices and ratings. We identify several factors that encourage rating harvesting and lead to less-informative ratings. We also explore implications for buyer and seller surplus.

A key feature of our setting is that consumers cannot distinguish whether a rater left a good rating because she got a product of high quality, or paid a low price. We have several motivations for this feature. First, rating systems are typically coarse: they allow for a finite number of ratings (typically 5 stars), but quality levels and prices are rather continuous variables. So there will be some bunching where sellers with different price-quality combinations will get the same ratings. Second, although platforms typically give consumers easy access to past ratings, they do not connect them to the prices that the raters paid.³³ This feature of many rating environments is a key reason, in our setting, why consumers cannot distinguish whether a given rating is the result of high quality, or a low price. If ratings, however, would reflect purchase prices, consumers may be better able to identify high-quality firms, which could discourage rating harvesting.

³³For example, platforms such as Amazon do not reveal past prices to consumers on their listings, nor do they reveal which price a rater paid. While there exist third party websites such as <https://camelcamelcamel.com/> and <https://keepa.com/> that track past prices on Amazon, such websites do not link prices to actual sales and reviews, so they do not allow users to infer if prices influenced ratings.

In practice, consumers may also read product reviews to learn about product quality. In principle, these reviews could help consumers disentangle how quality and price affect ratings. We argue that reviews do not allow consumers to fully disentangle the effect of past prices on the ratings they observe. First, even if consumers take the time to read reviews, they will only read a small and selected sample of user experiences. Second, even reviews that comment on value-for-money rarely mention the exact price they paid, which makes it impossible to judge relative to which price the value was good.

The key feature of our analysis is that consumers need to obtain a sufficiently high value-for-money to leave a good rating. But consumers may rate for other reasons, e.g. to help other consumers by signalling product quality through ratings, or because they are intrinsically motivated to reveal the true quality of the product. Because prices change over time, however, consumers who rate for such motives will rate only based on quality. Thus, even if some consumers have such other motivations, we would still expect price dynamics resembling rating harvesting, as long as at least some consumers rate based on their value-for-money.

In practice, fake ratings also undermine how informative ratings are. If low-quality firms are somewhat more inclined to acquire fake ratings, fake ratings will also make ratings less-informative of quality. In contrast, firms in our setting use lower prices to get better reviews, which benefits consumers and can explain evidence that lower prices induce better ratings.

References

- Abrate, G., Quinton, S., & Pera, R. (2021). The relationship between price paid and hotel review ratings: Expectancy-disconfirmation or placebo effect? *Tourism Management*, *85*, 104314. <https://doi.org/10.1016/j.tourman.2021.104314>
- Airbnb. (2022). *Why was my listing paused or suspended?* Retrieved April 2021, from <https://www.airbnb.com/help/article/1303/why-was-my-listing-paused-or-suspended>
- Amazon. (2021). *About amazon verified purchase reviews.* Retrieved June 2021, from <https://www.amazon.com/gp/help/customer/display.html?nodeId=202076110>
- Amazon Customer. (2012). *Reviews must contain at least 20 words...* Retrieved April 2021, from <https://www.amazon.com/review/R22R8P0JT7GHWN>
- Ananthakrishnan, U. M., Proserpio, D., & Sharma, S. (2019). I hear you: Does quality improve with customer voice? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3467236>
- Anderson, S. P., & Renault, R. (2006). Advertising content. *The American Economic Review*, *96*(1), 93–113.
- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, *37*(3), 668–691. <https://doi.org/10.1111/j.1756-2171.2006.tb00037.x>
- Armstrong, M., & Zhou, J. (2022). Consumer information and the limits to competition. *American Economic Review*, *112*(2), 534–77.
- Avery, C., Resnick, P., & Zeckhauser, R. (1999). The Market for Evaluations. *American Economic Review*, *89*(3), 564–584. <https://doi.org/10.1257/aer.89.3.564>
- Bajari, P., & Hortacısu, A. (2003). The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions. *The RAND Journal of Economics*, *34*(2), 329–355. <https://doi.org/10.2307/1593721>
- Bar-Isaac, H., & Tadelis, S. (2008). Seller Reputation. *Foundations and Trends® in Microeconomics*, *4*(4), 273–351. <https://doi.org/10.1561/07000000027>
- Bolton, G. E., Katok, E., & Ockenfels, A. (2004). How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation. *Management Science*, *50*(11), 1587–1602. <https://doi.org/10.1287/mnsc.1030.0199>
- Bolton, G. E., & Ockenfels, A. (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, *90*(1), 166–193. <https://doi.org/10.1257/aer.90.1.166>
- Cabral, L. (2000). Stretching Firm and Brand Reputation. *The RAND Journal of Economics*, *31*(4), 658–673. <https://doi.org/10.2307/2696353>

- Cabral, L., & Hortaçsu, A. (2010). The Dynamics of Seller Reputation: Evidence from Ebay*. *The Journal of Industrial Economics*, 58(1), 54–78. <https://doi.org/10.1111/j.1467-6451.2010.00405.x>
- Cabral, L., & Li, L. I. (2015). A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay. *Management Science*, 61(9), 2052–2063. <https://doi.org/10.1287/mnsc.2014.2074>
- Cai, H., Jin, G. Z., Liu, C., & Zhou, L.-a. (2014). Seller reputation: From word-of-mouth to centralized feedback. *International Journal of Industrial Organization*, 34, 51–65. <https://doi.org/10.1016/j.ijindorg.2014.03.002>
- Caillaud, B., & Jullien, B. (2003). Chicken & Egg: Competition among Intermediation Service Providers. *The RAND Journal of Economics*, 34(2), 309–328. <https://doi.org/10.2307/1593720>
- Carnehl, C., Schaefer, M., Stenzel, A., & Tran, K. D. (2021). Value for money and selection: How pricing affects airbnb ratings. *Working Paper*.
- Casner, B. (2020). Seller curation in platforms. *International Journal of Industrial Organization*, 72, 102659. <https://doi.org/10.1016/j.ijindorg.2020.102659>
- Chakraborty, I., Deb, J., & Oery, A. (2022). When do consumers talk? *Available at SSRN 4155523*.
- Chen, Y., Harper, F. M., Konstan, J., & Li, S. X. (2010). Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens. *American Economic Review*, 100(4), 1358–1398. <https://doi.org/10.1257/aer.100.4.1358>
- Chevalier, J. A., & Mayzlin, D. (2006). The effect of word of mouth on sales: Online book reviews. *Journal of marketing research*, 43(3), 345–354.
- Cho, I.-K., & Sobel, J. (1990). Strategic Stability and Uniqueness in Signaling Games. *Journal of Economic Theory*. [https://doi.org/10.1016/0022-0531\(90\)90009-9](https://doi.org/10.1016/0022-0531(90)90009-9)
- Competition and Markets Authority (UK). (2017). *Digital comparison tools market study*.
- Crawford, G., Crémer, J., Dinielli, D., Fletcher, A., Heidhues, P., Luca, M., Salz, T., Schnitzer, M., Scott Morton, F. M., Seim, K., & Sinkinson, M. (2023). Consumer protection for online markets and large digital platforms. *Yale Journal on Regulation*.
- crebel. (2017). *Woohoo! the 20-word minimum review requirement is back*. Retrieved April 2021, from <https://www.kboards.com/threads/woohoo-the-20-word-minimum-review-requirement-is-back.250908/>
- De Langhe, B., Fernbach, P. M., & Lichtenstein, D. R. (2016). Navigating by the Stars: Investigating the Actual and Perceived Validity of Online User Ratings. *Journal of Consumer Research*, 42(6), 817–833. <https://doi.org/10.1093/jcr/ucv047>

- Dellarocas, C., & Wood, C. A. (2008). The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias. *Management Science*, *54*(3), 460–476. <https://doi.org/10.1287/mnsc.1070.0747>
- Dickey, M. R. (2019). Uber will start deactivating riders with low ratings. *TechCrunch*. Retrieved April 2021, from <https://techcrunch.com/2019/05/29/uber-will-start-deactivating-riders-with-low-ratings/>
- Donaker, G., Kim, H., & Luca, M. (2019). Designing better online review systems. *Harvard Business Review*. Retrieved June 2021, from <https://hbr.org/2019/11/designing-better-online-review-systems>
- Dufwenberg, M., & Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior*, *47*(2), 268–298. <https://doi.org/10.1016/j.geb.2003.06.003>
- Ert, E., & Fleischer, A. (2019). The evolution of trust in Airbnb: A case of home rental. *Annals of Tourism Research*, *75*, 279–287. <https://doi.org/10.1016/j.annals.2019.01.004>
- Filippas, A., & Horton, J. J. (2022). Altruism can ruin reputation systems. *Working Paper*.
- Filippas, A., Horton, J. J., & Golden, J. M. (2022). Reputation inflation. *Marketing Science*, *41*, 305–317. <https://doi.org/10.1287/mksc.2022.1350>
- Fradkin, A., & Holtz, D. (2022). Do incentives to review help the market? evidence from a field experiment on airbnb. *American Economic Review: Insights*.
- Granados, N., Gupta, A., & Kauffman, R. J. (2012). Online and offline demand and price elasticities: Evidence from the air travel industry. *Information Systems Research*, *23*(1), 164–181.
- Gutt, D., & Herrmann, P. (2015). Sharing Means Caring? Hosts' Price Reaction to Rating Visibility. *ECIS*, *54*, 14.
- Gutt, D., & Kundisch, D. (2016). Money Talks (Even) in the Sharing Economy: Empirical Evidence for Price Effects in Online Ratings as Quality Signals. *Proceedings of the 37th International Conference on Information Systems (ICIS), Dublin, Ireland*, 10.
- Hagiu, A., & Jullien, B. (2011). Why do intermediaries divert search? *The RAND Journal of Economics*, *42*(2), 337–362. <https://doi.org/10.1111/j.1756-2171.2011.00136.x>
- Halliday, S. D., & Lafky, J. (2019). Reciprocity through ratings: An experimental study of bias in evaluations. *Journal of Behavioral and Experimental Economics*, *83*, 101480. <https://doi.org/10.1016/j.socec.2019.101480>
- He, S., Hollenbeck, B., & Proserpio, D. (2022). The market for fake reviews. *Marketing Science*. <https://doi.org/10.1287/mksc.2022.1353>
- Hefti, A., Liu, S., & Schmutzler, A. (2022). Preferences, confusion and competition. *The Economic Journal*, *132*(645), 1852–1881.

- Hu, N., Zhang, J., & Pavlou, P. A. (2009). Overcoming the J-shaped distribution of product reviews. *Communications of the ACM*, 52(10), 144–147. <https://doi.org/10.1145/1562764.1562800>
- Introducing steam reviews.* (2021). Retrieved June 2021, from <https://store.steampowered.com/reviews/>
- Jin, G. Z., & Kato, A. (2006). Price, quality, and reputation: Evidence from an online field experiment. *The RAND Journal of Economics*, 37(4), 983–1005. <https://doi.org/10.1111/j.1756-2171.2006.tb00067.x>
- Johnen, J., & Leung, B. T. K. (2022). Distracted from comparison: Product design and advertisement with limited attention.
- Johnen, J., & Ng, R. (2024). Harvesting Ratings Web Appendix. https://robinng.com/research/Harvesting_Ratings_Appendix_B.pdf
- Johnen, J., & Somogyi, R. (2022). Deceptive features on platforms. <https://ssrn.com/abstract=3886615>
- Jolivet, G., Jullien, B., & Postel-Vinay, F. (2016). Reputation and prices on the e-market: Evidence from a major French platform. *International Journal of Industrial Organization*, 45, 59–75. <https://doi.org/10.1016/j.ijindorg.2016.01.003>
- Jullien, B., & Park, I.-U. (2014). New, Like New, or Very Good? Reputation and Credibility. *The Review of Economic Studies*, 81(4), 1543–1574. <https://doi.org/10.1093/restud/rdy012>
- Kihlstrom, R. E., & Riordan, M. H. (1984). Advertising as a signal. *Journal of Political Economy*, 92(3), 427–450. <https://doi.org/10.1086/261235>
- Klein, T. J., Lambertz, C., & Stahl, K. O. (2016). Market Transparency, Adverse Selection, and Moral Hazard. *Journal of Political Economy*, 124(6), 1677–1713. <https://doi.org/10.1086/688875>
- Kovbasyuk, S., & Spagnolo, G. (2021). Memory and markets. *Available at SSRN 2756540*.
- Lewis, G., & Zervas, G. (2019). The supply and demand effects of review platforms. *Proceedings of the 2019 ACM Conference on Economics and Computation*, 197. <https://doi.org/10.1145/3328526.3329655>
- Li, L. I., Tadelis, S., & Zhou, X. (2020). Buying reputation as a signal of quality: Evidence from an online marketplace. *The RAND Journal of Economics*, 51(4), 965–988. <https://doi.org/10.1111/1756-2171.12346>
- Li, X., & Hitt, L. M. (2010). Price Effects in Online Product Reviews: An Analytical Model and Empirical Analysis. *MIS Quarterly*, 34(4), 809–831. <https://doi.org/10.2307/25750706>

- Livingston, J. A. (2005). How Valuable Is a Good Reputation? A Sample Selection Model of Internet Auctions. *Review of Economics and Statistics*, 87(3), 453–465. <https://doi.org/10.1162/0034653054638391>
- Luca, M., & Reshef, O. (2021). The effect of price on firm reputation. *Management Science*, 67(7), 4408–4419. <https://doi.org/10.1287/mnsc.2021.4049>
- Martin, S., & Shelegia, S. (2021). Underpromise and overdeliver? - online product reviews and firm pricing. *International Journal of Industrial Organization*, 79, 102775. <https://doi.org/10.1016/j.ijindorg.2021.102775>
- Masters, K. (2021). A short history of amazon’s product review ecosystem. *Forbes*. Retrieved April 2021, from <https://www.forbes.com/sites/kirimasters/2021/03/22/a-short-history-of-amazons-product-review-ecosystem/>
- McDonald, C. G., & Slawson, V. C. (2002). Reputation in an Internet Auction Market. *Economic Inquiry*, 40(4), 633–650. <https://doi.org/10.1093/ei/40.4.633>
- Milgrom, P., & Roberts, J. (1986). Price and advertising signals of product quality. *Journal of political economy*, 94(4), 796–821. <https://doi.org/10.1086/261408>
- Miller, N., Resnick, P., & Zeckhauser, R. (2005). Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51(9), 1359–1373. <https://doi.org/10.1287/mnsc.1050.0379>
- Nelson, P. (1974). Advertising as information. *Journal of political economy*, 82(4), 729–754. <https://doi.org/10.1086/260231>
- Neumann, J., Gutt, D., & Kundisch, D. (2018). A Homeowner’s Guide to Airbnb: Theory and Empirical Evidence for Optimal Pricing Conditional on Online Ratings. *Working Papers Dissertations 43, Paderborn University, Faculty of Business Administration and Economics.*, 23.
- Nosko, C., & Tadelis, S. (2015). The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment. *National Bureau of Economic Research*, w20830. <https://doi.org/10.3386/w20830>
- Peitz, M., & Sobolev, A. (2022). Inflated recommendations. <https://ssrn.com/abstract=4121443>
- Perez, S. (2019). Amazon tests a one-tap review system for product feedback. *TechCrunch*. Retrieved April 2021, from <https://techcrunch.com/2019/09/13/amazon-tests-a-one-tap-review-system-for-product-feedback/>
- Prodius, D. (2020). *Closing for quality control*. Retrieved April 2021, from <https://partner.booking.com/en-gb/community/just-joined-bookingcom/closing-quality-control>

- Proserpio, D., Xu, W., & Zervas, G. (2018). You get what you give: Theory and evidence of reciprocity in the sharing economy. *Quantitative Marketing and Economics*, 16(4), 371–407. <https://doi.org/10.1007/s11129-018-9201-9>
- Rabin, M. (1993). Incorporating Fairness into Game Theory and Economics. *The American Economic Review*, 83(5), 1281–1302.
- Raval, D. (2020). Whose voice do we hear in the marketplace? evidence from consumer complaining behavior. *Marketing Science*, 39(1), 168–187.
- Resnick, P., Zeckhauser, R., Swanson, J., & Lockwood, K. (2006). The value of reputation on eBay: A controlled experiment. *Experimental Economics*, 9(2), 79–101. <https://doi.org/10.1007/s10683-006-4309-2>
- Rey, J. D. (2020). Amazon can't end fake reviews, but its new system might drown them out. *Vox*. Retrieved April 2021, from <https://www.vox.com/recode/2020/2/14/21121209/amazon-fake-reviews-one-tap-star-ratings-seller-feedback>
- Rhodes, A., & Wilson, C. M. (2018). False advertising. *The RAND Journal of Economics*, 49(2), 348–369. <https://doi.org/10.1111/1756-2171.12228>
- Rochet, J.-C., & Tirole, J. (2003). Platform Competition in Two-Sided Markets. *Journal of the European Economic Association*, 1(4), 990–1029. <https://doi.org/10.1162/154247603322493212>
- Rushdie. (2018). *Info: Why was my account suspended, immediately after registration?* Retrieved April 2021, from <https://sellercentral.amazon.com/forums/t/info-why-was-my-account-suspended-immediately-after-registration/415889>
- Schneider, C., Weinmann, M., Mohr, P. N., & vom Brocke, J. (2021). When the Stars Shine Too Bright: The Influence of Multidimensional Ratings on Online Consumer Ratings. *Management Science*, 67(6), 3871–3898. <https://doi.org/10.1287/mnsc.2020.3654>
- Siering, M., Muntermann, J., & Rajagopalan, B. (2018). Explaining and predicting online review helpfulness: The role of content and reviewer-related signals. *Decision Support Systems*, 108, 1–12. <https://doi.org/10.1016/j.dss.2018.01.004>
- Sobolev, A., Stahl, K., Stenzel, A., Wolf, C., et al. (2021). Strategic pricing and ratings.
- Soper, S. (2019). Amazon to boost quality control after selling toxic kid products. *Bloomberg*. Retrieved April 2021, from <https://www.bloomberg.com/news/articles/2019-05-09/amazon-to-boost-quality-control-after-selling-toxic-kid-products>
- Stenzel, A., Wolf, C., & Schmidt, P. (2020). Pricing for the Stars: Dynamic Pricing in the Presence of Rating Systems. *Proceedings of the 21st ACM Conference on Economics and Computation*, 273–274. <https://doi.org/10.1145/3391403.3399522>
- Tadelis, S. (1999). What's in a Name? Reputation as a Tradeable Asset. *The American Economic Review*, 89(3), 548–563. <https://doi.org/10.1257/aer.89.3.548>

- TaobaoTranslate. (2021). *Taobao item review guide*. Retrieved June 2021, from <https://www.taobaotranslate.com/taobao-item-review-rating/taobao-item-rating.html>
- UberLyftDriver. (2017). *Ubereat's deactivation policy for delivery drivers/couriers*. Retrieved April 2021, from <https://www.ridesharingforum.com/t/ubereats-deactivation-policy-for-delivery-drivers-couriers/444>
- Wang, J. (2021). Do birds of a feather flock together? Platform's quality screening and end-users' choices theory and empirical study of online trading platforms. *International Journal of Industrial Organization*, 75, 102698. <https://doi.org/10.1016/j.ijindorg.2020.102698>
- Yelp. (2022). *Yelp elite squad*. Retrieved July 2022, from <https://www.yelp.com/elite>
- Zervas, G., Proserpio, D., & Byers, J. W. (2021). A first look at online reputation on Airbnb, where every stay is above average. *Marketing Letters*, 32(1), 1–16. <https://doi.org/10.1007/s11002-020-09546-4>
- Zodiak. (2021). *What to do if your airbnb listing is suspended*. Retrieved April 2021, from <https://zodiak.co.nz/what-to-do-if-your-airbnb-listing-is-suspended/>

Appendix A Proofs

Proof of Proposition 1

We proceed as follows. First, we pin down equilibrium prices in period 1 in Lemma 2 and equilibrium beliefs in Lemma 3. Afterwards, we use these lemmas to prove the remaining statements in Proposition 1.

Lemma 2. *In equilibrium, firm j plays the price p_{t,R_t}^j in period t , in order to receive the rating R_t .*

In the first period, firms play the following equilibrium prices with positive probability.

- *High-quality firm:* $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\}$
- *Low-quality firm:*

$$p_{1,1}^L = q^L - e$$

$$p_{1,0}^L = E[q_1|R_0, p_{1,1}^H]$$

Proof of Lemma 2.

We proceed in three steps. First, we consider the pricing strategy of the high-quality firms. Second, we look at the pricing strategy of low-quality firms receiving a rating. Third, we focus on prices of low-quality firms obtaining no rating.

To start, we look at the pricing strategy of high-quality firms. We show that their pricing strategy is unique and $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\}$ with probability 1.

Given Selection Assumption 1, we focus on equilibria where the high-quality firm sets a price which allows it to obtain a rating, $R_1 = 1$. Therefore, in equilibrium, the firm does not consider the pricing strategy that obtains no rating. We show that high-quality firms set a unique price in period 1 with probability 1. Suppose towards a contradiction that high-quality firms set more than one price with positive probability.

Without loss of generality, suppose that the high-quality firm sets a distribution of prices, $p \in [p', p'']$ such that $p'' > p'$, and the firm receives a rating $R_1 = 1$ with probability 1 for all $p \in [p', p'']$. Therefore, at all $p \in [p', p'']$, consumers purchase products with probability 1.

Notice that for any price $\hat{p} \in [p', p'']$ such that $\hat{p} > p$, we have $\pi^H(\hat{p}) > \pi^H(p)$. To see this, observe that both prices induce the same demand in period 1, but \hat{p} induces a larger margin and therefore strictly larger profits in period 1. Further, both prices induce $R_1 = 1$ with probability 1, and therefore the same expected profits in period 2. As a result, $\pi^H(\hat{p}) >$

$\pi^H(p)$. Shifting the probability mass of the entire price distribution in period 1 to one mass point, p'' , strictly increases profits for the high-quality firm, contradicting that the firm sets more than one price with positive probability. Essentially, the same argument implies that the high-quality firm does not set more than one price with strictly positive probability. We conclude that the high-quality firm sets a unique price in period 1 with probability 1.

Next, we prove that there exist an upper bound on prices, \overline{p}_t^j for $j \in \{L, H\}$ such that a firm j receives a rating. In order for a firm to induce a rating, the rating utility must be weakly positive, i.e. $v_t \geq 0$. Therefore,

$$[q^j - p_{t,1}^j] - e \geq 0 \Leftrightarrow p_{t,1}^j \leq \overline{p}_t^j \equiv q^j - e.$$

Therefore, the upper bound on prices such that the high-quality firm receives a positive rating is $\overline{p}_1^H = q^H - e$.

Finally, consider that this upper bound is restricted by consumer's beliefs, $E[q_1|R_0, p_{1,1}^H] < \overline{p}_1^H$. Under such scenarios, by Selection Assumption 1, high-quality firms prefer obtaining a rating. This can only be achieved if consumers buy. Therefore, $p_{1,1}^H$ has an upper bound of $E[q_1|R_0, p_{1,1}^H]$.

We next show that $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\}$ with probability 1. To see this, note that $\overline{p}_1^L = q^L - e$ is the cut-off price above which the low-quality firm receives no rating. See also that $\overline{p}_1^H > \overline{p}_1^L$ and $q^L > \overline{p}_1^L$. Thus, because $E[q_1|R_0, p_{1,1}^H] \geq q^L$, the high-quality firm sets its equilibrium price in period 1 strictly above \overline{p}_1^L , i.e. $p_{1,1}^H > \overline{p}_1^L$. By Selection Assumption 2, consumers have the same beliefs for all prices strictly above \overline{p}_1^L , and since $p_{1,1}^H > \overline{p}_1^L$, these beliefs are the correct equilibrium beliefs $E[q_1|R_0, p_{1,1}^H]$. Because consumers have the same beliefs $E[q_1|R_0, p_{1,1}^H]$ for all prices above \overline{p}_1^L , the high-quality firm optimally sets the largest price for which consumers purchase and rate with probability 1, which is $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\}$.

We conclude that high-quality firms set a unique price $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\}$ with probability 1.

We now proceed to the second step and characterize the pricing strategy of low-quality firms who receive a rating. First, we show that the price which it sets and receives a rating is unique. Then, that $p_{1,1}^L = q^L - e$.

Essentially the same argument as used for high-quality firms implies that—conditional on obtaining a rating—the low-quality firm sets a single price with probability 1.

By definition of \overline{p}_1^L , $\overline{p}_1^L = q^L - e$. Since this is strictly less than q^L , and since consumers beliefs

must be weakly above q^L , consumers are always willing to buy at any price weakly below \overline{p}_1^L . Since demand and ratings are the same for all prices weakly below \overline{p}_1^L , a low-quality firm that obtains a rating must optimally set \overline{p}_1^L with probability 1. We conclude that conditional on obtaining a rating—the low-quality firm sets \overline{p}_1^L with probability 1.

We now proceed to the third step of the proof and determine prices of low-quality firms who obtain no rating. We show that low-quality firms who obtain no rating optimally set $E[q_1|R_0, p_{1,1}^H]$. We have shown in step 2 that all prices above \overline{p}_1^L induce the same beliefs $E[q_1|R_0, p_{1,1}^H]$. Thus, low-quality firms who obtain a no rating optimally set the highest possible price, $E[q_1|R_0, p_{1,1}^H]$ with probability 1.

This concludes the proof. □

Lemma 3. *In the first period, consumer's beliefs for each equilibrium price p_1 is given by*

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} & \text{if } p_1 > \overline{p}_1^L \\ q^L & \text{if } p_1 \leq \overline{p}_1^L, \end{cases}$$

and in the second period,

$$E[q_2|R_1] = \begin{cases} \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} & \text{if } R_1 = 1 \\ q^L & \text{if } R_1 = 0. \end{cases}$$

Proof of Lemma 3.

We prove this Lemma by constructing expected quality using Bayes rule. We start by considering the second period, followed by the first period.

Before we begin, note that by Lemma 2, high-quality firms charge $p_{1,1}^H$ and obtain a rating with probability 1, and low-quality firms obtain a rating if and only if they charge a low price $p_{1,1}^L$. We denote the probability δ^* as the probability with which low-quality firms charge $p_{1,0}^L$ in equilibrium. Thus, $(1 - \delta^*)$ is the probability which the low-quality firm obtains a good rating.

Now consider the second period. Given the consumer's information set in the second period, they are aware of historical ratings R_1 , and current prices, p_2 . We now show that expected quality in period 2 is independent of second period prices, p_2 . Since period 2 is the final period, no firm obtains a rating, which, by Selection Assumption 2, implies that the expected quality in period 2 is independent of second period prices. We conclude that expected quality in period 2 only depends on past ratings, R_1 .

Next, we pin down consumers' expectations in period 2. In equilibrium, consumers observe $R_1 = 1$ from all high-quality firms and low-quality firms with low prices, i.e. with probability $\gamma + (1 - \delta^*)(1 - \gamma)$. Because only low-quality firms get no rating, the expected quality after observing $R_1 = 0$ is q^L . Thus, applying Bayes rule leads to

$$E[q_2|R_1] = \begin{cases} \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)} & \text{if } R_1 = 1 \\ q^L & \text{if } R_1 = 0. \end{cases}$$

We now consider the first period. First, recall that $R_0 = 0$.

We distinguish two cases, (i) $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\} = E[q_1|R_0, p_{1,1}^H]$ and (ii) $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\} = q^H - e$.

We start with case (i) and suppose $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\} = E[q_1|R_0, p_{1,1}^H]$. Then by Lemma 2, we have $p_{1,0}^L = p_{1,1}^H$ and $p_{1,1}^L = q^L - e$. Applying Bayes rule leads to

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} & \text{if } p_1 = p_{1,1}^H \\ q^L & \text{if } p_1 = p_{1,1}^L. \end{cases}$$

Now consider case (ii) and suppose $p_{1,1}^H = \min\{q^H - e, E[q_1|R_0, p_{1,1}^H]\} = q^H - e$. Then $p_{1,1}^H \neq p_{1,0}^L$, and Bayes rule implies $E[q_1|R_0, p_{1,1}^H] = q^H$. This is only consistent with the finding in Lemma 2 that $p_{1,0}^L = E[q_1|R_0, p_{1,1}^H] = q^H$ if $\delta^* = 0$, i.e. the low-quality firm sets $p_{1,0}^L$ with probability zero. Thus, beliefs are given again by

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} & \text{if } p_1 = p_{1,1}^H \\ q^L & \text{if } p_1 = p_{1,1}^L, \end{cases}$$

applied at $\delta^* = 0$. We conclude from cases (i) and (ii) that for equilibrium prices in period 1, beliefs are given by

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} & \text{if } p_1 = p_{1,1}^H \\ q^L & \text{if } p_1 = p_{1,1}^L. \end{cases}$$

Because $p_{1,1}^H > \bar{p}_1^L$, $p_{1,0}^L > \bar{p}_1^L$, and $p_{1,1}^L \leq \bar{p}_1^L$, this concludes the proof. \square

We prove a slightly more general statement than **Proposition 1**.

Proposition 5. *All perfect Bayesian equilibria satisfy the following. In period 1:*

1. High-quality firms receive a good rating with probability 1 and charge $\bar{p} \equiv E[q_1 | R_0, p_{1,1}^H]$.
2. Low-quality firms randomize their strategy.

a. They charge \bar{p} and obtain no rating with probability δ^* .

b. They charge $\underline{p} \equiv q^L - e$ and obtain a good rating with probability $1 - \delta^*$.

3. Consumers beliefs of equilibrium prices are given by

$$E[q_1 | p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} & \text{if } p_1 > q^L - e \\ q^L & \text{if } p_1 \leq q^L - e \end{cases}.$$

In period 2:

4. Prices are equal to expected quality conditional on ratings.

$$5. \text{ Consumer beliefs are given by } E[q_2 | R_1] = \begin{cases} \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} & \text{if } R_1 = 1 \\ q^L & \text{if } R_1 = 0 \end{cases}.$$

The equilibrium is unique up to off-equilibrium-path beliefs. $\delta^* = 1$ if and only if $(1-\gamma)(q^H - q^L) \leq e$, and $\delta^* \in (\frac{1}{2}, 1)$ otherwise. The equilibrium exists if $q^H - e \geq \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$.

Proposition 1 in the main text is obtained as a special case.

Proof of Proposition 5.

From Lemma 2 and 3, we have shown statements 2, 3 and 5. Hence, what remains is to prove statement 1 and 4, as well as existence and uniqueness up to off-equilibrium-path beliefs.

To prove statement 1, note that by Lemma 2, we know $p_{1,1}^H = \min\{q^H - e, E[q_1 | R_0, p_{1,1}^H]\}$, and it remains to show that $\min\{q^H - e, E[q_1 | R_0, p_{1,1}^H]\} = E[q_1 | R_0, p_{1,1}^H]$. Suppose towards a contradiction that $\min\{q^H - e, E[q_1 | R_0, p_{1,1}^H]\} = q^H - e$. As we argued in the proof of Lemma 3, we then have $p_1^H \neq p_{1,0}^L$, and $E[q_1 | R_0, p_{1,1}^H] = q^H$, and $p_{1,0}^L = E[q_1 | R_0, p_{1,1}^H] = q^H$ is played with probability $\delta^* = 0$. Because the low-quality firm charges $p_{1,1}^L = q^L - e$ with probability 1, all firms get a rating with probability 1. By Lemma 3, low-quality firms earn up to $q^L - e + \gamma q^H + (1-\gamma)q^L$. Low-quality firms can strictly increase profits by charging a first period price of q^H . By Selection Assumption 2, consumers believe $E[q_1 | R_0, p_1] = q^H$ for all prices $p_1 \geq q^L - e$, so they purchase in period 1. In period 2, the deviation earns q^L . Overall, the deviation earns $q^H + q^L$. Since $q^H \geq \gamma q^H + (1-\gamma)q^L$ and $e > 0$, this deviation profitable for low-quality firms, a contradiction.

We conclude that $p_{1,1}^H = \min\{q^H - e, E[q_1 | R_0, p_{1,1}^H]\} = E[q_1 | R_0, p_{1,1}^H]$, which proves statement 1. Because we have shown that $p_{1,1}^H = E[q_1 | R_0, p_{1,1}^H]$, and this is the same as $p_{1,0}^L$, to simplify

notation, we state that $\bar{p} = p_{1,1}^H = p_{1,0}^L = E[q_1 | R_0, \bar{p}]$. To further simplify notation, we label $\underline{p} = p_{1,1}^L = q^L - e$.

We now prove statement 4. We have shown in Lemma 3 that in period 2, firms are no longer incentivized by future ratings. We have also shown that consumers' beliefs only depend on past ratings. Thus, firms optimally charge prices equal to the expected profits conditional on the past ratings they received. We conclude that in period 2, prices equal expected quality conditional on ratings, which proves statement 4.

We conclude that statements 1 - 5 hold.

We now show that equilibria are unique up to off-equilibrium-path beliefs. To show uniqueness of equilibrium, consider that for some δ^* , low-quality firms are indifferent between getting a rating and no rating. From the proof of statement 1, we know that in equilibrium we must have $\delta^* < 1$ and $q^H - e > E[q_1 | p_{1,1}^H]$, implying that $\bar{p} = \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$.

Charging $\underline{p} = q^L - e$ induces a rating and, given correct equilibrium beliefs, the following is the total profits for the low-quality firm:

$$q^L - e + \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)}. \quad (4)$$

This is strictly increasing in δ^* for all $\gamma \in (0, 1)$ and $q^H > q^L$.

When the low-quality firm charges $\bar{p} = \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$ in period 1, it obtains no rating and earns

$$\frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + q^L, \quad (5)$$

which strictly decreases in δ^* for all $\gamma \in (0, 1)$ and $q^H > q^L$.

To start we show that $\delta^* = 1$ can only be an equilibrium if no mixed-strategy equilibrium exists. Suppose $\delta^* = 1$. Then (4) and (5) become $q^L - e + q^H$ and $\gamma q^H + (1 - \gamma)q^L + q^L$, respectively. For $\delta^* = 1$ to be optimal, it must be the case that

$$q^L - e + q^H \leq \gamma q^H + (1 - \gamma)q^L + q^L. \quad (6)$$

Since (4) strictly increases in δ^* and (5) strictly decreases in δ^* , this means whenever $\delta^* = 1$ is an equilibrium, we cannot have a mixed-strategy equilibrium. We conclude that $\delta^* = 1$

can only be an equilibrium if no mixed-strategy equilibrium exists.

We now show that $\delta^* = 0$ cannot be an equilibrium. Suppose towards a contradiction that $\delta^* = 0$. Then (4) and (5) become $q^L - e + \gamma q^H + (1 - \gamma)q^L$ and $q^H + q^L$, respectively. But since $q^H > \gamma q^H + (1 - \gamma)q^L$ and $q^L > q^L - e$, low-quality firms optimally set $\bar{p} = q^H$ when consumers believe they set this price with probability zero in period 1, a contradiction. We conclude that $\delta^* = 0$ cannot be an equilibrium.

We now characterize the mixed-strategy equilibrium. To have a mixed-strategy equilibrium, consumers must have beliefs such that (4) = (5) and low-quality firms must play some δ^* such that these beliefs are correct. Thus, in a mixed-strategy equilibrium, we have

$$q^L - e + \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)} = \frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + q^L. \quad (7)$$

We have two candidates that solve this equation:

$$\delta^* = \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} \pm \frac{(4\gamma^2(q^H - q^L)^2 + (1 + \gamma)^2 e^2)^{\frac{1}{2}}}{2(1 - \gamma)e}$$

Recall that as a probability, δ^* is bound by 0 and 1.

Consider the scenario where the last term is subtracted.

$$\begin{aligned} & \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} - \frac{(4\gamma^2(q^H - q^L)^2 + (1 + \gamma)^2 e^2)^{\frac{1}{2}}}{2(1 - \gamma)e} \\ & < \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} - \frac{((1 + \gamma)^2 e^2)^{\frac{1}{2}}}{2(1 - \gamma)e} \\ & < \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} - \frac{(1 + \gamma)}{2(1 - \gamma)} \\ & < \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} - \frac{(1 - \gamma)}{2(1 - \gamma)} = -\frac{\gamma(q^H - q^L)}{(1 - \gamma)e} < 0 \end{aligned}$$

. Therefore, we conclude that

$$\delta^* = \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} + \frac{(4\gamma^2(q^H - q^L)^2 + (1 + \gamma)^2 e^2)^{\frac{1}{2}}}{2(1 - \gamma)e}.$$

Further, notice that the optimal δ^* is lies strictly between $\frac{1}{2}$ and 1.

We see from (4) and (5) that $\delta^* < 1$ if $q^L - e + q^H > \gamma q^H + (1 - \gamma)q^L + q^L$. Also note that

$$\delta^* > \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} + \frac{(4\gamma^2(q^H - q^L)^2)^{\frac{1}{2}}}{2(1 - \gamma)e} = \frac{1}{2}.$$

We now show the equilibrium is unique up to off-equilibrium-path beliefs. This follows immediately from having shown that we cannot have $\delta^* = 0$, and that $\delta^* = 1$ can only be an equilibrium if no mixed-strategy equilibrium exists. Additionally, if a $\delta^* \in (0, 1)$ exists such that (7) holds, it must be unique, because (4) strictly increases, and (5) strictly decreases in δ^* . We conclude that if mixed-strategy equilibria exists, it is a unique mixed-strategy, $\delta^* \in (0, 1)$ Thus, we either have a unique pure-strategy equilibrium or a unique mixed-strategy equilibrium, but not both. We conclude that the equilibrium is unique up to off-equilibrium-path beliefs.

Therefore, we can conclude that $\delta^* < 1$ if

$$(1 - \gamma)(q^H - q^L) > e. \quad (8)$$

And this equilibrium is a unique interior solution where $\delta^* \in (\frac{1}{2}, 1)$ up to off-equilibrium-path beliefs. Otherwise, there is a unique corner solution at $\delta^* = 1$ up to off-equilibrium-path beliefs.

$$\delta^* = \begin{cases} \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1 - \gamma)e} + \frac{(4\gamma^2(q^H - q^L)^2 + (1 + \gamma)^2 e^2)^{\frac{1}{2}}}{2(1 - \gamma)e} & \text{if (8) holds} \\ 1 & \text{otherwise} \end{cases} \quad (9)$$

We now show that these equilibria exist.

To start, consider the case where (6) holds. In the candidate equilibrium in period 1, the high-quality firm sets $\bar{p} = \gamma q^H + (1 - \gamma)q^L$ with probability 1 and obtains a rating, and the low-quality firm charges \bar{p} with probability 1 and gets no rating. In period 2, the high-quality firm charges a price equal q^H , and the low-quality firm charges q^L . Consumers' beliefs are as follows. In period 1, they believe

$$E[q_1 | p_1] = \begin{cases} \gamma q^H + (1 - \gamma)q^L & \text{if } p_1 > q^L - e \\ q^L & \text{if } p_1 \leq q^L - e, \end{cases}$$

in the second period, beliefs are independent of prices and are

$$E[q_2|R_1] = \begin{cases} q^H & \text{if } R_1 = 1 \\ q^L & \text{if } R_1 = 0. \end{cases}$$

These beliefs follow Bayes rule on the path of play. The candidate equilibrium is also consistent with our Selection Assumptions. Because high-quality firms obtain a rating with probability 1, the candidate equilibrium is consistent with Selection Assumption 1. Because consumers have the same beliefs for all second period prices, and the same beliefs for all first period prices where the low-quality firm obtains no rating, the candidate equilibrium is consistent with Selection Assumption 2.

We now show that firms have no profitable deviations.

In the candidate equilibrium, the high-quality firm earns $\gamma q^H + (1 - \gamma)q^L + q^H$. Deviations in period 2 to a higher price would induce zero demand, and deviations to lower prices would reduce margins without increasing demand. There are no profitable deviations in period 2. In period 1, deviations to a higher price reduces demand to zero and earns a maximal total profit of $0 + q^L$. Deviations to a lower price in period 1 reduce margins without increasing demand or increasing ratings. Therefore, there is no profitable deviation in period 1. We conclude that high-quality firms have no profitable deviations.

We now show that low-quality firms have no profitable deviations. In the candidate equilibrium they earn $\gamma q^H + (1 - \gamma)q^L + q^L$. Deviations in period 2 to a higher price would induce zero demand, and deviations to lower prices would reduce margins without increasing demand. There are no profitable deviations in period 2. In period 1, deviations to a higher price reduces demand to zero and earns a maximal total profit of $0 + q^L$, which is not a profitable deviation. In period 1, deviations to a lower price above $q^L - e$ does not improve the rating and only reduces margins without raising demand, this is not a profitable deviation. Deviations to lower prices below $q^L - e$ leads to profits weakly below $q^L - e + q^H$, which is not a profitable deviation since (6) holds.

We conclude that if (6) holds, no profitable deviations exist for either type of firm.

Finally, we have shown above that $\bar{p} = E[q_1|\bar{p}]$, which requires $\gamma q^H + (1 - \gamma)q^L \leq q^H - e$.

We conclude that if $(1 - \gamma)(q^H - q^L) \leq e$ and $\gamma q^H + (1 - \gamma)q^L \leq q^H - e$, the candidate equilibrium exists. As we have shown above, it must be the unique equilibrium up to off-equilibrium beliefs.

Now consider the case where (8) holds. We have shown above that no pure-strategy equilibrium exists in this case. In the candidate equilibrium in period 1, the high-quality firm sets $\bar{p} = \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$ with probability 1 and obtains a rating. The low-quality firm charges \bar{p} with probability δ^* and gets no rating, and sets $\underline{p} = q^L - e$ with probability $1 - \delta^*$ and gets a rating. In period 2, all firms with a good rating charge $\frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)}$, and firms without a rating charge q^L . Consumers' beliefs are as follows. In period 1, they believe

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} & \text{if } p_1 > q^L - e \\ q^L & \text{if } p_1 \leq q^L - e. \end{cases}$$

in the second period, beliefs are independent of prices and are

$$E[q_2|R_1] = \begin{cases} \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} & \text{if } R_1 = 1 \\ q^L & \text{if } R_1 = 0. \end{cases}$$

These beliefs follow Bayes rule on the path of play. The candidate equilibrium is also consistent with our Selection Assumptions. Because high-quality firms obtain a rating with probability 1, the candidate equilibrium is consistent with Selection Assumption 1. Because second period beliefs are independent of prices, and consumers have the same beliefs for all first period prices where the low-quality firm obtains no rating, the candidate equilibrium is consistent with Selection Assumption 2.

We now show that firms have no profitable deviations.

We start with low-quality firms, who earn total profits $q^L - e + \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)}$. The firm is indifferent between charging \bar{p} and \underline{p} , where $\bar{p} > \underline{p}$. If the firm deviates to a price above \bar{p} , demand drops to zero and total profits are weakly below $0 + q^L$, this is not a profitable deviation. Deviations to a price $p_1 \in (\underline{p}, \bar{p})$, for which the firm gets the same rating as charging \bar{p} and therefore earns the same profit in period 2, but the firm earns a lower margin than when it charges \bar{p} without increasing demand in period 1, this is not a profitable deviation. Deviations to a price below \underline{p} lead to the same rating as when charging \underline{p} and therefore the same continuation profits, but decrease margins in period 1 without increasing demand, this is not a profitable deviation. In period 2, the low-quality firm extracts expected total surplus conditional on the rating, and cannot strictly increase profits. We conclude that low-quality firms have no profitable deviation.

We now show that high-quality firms have no profitable deviation. In the candidate equilibrium, high-quality firms earn $\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} + \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)}$. In the second period, high-

quality firms extract total surplus conditional on their good rating and therefore cannot profitably deviate. In the first period, a higher price reduces demand to zero and earns profits weakly below $0 + q^L$, which is not a profitable deviation. Deviating to a lower price does not improve the rating and therefore does not increase continuation profits, but reduces margins in period 1 without increasing demand, this is also not a profitable deviation. We conclude that high-quality firms have no profitable deviation.

We conclude that no firm has a profitable deviation.

Finally, we need to check $\bar{p} = \min\{q^H - e, E[q_1|\bar{p}]\} = E[q_1|\bar{p}]$, which requires $\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} \leq q^H - e$.

We conclude that if $(1-\gamma)(q^H - q^L) > e$ and $\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} \leq q^H - e$, the candidate equilibrium exists. As we have shown above, it must be the unique equilibrium up to off-equilibrium beliefs.

This concludes the proof. □

Proof of Corollary 1

Proof of Corollary 1.

In this proof, we show that $\frac{\partial \delta^*}{\partial e} > 0$ if $(1-\gamma)(q^H - q^L) > e$ holds.

We know from Proposition 1 that $\delta^* = \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1-\gamma)e} + \frac{(4\gamma^2(q^H - q^L)^2 + (1+\gamma)^2 e^2)^{\frac{1}{2}}}{2(1-\gamma)e}$, and that $\delta^* \in (\frac{1}{2}, 1)$ when $(1-\gamma)(q^H - q^L) > e$. We can then calculate the derivative with respect to e , which leads to:

$$\frac{\partial \delta^*}{\partial e} = \frac{\gamma(q^H - q^L)((2\gamma(q^H - q^L))^2 + ((1+\gamma)e)^2)^{\frac{1}{2}} - 2\gamma(q^H - q^L)}{\Delta(1-\gamma)e^2((2\gamma(q^H - q^L))^2 + ((1+\gamma)e)^2)^{\frac{1}{2}}} > 0$$

Thus we have shown that when low-quality firms would play a mixed-strategy, as the opportunity cost of rating increases, low-quality firms are less likely to participate in ratings harvesting.

This concludes the proof. □

Proof of Proposition 2

Proof of Proposition 2.

We show $\frac{\partial \delta^*}{\partial \gamma} \leq 0$ when γ is sufficiently small and $\frac{\partial \delta^*}{\partial \gamma} > 0$ when γ is sufficiently large. We then characterise this switching point, $\bar{\gamma}$, and show its uniqueness and existence.

To start, note from (7), δ^* is such that low-quality firms are indifferent between charging low prices and high prices,

$$\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} + q^L = q^L - e + \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)}.$$

Taking the derivative of the right-hand side leads to

$$\frac{\partial RHS}{\partial \gamma} = \frac{(q^H - q^L)((1-\delta^*) + (1-\gamma)\gamma\frac{\partial \delta^*}{\partial \gamma})}{(\gamma + (1-\delta^*)(1-\gamma))^2},$$

and for the left-hand side

$$\frac{\partial LHS}{\partial \gamma} = \frac{(q^H - q^L)(\delta^* - (1-\gamma)\gamma\frac{\partial \delta^*}{\partial \gamma})}{(\gamma + (1-\gamma)\delta^*)^2}.$$

Since both derivatives must be equal in equilibrium, we get

$$\frac{\partial \delta^*}{\partial \gamma} = \frac{(2\delta^* - 1)(\gamma^2 + \delta^{*2}(1-\gamma)^2 - \delta^*(1-\gamma)^2)}{\gamma(1-\gamma)((\gamma + (1-\delta^*)(1-\gamma))^2 + (\gamma + (1-\gamma)\delta^*)^2)} \quad (10)$$

In particular, we show that the sign of $\frac{\partial \delta^*}{\partial \gamma}$ switches at some point $\bar{\gamma}$. Since the denominator is strictly positive, the expression is negative if and only if the numerator is negative, i.e.

$$(2\delta^* - 1)(\gamma^2 - (1-\gamma)^2\delta^*(1-\delta^*)) < 0.$$

And positive when

$$(2\delta^* - 1)(\gamma^2 - (1-\gamma)^2\delta^*(1-\delta^*)) > 0,$$

and 0 if

$$(2\delta^* - 1)(\gamma^2 - (1-\gamma)^2\delta^*(1-\delta^*)) = 0.$$

Since $\delta^* > 0.5$, the when evaluating the sign of $(2\delta^* - 1)(\gamma^2 - (1-\gamma)^2\delta^*(1-\delta^*))$, we only need consider $(\gamma^2 - (1-\gamma)^2\delta^*(1-\delta^*))$

We show that there is only one γ such that $(2\delta^* - 1)(\gamma^2 - (1-\gamma)^2\delta^*(1-\delta^*)) = 0$, therefore there is a unique switching point.

$$\bar{\gamma} = \frac{(q^H - q^L)^2 - e^2}{3(q^H - q^L)^2 + e^2} \quad (11)$$

We show that $\bar{\gamma} \in (0, 1)$. The denominator is strictly positive. From (1), the numerator is strictly positive. Hence, $\bar{\gamma} > 0$. Moreover, because the denominator is always strictly larger than the numerator, $\bar{\gamma} < 1$.

Finally, we show that if $\gamma < \bar{\gamma}$, then $\frac{\partial \delta^*}{\partial \gamma} < 0$, and $\frac{\partial \delta^*}{\partial \gamma} > 0$ when $\gamma \in (\bar{\gamma}, 1)$.

Notice that when $\gamma \rightarrow 0$, and from (9) $\delta^* \rightarrow 1$. Thus $\delta^*(1 - \delta^*) \rightarrow 0$ and $\frac{\gamma^2}{(1-\gamma)^2} \rightarrow 0$. And when $\gamma \rightarrow 1$, $\delta^* = 1$. Further, when $\gamma = 1$, $\frac{\gamma^2}{(1-\gamma)^2} = \infty$.

We show that $\delta^*(1 - \delta^*) > \frac{\gamma^2}{(1-\gamma)^2}$ when γ is sufficiently small. Consider that $\frac{\gamma^2}{(1-\gamma)^2}$ is strictly convex, with $\frac{\partial \frac{\gamma^2}{(1-\gamma)^2}}{\partial \gamma} = \frac{2\gamma}{(1-\gamma)^3} > 0$ and $\frac{\partial^2 \frac{\gamma^2}{(1-\gamma)^2}}{\partial \gamma^2} = \frac{2+4\gamma}{(1-\gamma)^4} > 0$. Evaluated at $\gamma \rightarrow 0$, $\frac{\partial \frac{\gamma^2}{(1-\gamma)^2}}{\partial \gamma} = \frac{2\gamma}{(1-\gamma)^3} \rightarrow 0$. Further, consider that $\frac{\partial \delta^*(1-\delta^*)}{\partial \gamma} = \frac{\partial \delta^*}{\partial \gamma}(1 - 2\delta^*)$. Evaluated at $\gamma \rightarrow 0$, $\delta^* \rightarrow 1$ and $\frac{\partial \delta^*}{\partial \gamma} < 0$. Then, $\frac{\partial \delta^*(1-\delta^*)}{\partial \gamma} = \frac{\partial \delta^*}{\partial \gamma}(1 - 2\delta^*) > 0$. Hence, there exists a range of $\gamma \in (0, 1)$ where $\delta^*(1 - \delta^*) > \frac{\gamma^2}{(1-\gamma)^2}$ is satisfied. We label the first upper boundary (closest to 0) of this range as γ' , such that $\gamma' \in (0, 1)$ and when $\gamma \in (0, \gamma')$, $\delta^*(1 - \delta^*) > \frac{\gamma^2}{(1-\gamma)^2}$ is satisfied and $\frac{\partial \delta^*}{\partial \gamma} < 0$.

Additionally, when $\gamma \rightarrow 1$, $\delta^* \rightarrow 1$ and $\frac{\gamma^2}{(1-\gamma)^2} \rightarrow \infty$. This implies that $\frac{\gamma^2}{(1-\gamma)^2} > \delta^*(1 - \delta^*)$ for some range of γ and $\frac{\partial \delta^*}{\partial \gamma} > 0$. This implies that at $\gamma \rightarrow 1$, $\frac{\partial \delta^*(1-\delta^*)}{\partial \gamma} = \frac{\partial \delta^*}{\partial \gamma}(1 - 2\delta^*) < 0$. Therefore, we can conclude that there exists a range of $\gamma \in (0, 1)$ where $\frac{\gamma^2}{(1-\gamma)^2} > \delta^*(1 - \delta^*)$ is satisfied. We label the first lower boundary (closest to 1) of this range as γ'' , such that $\gamma'' \in (0, 1)$ and when $\gamma \in (\gamma'', 1)$, $\frac{\gamma^2}{(1-\gamma)^2} > \delta^*(1 - \delta^*)$ is satisfied and $\frac{\partial \delta^*}{\partial \gamma} > 0$.

Notice that by definition $\gamma'' \geq \gamma'$. We have shown before that there is a unique switching point, therefore $\bar{\gamma} = \gamma'' = \gamma'$.

We conclude that there is a unique switching point $\bar{\gamma}$, below which $\frac{\partial \delta^*}{\partial \gamma} < 0$, and above which $\frac{\partial \delta^*}{\partial \gamma} > 0$. Where $\frac{\partial \delta^*}{\partial \gamma} = 0$ only if $\gamma = \bar{\gamma}$, and $\bar{\gamma} \in (0, 1)$. \square

Proof of Proposition 3

Proof of Proposition 3.

In this proof, we show that if a social planner concerned with the welfare of consumers would set an optimal e , e^{cs} .

First, we find consumer surplus. This is the sum of the difference between actual price and

quality that consumers receive in each period.

$$\begin{aligned}
CS_1 &= \gamma \left[q^H - \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} \right] + (1-\delta^*)(1-\gamma)e + \delta^*(1-\gamma) \left[q^L - \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} \right] \\
&= (1-\delta^*)(1-\gamma)e \\
CS_2 &= \gamma \left[q^H - \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} \right] + (1-\delta^*)(1-\gamma) \left[q^L - \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} \right] + \\
&\quad \delta^*(1-\gamma)[q^L - q^L] = 0.
\end{aligned}$$

Therefore, consumer surplus arises only from the low-quality firm's attempt to receive a good rating, and total consumer surplus in our model is given by

$$CS = (1-\delta^*)(1-\gamma)e. \quad (12)$$

This is (2). From (2) we can evaluate the effects of changes to effort cost.

First, we show that consumer surplus is concave in e .

$$\frac{\partial CS}{\partial e} = \frac{1}{2} \left[1 - \gamma - \frac{(1+\gamma)^2 e}{\sqrt{4\gamma^2(q^H - q^L)^2 + (1+\gamma)^2 e^2}} \right]$$

This is positive if and only if

$$(1-\gamma)^2 \gamma (q^H - q^L)^2 > (1+\gamma)^2 e^2$$

, and $\frac{\partial CS}{\partial e} = 0$ with equality. This condition is (3).

We now look at the second derivative.

$$\frac{\partial^2 CS}{\partial e^2} = -\frac{2\gamma^2(1+\gamma)^2(q^H - q^L)^2}{\Delta^2((2\gamma(q^H - q^L))^2 + ((1+\gamma)e)^2)^{\frac{3}{2}}} < 0$$

Second, we solve for the optimal level of effort required to leave a rating.

$$e = \pm \frac{(q^H - q^L)(1-\gamma)\sqrt{\gamma}}{1+\gamma}$$

We may reject the negative as we assume that $e \geq 0$. Therefore,

$$e^{cs} = \frac{(q^H - q^L)(1-\gamma)\sqrt{\gamma}}{1+\gamma}$$

where e^{cs} is the level of effort cost that maximises consumer surplus. This e^{cs} is indeed positive when (3) holds. Which is the restriction required for $\frac{\partial CS}{\partial e} > 0$ at $e = 0$.

When (3) does not hold, then $\frac{\partial CS}{\partial e} < 0$ and this implies that $e^{cs} = 0$, the lower bound.

Therefore, for a planner maximizing the welfare of consumers,

$$e^{cs} = \frac{(q^H - q^L)(1 - \gamma)\sqrt{\gamma}}{1 + \gamma} \text{ when (3) holds, and 0 otherwise.}$$

This concludes the proof. □

Proof of Corollary 2

Proof of Corollary 2.

In this proof, we show that in expectation, sellers prefer $e = 0$, and thus completely uninformative ratings. Because high-quality firms prefer more-informative ratings, we argue that the expected sellers' preference for uninformative ratings is driven by low-quality firms.

First, we look at the average profit function of the firm,

$$\begin{aligned} \pi &= \gamma \left[\frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)} \right] + \\ & (1 - \gamma) \left[(1 - \delta^*) \left[q^L - e + \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)} \right] + \delta^* \left[\frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + q^L \right] \right] \\ & = 2[\gamma q^H + (1 - \gamma)q^L] - (1 - \delta^*)e \end{aligned}$$

Taking the derivative to e ,

$$\begin{aligned} \frac{\partial \pi}{\partial e} &= \frac{\partial \delta^*}{\partial e} e - \frac{(1 - \gamma)(1 - \delta^*)}{\Delta} \\ &= -\frac{1}{2} + \frac{(1 + \gamma)^2 e}{2(1 - \gamma)\sqrt{4\gamma^2(q^H - q^L)^2 + (1 + \gamma)^2 e^2}} \end{aligned}$$

and this is negative when (3) holds.

Therefore, on average firms prefer the smallest level of e , and the level of effort cost that maximises firm's profit is $e^s = 0$.

Second, we show that high-quality firms prefer informative ratings.

$$\begin{aligned}\pi^H &= \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} + \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} \\ \frac{\partial \pi^H}{\partial e} &= -\frac{(1-\gamma)^2 \gamma (1+\gamma)(q^H - q^L)(1-2\delta^*) \frac{\partial \delta^*}{\partial e}}{(\gamma + (1-\gamma)(1-\delta^*))^2 (\gamma + (1-\gamma)\delta^*)^2}\end{aligned}$$

Recall that $\delta^* \in (\frac{1}{2}, 1)$, and $\frac{\partial \delta^*}{\partial e} > 0$ when (1) holds. Therefore, $\frac{\partial \pi^H}{\partial e} > 0$.

This means that when low-quality firms would play a mixed-strategy, more-informative ratings benefits high-quality firms.

We can show that (3) is a stricter condition than (1).

From (3),

$$\begin{aligned}(1-\gamma)^2 \gamma (q^H - q^L)^2 &\geq (1+\gamma)^2 e^2 \\ \frac{\gamma}{(1+\gamma)^2} &\geq \frac{e^2}{(1-\gamma)^2 (q^H - q^L)^2}\end{aligned}$$

And from (1),

$$\begin{aligned}(1-\gamma)(q^H - q^L) &> e \text{ since both sides are positive,} \\ (1-\gamma)^2 (q^H - q^L)^2 &> e^2 \\ 1 &> \frac{e^2}{(1-\gamma)^2 (q^H - q^L)^2}\end{aligned}$$

Finally, notice that $\gamma \in (0, 1)$ and thus $\frac{\gamma}{(1+\gamma)^2} < 1$. Therefore, (3) is a stricter condition than (1).

We conclude that more-informative ratings leads to an decrease in the average profits of the firm. This decrease in profits is driven by low-quality firms, and high-quality firms benefit from more-informative ratings environments. \square

Proof of Proposition 4

Proof of Proposition 4.

We show that there is some γ^{cs} that maximises consumer surplus, and that consumer surplus is concave in γ .

To start, recall from Proposition 3 that total consumer surplus reduces to (2), i.e. $CS = (1-\delta^*)(1-\gamma)e$. We take the derivative of (2) with respect to γ and show that there is some

$\gamma^{cs} < \bar{\gamma}$ that maximises consumer surplus.

$$\frac{\partial CS}{\partial \gamma} = -(1 - \delta^* + (1 - \gamma) \frac{\partial \delta^*}{\partial \gamma}) e \quad (13)$$

Since $e > 0$, the sign of $\frac{\partial CS}{\partial \gamma}$ depends on $-(1 - \delta^* + (1 - \gamma) \frac{\partial \delta^*}{\partial \gamma})$.

We solve for γ^{cs} , the optimal choice of γ that maximises consumer surplus.

$$\gamma^{cs} = \frac{-e^2}{4(q^H - q^L)^2 + e^2} \pm \frac{((q^H - q^L)e(2(q^H - q^L) - e)^2)^{\frac{1}{2}}}{4(q^H - q^L)^2 + e^2}$$

Because the denominator is positive, we reject the negative, as $\gamma \in (0, 1)$. Therefore, we show that $\gamma^{cs} \in (0, 1)$ exists if $(q^H - q^L)(2(q^H - q^L) - e)^2 \geq e^3$.

We now show that $(q^H - q^L)(2(q^H - q^L) - e)^2 \geq e^3$ always holds when the low-quality firm plays a mixed strategy. Recall that (1) requires $(1 - \gamma)(q^H - q^L) > e$ for the low-quality firm to play a mixed strategy. Further, since $q^H - q^L > (1 - \gamma)(q^H - q^L)$, then

$$\begin{aligned} q^H - q^L &> e \\ 2(q^H - q^L) - e &> e \end{aligned}$$

Using this,

$$\begin{aligned} (q^H - q^L)(2(q^H - q^L) - e)^2 &\geq e^3 \\ (q^H - q^L)e^2 &> e^3 \\ q^H - q^L &> e. \end{aligned}$$

This provides a condition that is weaker than (1). Therefore, we conclude that γ^{cs} exists when low-quality firms play a mixed-strategy.

We now show that $\frac{\partial CS}{\partial \gamma}$ is strictly concave. First, we shall argue that $\gamma^{cs} < \bar{\gamma}$. Second, we show that when $\gamma > \gamma^{cs}$, $\frac{\partial CS}{\partial \gamma} < 0$. Finally, we show that $\frac{\partial CS}{\partial \gamma} > 0$ when $\gamma < \gamma^{cs}$.

Since $\bar{\gamma}$ solves $\frac{\partial \delta^*}{\partial \gamma}$ implies $\frac{\partial \delta^*}{\partial \gamma} = 0$ at $\bar{\gamma}$. Further, $\delta^* \in (\frac{1}{2}, 1)$ for $\gamma > 0$, implying that

$(1 - \delta^*) > 0$. Therefore, from (13), at $\bar{\gamma}$, $\frac{\partial CS}{\partial \gamma} < 0$. This implies that $\gamma^{cs} < \bar{\gamma}$.

Second, we show that when $\gamma > \gamma^{cs}$, $\frac{\partial CS}{\partial \gamma} < 0$. When $\gamma > \bar{\gamma}$, we know that $(1 - \delta^*) > 0$, $(1 - \gamma) > 0$, and from Proposition 2 $\frac{\partial \delta^*}{\partial \gamma} > 0$. This implies that $\frac{\partial CS}{\partial \gamma} < 0$ when $\gamma > \gamma^{cs}$.

Finally, we show that when $\gamma < \gamma^{cs}$, $\frac{\partial CS}{\partial \gamma} > 0$. Notice that $(1 - \delta^* + (1 - \gamma)\frac{\partial \delta^*}{\partial \gamma}) \rightarrow 1 - \frac{q^H - q^L}{e}$ when evaluated at $\gamma \rightarrow 0$. From (1) we know that this is strictly negative. Therefore, $\frac{\partial CS}{\partial \gamma} > 0$ at $\gamma \rightarrow 0$. And we have shown that there is only one point where the sign changes when $\gamma > 0$. Therefore, when $\gamma < \gamma^{cs}$, $\frac{\partial CS}{\partial \gamma} > 0$.

This concludes the proof and we have shown that there exist some $\gamma^{cs} \in (0, \bar{\gamma})$ that maximises consumer surplus,

$$\gamma^{cs} = \frac{-e^2}{4(q^H - q^L)^2 + e^2} + \frac{((q^H - q^L)e(2(q^H - q^L) - e)^2)^{\frac{1}{2}}}{4(q^H - q^L)^2 + e^2}$$

and consumer surplus is strictly concave in γ .

This concludes the proof. □

Proof of Lemma 1

Proof of Lemma 1.

To show that sellers of all types benefit from quality controls, we look at the lifetime profits of both high- and low-quality firms independently.

We begin with low-quality firms.

$$\pi^L = \delta^* \left(\frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + q^L \right) + (1 - \delta^*)(q^L - e + \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)})$$

We know that δ^* is such that the low-quality firm is indifferent between $\frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + q^L$ and $q^L - e + \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)}$.

Therefore,

$$\begin{aligned} \pi^L &= \frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + q^L \\ \frac{\partial \pi^L}{\partial \gamma} &= \frac{(q^H - q^L)(\delta^* - \gamma(1 - \gamma)\frac{\partial \delta^*}{\partial \gamma})}{(\gamma + \delta^*(1 - \gamma))^2} \end{aligned}$$

Since the denominator is positive, and $q^H - q^L > 0$, we evaluate $\delta^* - \gamma(1 - \gamma)\frac{\partial\delta^*}{\partial\gamma}$. $\delta^* - \gamma(1 - \gamma)\frac{\partial\delta^*}{\partial\gamma} = \frac{(1+\gamma)e}{2\sqrt{4\gamma^2(q^H - q^L)^2 + (1+\gamma)^2e^2}} + \frac{1}{2} > 0$ Therefore, profits of low-quality firms benefits from quality controls.

Moreover, we show that high-quality firms also benefit from quality controls.

$$\begin{aligned}\pi^H &= \frac{\gamma q^H + \delta^*(1 - \gamma)q^L}{\gamma + \delta^*(1 - \gamma)} + \frac{\gamma q^H + (1 - \delta^*)(1 - \gamma)q^L}{\gamma + (1 - \delta^*)(1 - \gamma)} \\ \frac{\partial\pi^H}{\partial\gamma} &= \frac{32\gamma^2(q^H - q^L)^3e^2\sqrt{4\gamma^2(q^H - q^L)^2 + (1 + \gamma)^2e^2}}{(\gamma + \delta^*(1 - \gamma))^2(\gamma + (1 - \delta^*)(1 - \gamma))^2} - \\ &\quad \frac{32\gamma^2(q^H - q^L)^3e^2(2\gamma(q^H - q^L))}{(\gamma + \delta^*(1 - \gamma))^2(\gamma + (1 - \delta^*)(1 - \gamma))^2}\end{aligned}$$

Because $\sqrt{4\gamma^2(q^H - q^L)^2 + (1 + \gamma)^2e^2} - 2\gamma(q^H - q^L) > 0$, it is immediate that $\frac{\partial\pi^H}{\partial\gamma} > 0$.

Therefore, both high- and low-quality firms benefit when platforms implements quality controls.

This concludes the proof. □