Université catholique de Louvain
Faculté des Science Économiques, Sociales, Politiques et de Communication
Département des Sciences Economiques
Center for Operations Research and Econometrics

Online Public Goods

Doctoral Dissertation presented by

Robin Ng

in fulfillment of the requirements
for Doctoral degree in Economics and Management

Thesis committee:

Paul Belleflamme (Supervisor/Secretary)	Université catholique de Louvain
Johannes Johnen (Supervisor)	Université catholique de Louvain
Vincent Vannetelbosch (Chair)	Université catholique de Louvain
Heiner Schumacher (External Evaluator)	Universität Innsbruck
Greg Taylor (External Evaluator)	University of Oxford

September 2023

In the fond memory of those who left too soon:

Patrick Tay Teck Kian
the memory of my childhood

张娘惜
the cornerstone whom I miss daily

Gillian Ng
the godmother I never had

Acknowledgements[1]

It is important to me to first acknowledge *Paul Belleflamme* and *Johannes Johnen*. It is impossible to overstate their tremendous patience, positivity, encouragement, endless comments, and constant reminders. This dissertation would not have been possible without the chance they took on me. I hope I have not disappointed them.

I am tremendously grateful for the mentorship and guidance of *Greg Taylor*,
Heiner Schumacher, and especially *Jonathan Lafky*.

I have also been fortunate to have had the hospitality of *Bernard Caillaud* and *Greg*. They provided the opportunity to meet and learn from countless fascinating academics[2], and to befriend the most amazing people[3]. 🚄

From my alma mater, my gratitude to *Prof Chow* for thinking of me, and for being a continued source of inspiration in my teaching. I also appreciate the continued guidance of *Prof Gea*, *Prof Yichong*, and *Aljoscha Janssen*. 😇

In Louvain, I have met the sweetest people. It is perhaps apt to mention first the first person I spoke with, *Catherine*, with whom I share a birthday, for I would not have survived without her initial assistance. 🇧🇪

To the many others:

Matteo, *Niran*, *Sefane*, *Tiziano*, *Fabrizio*, *Quentin*, *Erika*, *Arnaud*, *Thomas*, ...
My friends, you have blessed me with the comfort of home away from home.

Thank you. 🥰

For sharing the tears of laughter, *Muxin*, *Thomas*, *Antoine*, and *Jacques*. 🤣

Philippe Chevalier, *Alain Strowel*, *Vincent Vannetelbosch*,
Especially *Anne-Lise Sibony*, and *Amma Panin*:
My sincerest gratitude for your time and encouragement.

To the administrative team, particularly *Monique*, *Raphael*, *Nancy*, *Margaux*, *Marie*:
Thank you for being so open and helpful, ensuring everything ran smoothly.

I will especially miss the random conversations with *François Maniquet*. 🥲

It is tough to express 12% of my life thus far in one page. The list of others I would like to mention is long, and possibly endless.[4] I think the best way to sum it up is to mention *Boon Han* for reminding me to **pass it on**.

Matthew, *Rachel*, *Shuyao*, *Shenxi*, *Manpreet*, *Julius*, *Vivien*:
Words cannot describe how much I cherish our friendship.
Instead, I offer you an egg. 🥚

Finally, while it is plainly obvious, it must be stated:
Nothing would have been possible without the support and patience of my family.

To my mother and brother:

Yes, it is over.

(well... sort of 😁)

[1]This dissertation was made possible with financial support from the FWB ARC grant 'PROSEco', the FNRS ASP grant, my father, the SMU/SSRC GRF, and numerous other travel awards.

[2]Francis, Nikhil, Gabrielle, Nicolas, Lionel, Vili, Johannes, Alexei, Dan, ...

[3]Eric, Damien, Nakita, Lennart, Jacob, Manuel, Chris, Tim, Jorge, Jack, Mac, Steven, David, ...

[4]Sebastian, John, Martin, Julian, Tat-How, Greg, Costas, Christian, Jana, Nisvan, the many Pauls, ...

Contents

Introduction

This dissertation aims to understand why and how people participate in online communities. In most online communities, individuals are anonymous to each other. Despite their anonymity, they continue to interact and even help each other. One such interaction is the contribution to and use of public goods. I will focus on two such public goods: (i) online ratings and (ii) open-source software development.

Online Ratings

We often rely on the experiences of others to make everyday decisions, from where to eat, what shows to watch and which hotels to stay in. Where we have no experience with a product, recommendations provide a signal for quality, helping us to choose what to buy. Suggestions from others become especially useful when we cannot inspect a product, such as on the Internet. Hence, ratings have been introduced as a recommendation system to guide consumers' choices. But are ratings able to inform us of product quality?

The answer to this question is not obvious and depends on two things. First, what motivates people to rate? Second, do consumers correctly anticipate what is being rated?

Why do people rate? The literature purports two reasons for why people rate. First, individuals may provide ratings as an expression of the utility they receive. Self-expression means raters rate based on the value-for-money of a product. Second, individuals may choose to rate out of kindness to future consumers. Kindness to future consumers implies that raters select ratings that help improve the consumption decisions of others.

The first chapter investigates what motivates value-for-money rating. Value-for-money means raters leave better ratings for products when there is a combination of higher quality and lower prices.

As an illustration, consider the act of rating a hotel. When rating a hotel, the comfort of the room and interactions with staff likely play an important role in the decision to rate. These are measures of the quality of the hotel. However, prices are just as likely to factor into this decision. If we book a cheap hotel, we do not expect too much in quality and may be willing to give a good rating, even if it is not the same quality as a more expensive hotel.

Chapter One In the first chapter of this dissertation, I attempt to shed light on how ratings are informative of product quality if consumers rate value-for-money.

To study this, we first have to understand what motivates raters to rate based on value-for-money. Value-for-money implies that raters provide better ratings when product quality is high and/or when prices are lower. One explanation for why consumers rate this way is reciprocity.

Reciprocity implies that consumers want to help firms that provide them with a higher value-for-money obtain a higher revenue. We show that a model of reciprocity can explain many observations. For example, why firms may initially set lower prices, why lower-quality products can obtain good ratings, and why the average rating is increasing over time. Since reciprocity can capture many phenomena in online ratings, it is potentially a fitting model to explain why and how people rate.

Using reciprocity, we can understand when and how ratings are informative of product quality. Because firms may choose to set lower prices to allow consumers to obtain a higher value-for-money and receive a good rating, firms can strategically influence the ratings they receive. Indeed, firms choose to do so because better ratings can translate to higher demand or the ability to set higher prices.

Since prices can influence ratings, when consumers cannot observe past product prices, they cannot know how much of a role price plays in the rating they see. Hence, firms with lower-quality products can benefit from setting a lower price, giving up some profit today to obtain a good rating and receive a higher profit in future. As a result, ratings become less informative of product quality when firms can influence the ratings they receive.

Do consumers correctly anticipate how raters rate? The first chapter studies the informativeness of ratings as they relate to product quality. However, when consumers prefer different aspects of a product, it is no longer clear what it means to be a high-quality product.

Returning to the previous illustration of a hotel, some guests may consider the comfort of the room to be of utmost importance, while others may prefer the nuances and touch of the staff. A stark distinction arises when considering business travellers - who may emphasise comfort and rest, and holidaymakers - who may consider proximity to transport and sights more important. Hence, the definition of quality may no longer be straightforward.

The second chapter investigates this.

Chapter Two In this chapter, I explore how raters and future consumers interpret ratings, focusing on the second motivation of why people rate: out of kindness to future consumers. If ratings are provided out of kindness to future consumers, this should minimise the role prices play in the choice of rating. Hence, allowing us to focus on how heterogeneous preferences (difference in preferences) between consumers affect the interpretation of a rating.

We adopt a model of altruism, where a rater concerns himself with the expected utility of future consumers. We find that a rater's decision to rate in this circumstance depends on what they believe a consumer interprets a rating to represent. Likewise, how consumers interpret ratings depends on what the consumer thinks the rater is rating. Hence, higher-order beliefs come into play. In other words, raters need to anticipate what consumers believe the rater is rating. Similarly, consumers anticipate that raters rate based on how raters believe consumers interpret ratings.

Higher-order beliefs mean multiple equilibria exist, each depending on the beliefs of raters and consumers. We set out to investigate if raters and consumers anticipate the beliefs of the other correctly. Using an experiment, we show that raters and consumers agree on an equilibrium in different ratings environments. Hence, this suggests that despite heterogeneous preferences leading to multiple equilibria, raters and consumers interpret ratings similarly.

Open-Source Software Development

In the third chapter, I turn my attention to a different phenomenon, the development of open-source software (OSS), software for which the underlying source code is available ('open') for everyone to see. Specifically, discussing how managers coordinate OSS development, taking into account that OSS developers are essentially volunteers.

The development of open-source software (OSS) is rather interesting: Individuals get together to contribute as much (or as little) as they want to develop software. This means that developers only contribute towards software features that interest them, often because it helps them at work or allows them to develop specific skills. Moreover, because developers work independently and often anonymously, coordinating their efforts is necessary for a coherent and successful outcome.

However, the role of a manager has its challenges. First, managers have to consider their personal objectives. Do they want to develop a software which caters exclusively to their interests? Or are they concerned with how their software can benefit others? Or perhaps they are driven by profit, seeking to monetise the OSS. Second, managers have to attract volunteers to develop the OSS. This means that the design of the software must appeal to developers, who may have different preferences from the manager. Finally, OSS are often seen as an alternative to existing proprietary software, and Managers have to consider if the OSS is still worth developing given the existing software.

In this chapter, I explore how managers balance their private goals and the incentives of individual developers when coordinating the development of OSS. I show that this balance depends on two key factors. The first factor is the competition in the market. When the existing proprietary software has a low quality, developers are motivated to develop the OSS if they find using the proprietary software sufficiently frustrating. This raises the quality of the OSS, which allows the OSS to attract more users. Since developers are already intrinsically motivated to contribute to the OSS, managers can more effectively capture the benefits of this positive network effect by exerting additional effort into coordinating the development of the OSS.

The second factor is the OSS software license. Software licenses grant rights to their holder. If a software license is permissive, it limits the manager's control over the OSS's development and allows others to repurpose the source code as they see fit. This means that other developers may choose to incorporate part, or all, of the source code into their own projects. To prevent this from happening, managers tend to be more accommodating of features that interest other developers. I show that when an OSS has a more permissive license, managers exert less effort in coordinating its development. Managers choose to do so because they want to avoid coordinating the development of features they have little interest in.

I study managers with a variety of incentives: self-interested managers who care exclusively about their personal objectives, altruistic managers who concern themselves with how others benefit from the software, and profit-driven managers looking to monetary rewards from the development of the OSS. While they respond to competition similarly, self-interested managers always exert more effort when coordinating OSS development. As a result, they generate a higher level of surplus than altruistic or profit-driven managers.

In summary, the first chapter asks how far consumers can trust ratings to be informative of product quality. The second chapter discusses whether raters and consumers interpret ratings similarly. The third chapter looks at the motivations and issues surrounding the management of open-source software development. These three chapters study how and why people are motivated to interact in online communities, contributing to various types of public goods. Hence, the title "Online Public Goods".

Chapter 1
Ratings and Reciprocity

This paper was written in collaboration with Johannes Johnen.

Abstract

Evidence suggests online ratings and reviews are motivated by reciprocity. We incorporate a standard model of reciprocity into a model of ratings to capture 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' reciprocal behaviour and obtain a good rating and larger profits in the future. Because also low-quality firms harvest ratings, reciprocity makes ratings less-informative about quality. Based on this mechanism, (i) we argue that reciprocity-based ratings cause 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. We show that even as ratings become less-informative, consumers can benefit from lower prices. 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 thereby undermines the informativeness of ratings and harms consumers.

JEL: D21, D83, D90, L10

Keywords: Reciprocity, Ratings and Reviews, Digital Economy, Reputation

We thank Paul Belleflamme, Maren Hahnen, Xavier Lambin, François Maniquet, 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 concerée grant 19/24-101 'PROSEco'. Robin acknowledges funding from the FRS-FNRS Aspirant grant and the SSRC/SMU graduate research fellowship.

1.1 Introduction

Ratings feature prominently in the decisions we make each day. We often rely on the experience of others, through ratings and reviews, 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). Empirical work highlights that prices influence ratings. For example, Li and Hitt (2010a) 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 have the ability to set higher prices (Luca & Reshef, 2021a). 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. 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 address the question: "are value-for-money based ratings informative of product quality?"

To develop a framework where prices influence ratings, we build on the well-established concept of reciprocity (Bolton & Ockenfels, 2000; Dufwenberg & Kirchsteiger, 2004; Rabin, 1993). Evidence on reciprocity suggests that firms may choose to lower prices or provide rebates to offer a higher value-for-money to consumers. These consumers interpret a sufficiently high value-for-money as a kindness, which they reciprocate with a good rating (Fradkin et al., 2021a; Halliday & Lafky, 2019a; Li & Hitt, 2010a). Reciprocity allows us to capture that consumers rate based on observed quality and price, and reward 'good' deals with a good rating. Given the true quality and price, the consumer is willing to make the effort to rate if a seller made a *sufficiently* kind offer, measured in terms of value-for-money. If the consumer receives a sufficiently high value-for-money, then they are willing to reciprocate the firm's kindness with an act of kindness of their own in the form of a good rating. Thus, 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 in each period—a long lived monopolistic firm and a short lived consumer. The

[1] They 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.

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. 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 is representative of 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.[2]

In line with evidence that high-quality firms are more likely to get good ratings (Ananthakrishnan et al., 2019; Li et al., 2020b), we focus on equilibria where high-quality firms set sufficiently low prices to obtain a good rating.

Our key mechanism is that low-quality firms face the following trade-off. First, in period 1 they can set a price which is sufficiently low to trigger reciprocity. 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 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'*.

This trade-off of low-quality firms influences how well consumers can infer quality from ratings. If low-quality firms harvest ratings, both types of firms get a good rating, and future consumers cannot use ratings to distinguish firms. But if low-quality firms mimic prices, only high-quality firms offer a sufficiently high value-for-money to obtain a good rating, allowing future consumers to distinguish firms.

In equilibrium, if consumers have a weak tendency to reciprocate, only high-quality firms obtain a rating and ratings perfectly signal quality. But if consumers are more inclined to reciprocate value-for-money, low-quality firms play a mixed-strategy, and do both price mimicking and ratings harvesting with positive probability. Intuitively, low-quality firms harvest ratings to get a good rating and free-ride on the reputation of high-quality firms.

[2]While 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.

But this undermines the value of a good rating until, in equilibrium, low-quality firms are indifferent between ratings harvesting and price mimicking. This mixed-strategy equilibrium endogenously determines how well ratings signal 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 well ratings signal quality. We leverage our key trade-off to understand how changes in the ratings environment impact the informativeness of ratings.

The dynamic pricing in our mixed-strategy equilibrium closely resembles evidence on rating harvesting: in online marketplaces, firms build reputation in early periods by obtaining good ratings. Subsequently, they increase prices (Cabral & Hortaçsu, 2010; Cabral & Li, 2015b; Li et al., 2020b).

Our main mechanism implies that reciprocity makes ratings less-informative and causes rating inflation.[3] Intuitively, when consumers have a stronger sense of reciprocity, firms need only leave a smaller surplus to trigger consumer's reciprocal behavior. This means that firms can set a higher price and still receive a good rating. Therefore, low-quality firms harvest ratings more often and receive better ratings. But if low-quality firms get better ratings, ratings become less-informative about quality.

Even though ratings 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 'reciprocity 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 reciprocity 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.

Next, 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 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

[3]Rating inflation here refers to the observation that ratings 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.

accurately [...] reflect the experience of all purchasers".[4,5] We show that this logic ignores how a lower effort to rate impact how sellers harvest ratings. Making it easier to rate means that firms need only transfer a smaller reciprocity 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 (2015b) 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. (2021) propose that sellers should not be able to pay raters conditional on the content of their ratings or reviews. We show that even unconditional payments for ratings—even when not discriminating between worse or better ratings—leads to less-informative 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 reciprocity 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 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 ratings, but less-informative ratings.

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.

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

[5]Prior 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).

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

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

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 reciprocity can explain why we observe extreme (very positive or negative) ratings in practice. 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, 2006a). We also show that our results persist when we introduce a continuum of firms.

We introduce the basic model in Section 1.2, and discuss the equilibrium in Section 1.3. Section 1.4 shows how various features in the rating system influence how well ratings reflect quality. We then discuss implications on surplus in Section 1.5. We present extension and robustness checks in Section 1.6. Section 1.7 connects our results to the literature, and Section 1.8 concludes.

1.2 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 $j \in \{L, H\}$ has quality q^j, where $q^L < q^H$. The probability that the firm is of type q^H or q^L is common knowledge 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, participate in only one of the two periods, with a new unit mass of consumers arriving in each period. We normalize the value of their outside option to zero. 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 past prices. Consumers may choose to buy or not to buy. If they choose not to buy, they exit the market. If they consume, they observe the firm's true

[8]We only require that low-quality firms face a smaller marginal cost than high-quality firms. Without loss of substance, it simplifies the analysis to assume that both firms face zero cost of production. This is similar to allowing for costless signals in a cheaptalk game (see Kreps and Sobel (1994) and Crawford and Sobel (1982)). Additionally, our model features a concept of endogenous cost such as that in Martin and Shelegia (2021). In our setting, by harvesting ratings, low-quality firms reduce the benefit of ratings harvesting.

quality and decide on leaving a rating. For simplicity, 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 $R_t = 0$, consumers choose not to provide any rating. 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.[9] 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.

The rating utility captures intrinsic reciprocity models as in Dufwenberg and Kirchsteiger (2004) and Rabin (1993).[12] For consumers in period t it is given by $v_t = [\kappa q^j - p_t^j]\Delta - e$ if $R_t = 1$ and $v_t = 0$ if $R_t = 0$. $\kappa \in [0, 1]$ represents the proportion of surplus which consumers think is equitable for firms to receive; $\Delta > 0$ represents the warm glow a consumer enjoys from the kindness of leaving a rating to the firm; and $e \geq 0$ reflects the opportunity cost of providing a rating.

The first term $[\kappa q^j - p_t^j]$ captures the consumer's perception of the firm's kindness. Consumers perceive a price equal κq^j as fair. Thus, they perceive any price below κq^j as a kindness, and $[\kappa q^j - p_t^j]$ is positive. Otherwise, if $[\kappa q^j - p_t^j]$ is weakly negative, firms keep more surplus than what consumers deem as equitable, and consumers perceive firms as unkind.

The second term Δ describes the consumers' sense of reciprocity and captures the warm glow that consumers receive from being kind to a firm by leaving a good rating. Finally, the consumer faces some costs when leaving a rating e, such as time, effort, attention etc. Below, we consider that e may depend on the design of a ratings systems, for example, one-click ratings, constant reminders, and purchase verification.

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

[9] We start with an environment where negative ratings do not exist to simplify exposition of results. We discuss how our results extend to a setting with negative ratings that punish bad deals in Section 1.6.

[10] This captures cases where the firm is new in the market. More generally, our model captures cases where two firms of different quality have the same rating at the beginning of period 1.

[11] This is highlighted by Cabral and Li (2015b), who show that incentivizing consumers to rate does not change the number of bids or bid levels on eBay, indicating that consumers do not consider the rating incentives when making a purchase.

[12] Derivation can be found in the Web Appendix A.2.3.

$t = 1$

Firm draws quality, q^j.

Firm sets price, p_1.

Consumers observe price and form expectations over consumption utility, u_1.

Consumers make purchase decision, if $u_1 \geq 0$.

Consumers that purchase observe true quality.

Consumers evaluate rating utility, and rate if $v_1 \geq 0$.

Period 1 consumers exit the market.

$t = 2$

Period 2 consumers arrive.

Firm observes historical rating, R_1, and sets new price, p_2.

Consumers observe past ratings and current price, forming expected quality.

Consumers make purchase decision.

In the main model, we make some simplifying assumptions. We discuss in Section 1.6 how results are robust to various extensions, namely a game with more than 2 periods, a continuum of quality types of the firm, and a rating system with positive, and negative ratings.[13]

This model is representative of markets such as Amazon, Taobao, eBay, AirBnB[14] and Google reviews. These platforms help to facilitate matches between consumers and firms. Consumers rely on ratings to form or update their expectations of product quality. Firms use ratings to differentiate themselves from lower quality firms, allowing them to build trust and gain patronage.

[13]Lewis 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 & Hortaçsu, 2003; Cabral & Hortaçsu, 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.

[14]Although we do not explicitly capture the mutual nature of AirBnB's rating system, ratings on AirBnB are only revealed after both host and guest have provided a rating. This removes the threat of retaliation rating-for-rating in response to a negative rating, which is why AirBnB and other similar mutual rating system resemble our one-directional rating system.

1.3 Equilibrium

We look for a perfect Bayesian equilibrium. We apply two additional restrictions as equilibrium selection assumptions.[15]

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

The restriction implies that consumers who observe a rating expect weakly higher quality of the product. We focus on such equilibria in line with evidence that high-quality firms are more likely to receive good ratings (Ananthakrishnan et al., 2019; Li et al., 2020b).[16] Importantly, Restriction 1 allows us to focus on the strategic decisions of low-quality firms.

We make a second selection assumption to limit off-the-path beliefs.

Restriction 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.*[17]

Restriction 2 implies that in a given period t, consumers have the same beliefs about quality for any price that induces the same average consumer rating. The restriction captures that firms can mimic each others' prices, making prices a poor signal of quality. More precisely, because all firms have zero marginal cost, low-quality firms can always deviate to any price set by a high-quality firm. This makes it difficult for high-quality firms to use price signals to differentiate themselves from low-quality firms.

Together, these restrictions give ratings the best shot at being an informative signal for product quality. Restriction 1 focuses on equilibria where high-type firms always receive a good rating, and Restriction 2 ensures that the information signal that future consumers receive comes from ratings alone.

The following proposition characterizes equilibria in this game.

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

1. *In period 1, high-quality firms charge* $\overline{p} \equiv \frac{\gamma q^H + \delta^*(1-\gamma) q^L}{\gamma + \delta^*(1-\gamma)}$ *with probability 1 and receive*

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

[16] Formally, equilibria where high-quality firms prefer to get a rating with probability 1 exist if κ is sufficiently large. Intuitively, when κ is large, the high-quality firm can charge a high price and still receive a good rating. This reduces their incentive to deviate to prices even closer to or at q^H, at which they receive no rating.

[17] Alternatively, we could use the D1 refinement, but we choose to use Restriction 2 instead, as it is weaker restriction and more intuitive (Cho & Sobel, 1990).

a good rating.

2. *In period 1, low-quality firms randomize their price.*

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

    b. *They charge* $\underline{p} \equiv \kappa q^L - \frac{e}{\Delta}$ $(< \overline{p})$ *with probability* $1-\delta^*$ *and obtain a good rating.*

    c. $\delta^* \in (\frac{1}{2}, 1)$ *if and only if*

$$(1-\gamma)(q^H - q^L) - (1-\kappa)q^L > \frac{e}{\Delta}, \tag{1.1}$$

    and $\delta^* = 1$ *otherwise.*

3. *In period 2, prices equal expected quality conditional on ratings.*

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

The equilibrium is unique up to off-equilibrium-path beliefs and exists if $\kappa q^H - \frac{e}{\Delta} \geq \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$.[18]

Firms charge one of two prices in equilibrium. High-quality firms always set $\overline{p}$ to extract all expected surplus (conditional on observing $\overline{p}$).[19] 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 firm and charge $\overline{p}$. This, in equilibrium, can lower the conditionally expected surplus below q^H. The low price $\underline{p}$ is just low enough so that consumers perceive low-quality firms as sufficiently kind to give them a good rating, i.e. $\underline{p}$ is such that $[\kappa q^L - \underline{p}]\Delta - e = 0$.

How do low-quality firms choose between these two prices? Low-quality firms benefit from ratings by earning higher prices in period 2 (point 3). But to get a rating, they have to charge a low price $\underline{p}$ $(< \overline{p})$ in period 1 so that consumers reciprocate with a good rating, which is why we call this strategy 'ratings harvesting'. Alternatively, in period 1 low-quality firms can mimic the high price of high-quality firms $\overline{p}$, and get no ratings and lower profits in period 2. 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. This δ^* also captures the informativeness of ratings: as δ^* increases, low-quality firms obtain

[18] The existence condition captures that κ should be sufficiently large. Intuitively, if κ is too small, consumers believe that firms should get only a very small share of the surplus and never leave a rating. Formal proofs can be found in Appendix A.1.

[19] Formally, Restriction 2 implies that $\overline{p}$ extracts all conditional expected surplus. Without this restriction, there can be equilibria where the high price is lower than the conditional expected surplus.

a rating less often, and the ratings better help consumers distinguish between high-and low-quality firms.

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.[20]

We now explain the impact of reciprocity on low-quality firms' pricing and ratings. To begin, suppose the kindness Δ is sufficiently small such that $\delta^* = 1$ in equilibrium. For a given quality of ratings δ^*, low-quality firms need to offer a discount $\overline{p} - \underline{p}$ to get a good rating; this discount $\overline{p}-\underline{p}$ decreases as Δ increases. Intuitively, when Δ is small, the rating utility $v_1 = [\kappa q^L - p_1^L]\Delta - e$ is quite 'small' and consumers have a low level of intrinsic motivation to leave a good rating. To encourage consumers to rate and compensate them for the cost of ratings e, low-quality firms need to offer a lower price to provide a higher value-for-money for consumers. When Δ is sufficiently small, low-quality firms find it too costly to obtain a rating, so they charge the higher price of $\overline{p}$ with probability $\delta^* = 1$. Because only high-quality firms obtain a rating, ratings perfectly signal quality.

Now suppose Δ increases. As Δ increases, consumers are more inclined to reciprocate a given value-for-money with a rating, which is why the price for which low-quality firms can obtain a good rating, $\underline{p}$, increases. In other words, more kindness (Δ) lowers the opportunity cost of obtaining a good rating. Low-quality firms set the lower price of $\underline{p}$ more often, and obtain a good rating in equilibrium with positive probability.

These arguments explain why reciprocity reduces the quality of ratings: low-quality firms can offer a higher value-for-money to trigger consumers' reciprocity and obtain good ratings. But when low-quality firms start to harvest ratings, both high- and low-quality firms receive good ratings with strictly positive probability, which makes ratings less-informative about product quality. Thus, low-quality firms harvest ratings to trigger the consumers' kindness, but thereby undermine the quality of ratings. The following proposition summarizes this result.

Proposition 2. $\frac{\partial \delta^*}{\partial \Delta} < 0$ *when* (1.1) *holds, and* $\frac{\partial \delta^*}{\partial \Delta} = 0$ *otherwise.*

When consumers exhibit a larger sense of reciprocity, low-quality firms harvest ratings more frequently. Two clear implications arise form this proposition. First, firms use prices to induce consumers to be kind. Second, when consumers are more reciprocal, ratings become less-informative.

[20] In equilibrium, $\delta^* > \frac{1}{2}$, because low-quality firms will only find it profitable to free ride on the reputation of high-quality firms, if their reputation is sufficiently good, i.e. if δ^* is sufficiently large.

Growing evidence suggests that ratings are influenced not just by quality, but also by prices. 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 Hitt (2010a) highlight that a 1% increase in price 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 (2021a) provides evidence that a 1% increase in prices leads to a 3-5% decrease in average rating. Abrate et al. (2021) suggests that a 1% increase in hotel prices leads to a decrease of 1 star (out of 10) in overall ratings. Because these articles control for product characteristics, they suggest that it is value-for-money that influences ratings, and not quality alone. Our mechanism explains how both quality and prices affect ratings through reciprocity: consumers perceive a higher value-for-money as a kindness, which they reciprocate with positive ratings.

Our results connect to evidence on rating inflation. Rating inflation is well documented in the literature (Filippas et al., 2022), but not well understood. We offer novel explanations for this phenomenon: 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.

We have shown how reciprocity can impact the informativeness of ratings. We now explore how the design of ratings systems can impact the informativeness of ratings, and ask how beneficial ratings are for consumers.

We restrict the remainder of our analysis to situations where the low-quality firm plays a mixed-strategy, i.e. where (1.1) holds. First, the condition is satisfied when the difference in quality is sufficiently large, i.e. when ratings are more relevant in the first place. Second, doing so allows us to directly study how changes to the ratings environment impacts the informativeness of ratings, captured by δ^*.

1.4 Designing ratings environments

In this section, we discuss how common features of a platform, such as the effort needed to leave a rating and the aggregate product quality on a marketplace, influence the informativeness of ratings.

Facilitating ratings

The design of ratings systems can make it easier, or more difficult, 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 and consumers respond when ratings become easier.

In our setting, less-costly ratings lead to less-informative ratings. The reason is closely related to reciprocity: if a kind action is less effort, consumers are more inclined to be kind. Thus, as becomes easier for consumers to rate (i.e. e decreases), firms need to leave less surplus to induce consumers to reciprocate with a good rating.[21] In equilibrium, this encourages especially low-quality firms to harvest ratings: they set $\underline{p}$ with a higher probability $(1 - \delta^*)$, which makes ratings less-informative about quality.

Corollary 1. *If* (1.1) *holds, then* $\frac{\partial \delta^*}{\partial e} > 0$*.*

The result suggests that making ratings less costly for consumers, even though it induces more ratings in equilibrium, encourages low-quality firms to harvest ratings and makes ratings less-informative. This provides another channel through which reciprocity-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 (2015b) measure quality using shipping speed. They find that for low-quality products, higher rebates for ratings decrease the proportion of negative ratings. Since rebates offset the cost of rating, their result highlights—as in our model—that lower cost of rating makes ratings less-informative about quality. Lafky (2014a) shows that when it becomes more costly to rate, ratings become more extreme. This follows our prediction that it is easier to differentiate between firms when ratings are costly.

Anecdotal evidence suggests that the cost to leave a rating has decreased over time on many ratings platforms. Yelp and Google encourage ratings with various perks, such as invitations to exclusive events and discount codes.[22] On google reviews, users receive constant reminders to leave ratings and can leave one-click ratings on their smartphone. This trend of fuss-free ratings is also gaining traction on Amazon. 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.[23] Publicly, they stated that more

[21] Formally, $\underline{p}$ increases and the difference between $\overline{p}$ and $\underline{p}$ shrinks.
[22] Yelp Elite Squad and Google Local Guides perks as described by Yelp (Yelp, 2022) and the Harvard Business Review article Donaker et al. (2019).
[23] A timeline of Amazon's rating system as documented by Forbes (Masters, 2021).

feedback would lead to more accurate ratings, relying on the law of large numbers to drown out fake reviews.[24] Our results suggest that Amazon's effort to encourage ratings to make them more-informative may backfire. Although more 'real' customers may be rating, which could solve the problem of fake ratings on Amazon, easier ratings also affect pricing incentives of firms. By reducing the cost to leave a rating, low-quality firms firms find it easier to trigger reciprocity and obtain a good rating, which can ultimately lead to less-informative ratings.

Quality control

We have shown that (i) a stronger sense of reciprocity (Δ), and (ii) lower cost to leave a rating (e) encourage low-quality firms to harvest ratings, triggering rating inflation. In both cases, rating inflation is the result of less-informative ratings. In practice, however, an increase in the average product quality also improves ratings. We now shed light on this scenario and explore how changes in the aggregate quality of sellers in the market affect the informativeness of ratings.

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 pre-screen and weed out low-quality firms to control the quality of 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. This way, better average seller quality may render ratings more- or less-informative.

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.

The starting point to understand this result is that low-quality firms only harvest ratings, when ratings are useful to distinguish seller quality. Intuitively, low-quality firms can only free-ride on the reputation of high-quality firms, when ratings are somewhat useful to distinguish sellers so that high-quality firms indeed have a reputation.

Let us first consider the case where the aggregate quality in the market is low. For simplicity, suppose γ is close to zero. Almost all firms are of low quality, so ratings are somewhat useless to help consumers distinguish sellers. But as γ increases, ratings become more useful, which encourages low-quality firms to harvest more ratings. Thus, for 'small' γ, an increase in the aggregate quality leads to less-informative ratings.

Now consider the opposite scenario where the aggregate quality is high. For simplicity,

[24]As reported by TechCrunch (Perez, 2019).

suppose γ is close to 1. Again, ratings are not useful to distinguish sellers. But as γ increases, they become even less useful to distinguish sellers so that low-quality firms harvest ratings less. This is why for 'large' γ, an increase in the aggregate quality leads to more-informative ratings.

The following proposition summarizes this result.

Proposition 3. *There exists a unique* $\overline{\gamma} \in (0,1)$ *such that* $\frac{\partial \delta^*}{\partial \gamma} < 0 \iff \gamma < \overline{\gamma}$, *and* $\frac{\partial \delta^*}{\partial \gamma} > 0 \iff \gamma > \overline{\gamma}$.[25]

The proposition works out how changes in aggregate quality interacts with the informativeness of ratings. When aggregate quality is low, quality increases make ratings less-informative. But when aggregate quality is high, improvements in quality have a double dividend: aggregate quality increases and ratings become more-informative.

Our 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.[26] Also Uber has announced that it will remove both riders and drivers with consistently poor ratings,[27] and both Uber and their subsidiary Uber Eats suspend drivers who fall below a minimum rating.[28] Booking.com suspends properties for quality control purposes;[29] Airbnb bans hosts based on a combination of factors, including being in the bottom 1% of overall ratings and guest feedback.[30]

We explain how such measures may affect the informativeness of ratings. First, 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.[31] Thus, even in scenarios when good ratings are uninformative, platforms can indeed use bad ratings to weed out low-quality firms. Second, when quality controls make ratings less-informative, that platforms can counter such adverse affects by adjusting the effort to rate accordingly.

[25] $\overline{\gamma}$ is given by (A.8).

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

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

[28] Leaked 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).

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

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

[31] In 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.

1.5 Surplus Analysis

So far, we studied how reciprocity affects the informativeness of ratings. Conventional wisdom suggests that more-informative ratings help consumers make more-informed purchase decisions and ultimately benefit consumers. We now explore the link between the informativeness of ratings and consumer surplus more carefully. The key insight is that consumer-optimal rating systems are often somewhat, but never fully, informative.

To start, we investigate how ratings harvesting affects consumer surplus. In order to induce consumers to reciprocate, 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. To see this rent, we write down the expected consumer surplus

$$CS = (1-\gamma)(1-\delta^*)(q^L(1-\kappa) + \frac{e}{\Delta}). \tag{1.2}$$

In period 2, and in period 1 when consumers face a high price of $\overline{p}$, 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 $q^L - \underline{p} = (q^L(1-\kappa) + \frac{e}{\Delta})$. Low-quality firms leave this surplus so that consumers reciprocate with a good rating, allowing low-quality firms to free-ride on the reputation of high-quality firms and charge a higher price in period 2.

We have shown above that low-quality firms harvest ratings and make ratings less-informative. But in order to harvest ratings, low-quality firms need to offer low prices. This is why rating harvesting benefits consumers. This, however, does not imply that consumers prefer uninformative ratings. 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 lower cost of ratings (e) has two opposing effects on consumer surplus. First, and following directly from Corollary 1, low-quality firms harvest ratings and charge the low price $\underline{p}$ more often; this tends to increase consumer surplus. Second, however, the level of the low price, $\underline{p}$, increases: when consumers can rate more easily, they are more inclined to reciprocate kindness. This is why low-quality firms who harvest ratings can now charge a higher $\underline{p}$ and still receive a good rating, reducing consumer surplus. Overall, because easier ratings encourage low-quality firms to harvest ratings, they harvest ratings more often, which tends to benefit consumers. But low-quality firms do so with a higher price, which tends to harm consumers.

We show that these two opposing effects pin down a positive level of ratings effort that maximizes consumer surplus.[32]

Proposition 4. *There exists a level of effort $e^{cs} > 0$ that maximizes consumer surplus. At this level of effort, $\delta^* \in (\frac{1}{2}, 1)$. This is true if*

$$(1-\gamma)^2\gamma(q^H - q^L)^2 \geq (1+\gamma)^2(q^L(1-\kappa) + \frac{e}{\Delta})^2. \tag{1.3}$$

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

The proposition characterizes when the conventional wisdom that more-informative ratings benefit consumers is true. 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.

A key implication of the Proposition is that even though consumers benefit when low-quality firms harvest ratings, 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 (1.3) consumers prefer somewhat-informative ratings if the difference between high- and low-quality firms is sufficiently large. This is similar in spirit to (1.1) 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. While buyers prefer somewhat-informative ratings, sellers on average prefer uninformative ratings. Intuitively, seller surplus is largest when sellers leave the smallest reciprocity rent to buyers. This is

[32] Note that if we would incorporate rating utility into consumer surplus, consumers would get an additional benefit and cost from rating, so the consumer optimal level of effort could be higher or lower. But crucially, it can still be positive.

true when $e = 0$ and ratings are least-informative. In this case. low-quality firms have the highest incentive to harvest ratings, but face the smallest opportunity cost for doing so.

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, this 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* (1.3) *holds, average seller surplus is maximal at* $e^s = 0$ ($< e^{cs}$). *Moreover, profits of high-quality firms increases in* e*, and profits of low-quality firms decreases 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. This suggests that the aforementioned efforts of Google reviews and Amazon to facilitate and encourage ratings might not just lead to less-informative ratings, but also harm consumers through higher prices. They do, however, benefit sellers.

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 3 that low-quality firms harvest ratings when they can free-ride on the reputation of high-quality firms. Thus, when γ is large and there are more high-quality firms to free-ride on, an individual low-quality firm harvests ratings more and benefits consumers. On the other hand, only low-quality firms harvest ratings and leave reciprocity rent to consumers, suggesting consumers benefit when γ is sufficiently small. As a result of these opposing forces, consumer surplus is concave in γ and an intermediate $\gamma^{cs} \in (0, 1)$ maximizes consumer surplus.

The second key insight resembles a previous one on the cost of ratings: consumers prefer quality levels that lead to somewhat-informative ratings ($\gamma^{cs} < \overline{\gamma}$, where $\overline{\gamma}$ is the same defined in Proposition 3), but not fully-informative ratings ($\gamma^{cs} \to 1$). The intuition is familiar from above: with somewhat-informative ratings, low-quality firms harvest ratings, which puts pressure on prices.

The next proposition summarizes these results.

Proposition 5. *Equilibrium consumer surplus is strictly concave in γ. There exists an aggregate quality level, denoted by γ^{cs}, that maximises consumer surplus, where $\overline{\gamma} > \gamma^{cs}$ and $\gamma^{cs} > 0$.*

Proposition 5 implies that, even when larger aggregate quality leads to more-informative ratings, consumers can be worse off. This is the case when $\gamma \notin [\gamma^{cs}, \overline{\gamma}]$. For $\gamma < \gamma^{cs}$, better quality encourages low-quality firms to harvest ratings, which benefits consumers but undermines the informativeness of ratings. For $\gamma > \overline{\gamma}$, 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 $\gamma^s \to 1$*

Because all firms benefit from a higher average quality, this also implies that (remaining) firms prefer a higher level of quality controls than consumers. Together with Proposition 3, Lemma 1 implies that sellers prefer ratings that are uninformative. Additionally, when considering Proposition 5, Lemma 1 suggests that firms prefer ratings that are less-informative than what consumers prefer. The intuition is familiar from above: firms on average prefer uninformative ratings, and consumers prefer somewhat-informative ratings. Hence, firms prefer ratings that are less-informative than consumers.

1.6 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. To do so, the platform sets a royalty to sellers, and chooses how easy consumers can rate (through e, e.g. by allowing one-click ratings or by sending reminders). 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. 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 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 the ever changing rating system on Amazon, from introducing new verification methods to allowing ratings to be provided without reviews Amazon, 2021. For details, see Web Appendix A.2.1.

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. For details, see Web Appendix A.2.1.

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 (negative reciprocity). When firms leave a sufficiently small surplus to consumers, consumers perceive this as unkind and respond with an unkindness, retaliating 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.[33] But—different from the main model—ratings are more extreme.[34] For details, see Web Appendix A.2.1.

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

[33] Fradkin 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.

[34] When consumers exhibit biased reciprocity, that is the warm glow they receive from leaving a good rating exceeds the warm glow they receive from leaving a bad rating, our result provides theoretical support for suggestions by Dellarocas and Wood (2008) and Filippas and Horton (2022) that reciprocity bias leads to a J-shaped ratings distribution, i.e. extreme ratings where negative ratings are less common than good ratings.

less-informative ratings as in the main model. For details, see Appendix Web Appendix A.2.1.

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. Interestingly, and in line with evidence by Cabral and Hortaçsu (2010) and Jin and Kato (2006a), low-quality firms are less likely to sustain a good rating. For details, see Web Appendix A.2.1.

1.7 Related Literature

Our key results connect **evidence on ratings.** Ratings aid consumers by reducing uncertainty about the quality of the product. Much of the literature suggests that ratings do indeed signal quality on major platforms like eBay (Hui et al., 2021), Taobao (Zhang et al., 2012a), and Airbnb (Proserpio et al., 2018a). The broader empirical literature, however, suggests that other factors also influence ratings (Gao et al., 2018a; Masterov et al., 2015a; Nosko & Tadelis, 2015; Zervas et al., 2021a).

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., 2014a; Carnehl et al., 2021; Li & Hitt, 2010a; Luca & Reshef, 2021a; Neumann et al., 2018). Second, firms with better ratings charge higher prices in the future, a pattern frequently referred to as 'rating harvesting' (Cabral & Hortaçsu, 2010; Cabral & Li, 2015b; Cai et al., 2014a; Carnehl et al., 2021; Ert & Fleischer, 2019; Gutt & Herrmann, 2015; Jin & Kato, 2006a; Jolivet et al., 2016; Lewis & Zervas, 2019; Li et al., 2020b; Livingston, 2005; Luca & Reshef, 2021a; McDonald & Slawson, 2002; Neumann et al., 2018; Proserpio et al., 2018a). We can explain both of these patterns in a single framework, i.e. low-quality firms charge lower prices to induce consumers to reciprocate with a good rating. Firms harvest these good ratings by charging higher prices in the future.

We also contribute to the **literature on reciprocity in ratings**. A series of empirical articles argue that reciprocity is a key driver of rating behavior, (Cabral & Li, 2015b; Diekmann et al., 2014a; Fradkin et al., 2021a; Li & Xiao, 2014b; Zervas et al., 2021a). Some experimental work directly identifies that reciprocity drives rating behavior (Bolton et al., 2013a; Halliday & Lafky, 2019a; Lafky, 2014a). Taken together, these experiments suggest: (i) reciprocity biases ratings upward in mutual-rating systems where buyers and sellers rate each other. But double-blind feedback strongly reduces this bias. Also following the work of Dellarocas and Wood (2008) and others on eBay, most online marketplaces adopted a double-blind approach to feedback, which is why we do not look at mutual-rating systems. (ii) Also with one-sided or double-blind rating systems, sellers take advantage of their ability to influence ratings, and they use prices to do so. We contribute to this literature by modelling how sellers use prices to trigger reciprocity and thereby influence ratings. By doing so, we derive novel predictions on how reciprocity causes rating inflation, namely that reciprocity induces ratings harvesting, reciprocity makes ratings less-informative, and that encouraging consumers to rate (e.g. through one-click ratings or reminders to rate) can backfire by making ratings less-informative.

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., 2021a).[35] We propose multiple channels through which reciprocity 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.[36]

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, Weinmann, Mohr, & vom Brocke, 2021a).

We connect to the wider theoretical literature on **trust and information transmission in the digital economy**. Platforms may recommend products (Hagiu & Jullien, 2011; Peitz & Sobolev, 2022), shroud additional fees and features of third-party sellers (Johnen & Somogyi, 2021), 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.

We closely connect to the **theoretical literature on reputation** (e.g. Cabral (2000a),

[35] In particular, Nosko and Tadelis (2015) study the binary rating environment on eBay and find that more than 99% of ratings are positive. Filippas and Horton (2022) study a marketplace that kept ratings unpublished at first and then started to publish ratings. Once ratings were published, they got better. This is in line with our mechanism: consumers can only reciprocate a good value-for-money once ratings are published, so our mechanism predicts more rating harvesting and therefore rating inflation once ratings are published.

[36] Research on information systems propose that social norm of good ratings increases over time (Qiu et al., 2012). Our mechanism can illustrate how social norms changed the way they did: platforms encourage consumers to rate (lowering e and/or raising Δ), which leads to a new equilibrium where reciprocity plays a more important role and ratings increase.

Jullien and Park (2014a), Kovbasyuk and Spagnolo (2021), Martin and Shelegia (2021), and Tadelis (1999a); see also Bar-Isaac and Tadelis (2008) for a survey), and **word-of-mouth** (Chakraborty, Deb, 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 are independent of price. While some papers relax some of these assumptions (e.g. Chakraborty, Deb, et al. (2022) relax (i), Sobolev et al. (2021) and Stenzel et al. (2020) relax (ii)), no article seems to feature that buyers choose if and how to rate strategically, and prices influence ratings. We capture both of these features, which allows us to capture empirical patterns of prices and ratings we discuss earlier in this section.

Martin and Shelegia (2021) study a signalling model where consumers leave a good rating if the product was better than expected. Initially, high-quality sellers offer a lower price to overdeliver to get a good rating. In contrast, we show how low-quality firms use prices to improve their ratings.

Some recent papers study the impact of **value-for-money ratings**. Stenzel et al. (2020) focus on prices in the long-run equilibria. Instead, we focus on short-term intertemporal trade-offs and derive novel dynamic effects like rating harvesting in our framework.Sobolev et al. (2021) start with the premise that more sales can lead to more or less informative ratings. The key novelty in our setting is that we endogenize consumers choices of if and whom to rate. To do so, we use a well-established theory of reciprocity. In our setting this endogenous choice implies that sellers trade-off ratings harvesting and price mimicking, and leads to different comparative statics for how the cost of rating or aggregate quality affect ratings.

Some researchers argue that consumers should be **paid to rate**. One argument is that sellers should be allowed to pay for feedback, because only high-quality firms are willing to pay for feedback, making this a credible signal (Halliday & Lafky, 2019a; 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., 2004a; Chen et al., 2010a). In contrast, we show that encouraging ratings from all consumers (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 is in line with evidence by Cabral and Li (2015b) and Lafky (2014a) that we discussed above.

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.

We also connect to recent work by Rhodes and Wilson (2018) on **false advertising**. Low-quality firms falsely advertise high quality to free-ride on the reputation of high-quality firms. Also in our setting, low-quality firms free-ride on the reputation of others and undermine information transmission. But our mechanism is very different. In contrast to ads, firms in our setting do not choose their own ratings, but need to charge a low price to trigger reciprocity and get a good rating. This leads to novel and inherently dynamic effects like rating harvesting.

1.8 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 use a model of consumer reciprocity to 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.

In practice, consumers may also read product reviews of previous customers. In principle, these reviews could help consumers disentangle how quality and price affect ratings. But we argue that reviews do not solve the problem. First, even if consumers take the time it takes to read reviews, they will only read a small and selected sample of user experience. 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. In these motivations, however, ratings are not affected by price. 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.

Many platforms give consumers easy access to past ratings, but do not connect them to the prices that the raters paid.[37] 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. We do, however, leave this and other questions for future research.

[37]For 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.

Chapter 2
Ratings with Heterogeneous Preferences

This paper was written in collaboration with Jonathan Lafky.

Abstract

We ask how ratings are interpreted in the presence of heterogeneous preferences among both raters and consumers. Altruistic raters rate for the benefit of future consumers. However, an ambiguity arises in the presence of heterogeneous preferences: on what basis are products evaluated? Theory suggests multiple equilibria exist, and ratings may reflect either the preferences of rater or the expected utility of future consumers. In an online experiment, we examine how ratings are selected by raters and interpreted by consumers, studying how two types of information helps to guide equilibrium selection. In each environment, we show raters and consumers update their evaluation of what a rating represents in similar ways.

JEL: C91, D64, D83, L86
Keywords: Ratings and Reviews, Altruism

We thank Nicolas Jacquemet, Aljoscha Janssen, Johannes Johnen, Amma Panin, Heiner Schumacher, and many seminar participants for their useful comments. Robin acknowledges funding from the Action de recherche concerée grant 19/24-101 'PROSEco', the FRS-FNRS Aspirant grant and the SSRC/SMU graduate research fellowship.

2.1 Introduction

Ratings feature heavily in everyday decisions, from where we eat, what shows we watch, to which toys we buy our children. When placed in situations we seldom face, such as replacing a phone or planning a vacation, ratings become even more valuable. But what can consumers actually learn from ratings? We might reasonably expect better rated products to be of higher quality, but it is not obvious how we interpret ratings when others' tastes may differ from our own. When consumer tastes are heterogeneous, whose preferences do ratings represent?

One illustration of this ambiguity is seen in consumer ratings for hotels, a market where customers can have starkly different preferences. Some consumers prioritise location - proximity to conference venues, theme parks, or entertainment districts. Others are less concerned with location and instead focus on decor, cleanliness, or even the firmness of their mattress. Recognizing these diverse preferences, how does a consumer rate the hotel at the end of their stay? Does a business traveler rate primarily based on proximity to conference centers, or also account for the preferences of those traveling with small children? In general, do raters rate based on their own preferences, or based on the preferences of the average consumer?

A common solution to resolving the uncertainty of a rating is to provide consumers with information about the rater's preferences. Suppose that, in addition to rating the hotel, the rater also provides a written review. This allows other consumers to learn of their preferences. Does revealing the preferences of a rater change the way they rate, or the way consumers interpret the rating? We address both sources of ambiguity in this paper, theoretically and experimentally examining how ratings are created and interpreted in the presence of heterogeneous preferences, as well as how changing consumer information about that heterogeneity influences behavior.

In many settings, it is also possible for consumers to observe the quality of some attributes of a product prior to purchase. In the case of hotels, its location and proximity to other attractions are easily verifiable, while the upkeep of the hotel and its services are less easily observed. Since consumers can already assess a hotel's accessibility, raters incorporating the hotel's location into their rating do not provide additional information. Instead, maximally informative ratings should focus on the less easily observable attributes of the hotel. We ask if raters provide maximally informative ratings, and if consumers anticipate that such ratings only reflect the quality of amenities?

Our question involves two closely related behaviors: how raters choose ratings and how consumers interpret those ratings. We first study how raters choose to leave ratings. A growing literature demonstrates that raters are motivated by social preferences (Bolton et al., 2013b; Cabral & Li, 2015a; Fradkin et al., 2021b; Halliday & Lafky, 2019b; Li & Xiao, 2014a; Zervas et al., 2021b). In particular, Chen et al. (2010b), Diekmann et al. (2014b), and Lafky (2014b) show that raters are motivated by altruism toward future consumers and seek to provide informative ratings for those consumers. However, in the presence of heterogeneous preferences, it is unclear when ratings are most informative. For example, in the context of hotel ratings described above, how should a rater who solely cares about a room's physical amenities choose to rate their room? Does their rating describe their own satisfaction with the room and reflect only the quality of those amenities, or do they consider the distribution of other consumer preferences accounting for both amenities and location?

On the receiving end of ratings, if consumers correctly anticipate altruistic raters and rely on ratings for information, they are likely to expect that ratings reflect a generic view of the world, representing the utility of an average individual. However, just as raters, consumers may project their own preferences upon ratings. As a result, misinterpreting the rating. Hence, for ratings to be informative, raters and consumers must form mutually consistent beliefs. Even when they do, can other information, such as revealing the rater's preference, or verifying some aspect of product quality, guide belief formation?

Even if raters are perfectly altruistic, when raters and consumers have heterogeneous preferences, theory predicts that ratings are only useful if there is an agreement as to what they represent. This leads to multiple equilibria. Hence, despite altruistic intentions, ratings may be ineffective if raters and consumers disagree on what is being rated. Although multiple equilibria exists, if additional information helps raters and consumers to form the same beliefs, this can help them to converge upon an equilibrium. We design an experiment to see if there is a convergence upon any particular equilibrium.

First, we construct a baseline setting where consumers only see ratings, showing that raters provide better ratings for products with a higher overall quality, but favour their own preferences when leaving a rating. Consumers anticipate this, and report a higher willingness to pay when they see better ratings.

Second, we inform consumers of the rater's preferences associated with each rating. Despite the prediction of multiple equilibria, we find that raters and consumers agree upon an equilibrium. We show that raters choose to rate based only on their own preferences and ignore other aspects of a product. Consumers anticipate this and report a higher willingness to pay only if they have the same preferences as the rater. When consumers' preferences are different from the rater, consumers choose to ignore ratings completely.

Third, we allow consumers to directly observe the quality of some attributes of a product. We find that this does not change how raters choose to rate. This contradicts our prediction that raters should only rate unobservable attributes of the product. This also suggests that ratings are an inefficient signal when the quality of some attribute of a product is directly observable by consumers. However, we also show that rates and consumers are able to come to a mutual understanding of what a rating represents.

Finally, having understood how raters and consumers interpret ratings, we ask if information about rater's preferences, or the quality of some attribute of a product, can lead to more efficient consumption decisions. Although consumers are able to better predict the value of the product when they have more information, they tend to overestimate the value when they learn the rater's preference. This indicates that consumers do not necessarily make better consumption decisions when they have more information.

We conclude that information does not always allow consumers to make better decisions. Consumers only make better decisions if the information coincides with what they explicitly need to know. Otherwise, relying on word of mouth (ratings) does not benefit their decision making process. At best, it does not alter their average outcomes but may potentially cause consumers to become over confident of the value of a product - which leads to worse consumption outcomes.

The rest of the paper is structured as follows. Section 2.2 discusses the related literature and our contribution to it. We describe our theoretical model in Section 2.3 and our experiment design in Section 2.4. In Section 2.5 lay out our hypotheses, Section 2.6 discuses the results of our experiment, and Section 2.7 concludes.

2.2 Related Literature

The literature on online ratings and reviews has largely focused on implications for firm pricing decisions and consumer purchasing decisions (Cabral & Hortacsu, 2010; Dellarocas et al., 2007; Li et al., 2020a; Luca & Reshef, 2021b; Mayzlin et al., 2014). Experimental evidence on consumer's response to ratings has shown that ratings increases trust between sellers and buyers, as shown in online marketplaces (Cai et al., 2014b; Jin & Kato, 2006b), and in the laboratory (Bolton et al., 2004b; Halliday & Lafky, 2019b; Lafky, 2014b). We contribute to the understanding of how consumers interpret ratings in two ways: (i) we study agents with heterogeneous preferences over a multi-dimensional product; (ii) we explore how different types of information influence how consumers interpret ratings.

Many papers have also studied the generation of ratings, examining when and how raters choose to rate. Unsurprisingly, quality has been shown to be an important driver of ratings (Hui et al., 2023; Li et al., 2020a; Proserpio et al., 2018b; Zhang et al., 2012b).

Lower prices also motivate raters to provide better ratings (Cai et al., 2014b; Li & Hitt, 2010b; Luca & Reshef, 2021b; Neumann & Gutt, 2017). Several other factors can influence rating decisions, this includes reference prices (Gesche, 2022), hidden fees (Chiles, 2021), social distance (Masterov et al., 2015b), and power distance (Gao et al., 2018b). We examine how raters resolve the ambiguity that arises in the presence of heterogeneous preferences, and how that behavior changes as the information available to raters and consumers varies.

Previous literature suggests that social preferences motivates the decision to provide ratings (Bolton et al., 2013b; Chakraborty, Kim, et al., 2022; Chen et al., 2010b; Fradkin et al., 2021b; Halliday & Lafky, 2019b; Qiao et al., 2020). Experiments have more precisely shown that altruism is one such social preference (Hoyer & van Straaten, 2022; Lafky, 2014b). Since ratings are a public good, this connects well to the broader literature suggesting altruism as one motivation for public goods contributions (see Bowles and Polania-Reyes (2012) for a survey). We use altruism as the basis for our theory and design an experiment where other incentives are unlikely to influence subjects to rate.

Zervas et al. (2021b) shows that identical properties on Airbnb and Tripadvisor receive different ratings. There are multiple potential explanations for this, Chevalier and Mayzlin (2006) suggests that there is a buyer-self-selection onto platforms, Zhang et al. (2012b) finds cultural differences in how people rate, and Schneider, Weinmann, Mohr, and vom Brocke (2021b) shows that the way ratings are collected can affect how people rate. We provide a novel explanation, suggesting that the type of information available on the platform can explain the difference in ratings observed on different platforms.

We study how preference similarities between raters and consumers may influence a rating, showing that consumers trust ratings more if they are sent by raters with similar preferences. Previous research has been inconclusive on this matter. Starting with Woodside and Davenport Jr (1974), which shows that similarities between an information provider and user improve trust in the information. However, in an experiment, Connors et al. (2011) shows that proclaimed similarities to the consumer by the rater do not improve the trustworthiness of a rating.

Our theory relates to the literature on cheap talk in communications games following Crawford and Sobel (1982). Few have studied communication games where agents are motivated by altruism. Cooper et al. (1992) and Palfrey and Rosenthal (1991) discusses how altruists can use pre-decision communication to influence the outcomes of one-shot coordination games. Lagunoff and Matsui (2004) studies the role of altruism in intergenerational communication. Most related, Chen et al. (2008) applies their concept of No Incentive To Separate to a model of altruism but considers a situation where receivers pay for a signal. We provide a simple model of communication with no side payments, and where both senders and receivers have the *same objective function*. We show that an altruistic sender's incentive to send a message depends on their beliefs of what receivers believe about their intentions and how receivers interpret the message.

We also contribute to the theoretical literature on reputation (Jullien and Park (2014b) and Peitz and Sobolev (2022), see Bar-Isaac, Tadelis, et al., 2008 for a survey). Most of the literature abstracts from the decision to rate, taking ratings as a given and modelling ratings as increasing in quality (Cabral, 2000b; Tadelis, 1999b). Carnehl et al. (2023) and Sobolev et al. (2021) also consider the impact of prices on ratings. Johnen and Ng (2023a) suggests reciprocity to firms captures why and how raters rate. Here, we propose that a model of altruism to future consumers can explain why and how people rate.

2.3 Theoretical Model

We provide a simple two-player, two-stage sequential game in which raters derive utility not only from their own consumption, but also from the welfare of a future consumer. To fix terminology, we label a consumer who generates ratings the "rater" and a future consumer who uses ratings the "consumer". In the first stage, the rater chooses to leave a rating, and in the second stage, a consumer decides to purchase the product.

Product. We construct a two-dimensional product with attributes, X and Y, with realisations x and y, respectively. The realisations are independently drawn from a commonly known distribution, and reflect the quality of each attribute.

Consumption Utility. Both the rater and consumer have preferences over each of the attributes. They draw two weights, α_i and β_i from some known distribution F_α and F_β respectively. These weights correspond to their preference for the dimensions X and Y respectively, where $i \in \{r, c\}$ represents the rater or the consumer. Consumption utility is given by $u_i = \alpha_i x + \beta_i y - p_i$, where p_i denotes the price of the product. For the consumer, consumption of the product is the only source of utility in the game. Hence, their utility is simply $u_c = \alpha_c x + \beta_c y - p_c$ if they choose to consume the product.

Rater Utility. In addition to their own consumption, the rater also has altruistic concern for the consumer's welfare. Following the literature on social preferences, we represent rater's concern for the consumer by a weight κ on the consumer's consumption utility. Although the rater has an incentive to help the consumer make better choices, we assume the actual act of rating to be burdensome due to the time and effort required. We model this as a cost of rating, e. Finally, $R \in \{R_p, R_n, R_\emptyset\}$ represents a positive, negative, or no rating respectively, and I captures all other information available to the consumer. Therefore, the rater's utility is:

$$u_r = \alpha_r x + \beta_r y - p_r + \mathbb{1}_{\text{E[consumer buys—R,I]}} \cdot \kappa E[u_c|x,y] - \mathbb{1}_{\text{rater rates}} \cdot e.$$

This captures that raters utility takes two parts. First, they receive some consumption utility. Second, they receive an additional utility that depends on the expected utility of future consumers.

Notice that the additional utility only exists if the rater expects the future consumer to purchase the product. This utility may either be positive or negative, and corresponds to the expected utility that a future consumer would receive from purchasing the product. If raters anticipate that consumers would not purchase the product, then the additional utility is zero. This means higher order beliefs are important here as a rater's decision to rate depends on their beliefs over how consumers interpret ratings. This belief dependent framework is a representation of psychological games, and we adopt a psychological nash equilibrium (Geanakoplos et al., 1989).

Consumer's Decision. Consumers decide to purchase a product if their expected utility, conditional on the information available, is weakly positive. That is $E_c[u_c|R,I] \geq 0$. This means that consumers decisions also depend on their beliefs over how raters interpret ratings. Hence, in equilibrium, both raters and consumers decision depends on higher order beliefs of how the other thinks they interpret a rating.

Rater's Decision. Prior to the beginning of the game, the rater consumes the product and learns the true quality of both attributes, observing the exact values of x and y. The rater's choice is to either leave a rating, $R \in \{R_p, R_n\}$ at some cost e, or do nothing $R \in \{R_\emptyset\}$. For simplicity we model a binary rating choice (or trinary, including the non-rating $R_\emptyset$) with R_p representing a positive signal and R_n a negative one.

An altruistic rater rates for one of two reasons: to mitigate harm; or generate benefit. Raters mitigate harm when a consumer would purchase a product in the absence of ratings, but the rater anticipates that such a purchase would generate a negative utility. Hence, the rater chooses to leave a bad rating to warm consumers away from the product. Conversely, raters generate benefit when a consumer would not purchase a product in the absence of ratings, but the rater anticipates that such a purchase would generate a positive utility. Here, raters choose a good rating to encourage consumers to purchase the product.

To mitigate harm and warn the consumer away from a "bad" product, three equations must be met for the rater to rate negatively:

$$\kappa E_r[u_c|x, y] < -e \tag{2.1a}$$

$$R_r[E_c[u_c|R_\emptyset, I]] \geq 0 \tag{2.1b}$$

$$E_r[E_c[u_c|R_n, I]] < 0 \tag{2.1c}$$

Equation (2.1a) captures that raters expect consumers to receive a negative consumption utility, and their concern for the consumer's utility exceeds the cost of leaving a rating. Equation (2.1b) shows that in the absence of a rating, the rater expects that consumers choose to consume the product based on all other information available. Equation (2.1c) reflects how raters expect consumers not to consume the product if consumers observe a negative rating alongside all other available information.

These three equations can be further summarised into two conditions. The equation (2.1a) represents the expected utility gained by the rater from leaving a rating. And equations (2.1b) and (2.1c) represent the ability of a rating to change the purchase decision of consumers.

There are three corresponding equations that must be met for the rater to rate positively:

$$\kappa E_r[u_c|x, y] > e \tag{2.2a}$$

$$E_r[E_c[u_c|R_\emptyset, I]] < 0 \tag{2.2b}$$

$$E_r[E_c[u_c|R_p, I]] \geq 0 \tag{2.2c}$$

Likewise, equation (2.2a) shows the rater expects that the consumer receives positive consumption utility, and their concern for the consumer's potential gain in utility exceeds the cost of leaving a rating. Equation (2.2b) shows the rater expects that if the consumer does not observe a rating they will choose not to consume the product based on all other information available. And equation (2.2c) reflects that the rater expects the consumer will choose to consume the product if they observe a positive rating alongside all other information available.

Information Environment. We vary the information available to the consumer on two dimensions, creating four distinct information environments. Our first environment serves as a minimal information baseline setting in which all information about preferences and product quality is private, and the only potential signal the consumer receives is through a rating. This means $I = \emptyset$. Our second information environment is identical to the first except that the consumer also observes the preferences (α_r, β_r) of the rater. This means that $I = \{\alpha_r, \beta_r\}$. This is akin to consumers being able to read reviews to learn about why a rater chose a certain rating.

The third environment is identical to the first, except that the quality level of one dimension of the product is common knowledge. This means $I \in \{x, y\}$. This is akin to products having a verifiable component, such as the location of a hotel. The fourth and final environment simply combines the previous two, allowing consumers to observe both the preference of the rater, and the quality level of one dimension of the product.

2.3.1 Theoretical Findings

From this theoretical environment, we provide some intuitive claims without formal proof. We rely upon these claims to guide the formation of our hypotheses.

Observe that three equations need to be simultaneously satisfied for raters to provide either a positive or a negative rating. The expected utility gained by the rater from leaving a rating has to be weakly larger than the cost of leaving a rating, (2.1a) and (2.2a). By leaving a rating, the rater is able to change the decision of the consumer, (2.1b) and (2.1c), and (2.2b) and (2.2c).

In all information environments, (2.1a) and (2.2a) are more readily satisfied when the quality level of each dimension is more extreme as the potential gain in utility for the rater is larger the further the quality level is from its unconditional expectation.

Second, raters only incur the cost of sending a rating if they can change the decision of consumers. Hence, ratings are only sent if raters anticipate that positive (or negative) ratings are interpreted by consumers to represent a product of significantly higher (lower) quality than no rating. In equilibrium, this means positive (negative) ratings should only be sent for products of sufficiently high (low) quality.

Hence, this leads to our first claim:

Claim 1. *In every information environment, positive (negative) ratings are sent when a product's attributes are of a sufficiently high (low) quality.*

We now discuss our predictions for how the rater's decisions would be influenced by the information environment. We focus on how both rater and consumer behavior responds to having more information.

We first consider the effect of making rater preferences public knowledge. Here, our theory suggests the existence of multiple equilibria. When consumers learn raters preferences, this does not directly improve their knowledge about product quality. Instead, learning raters preferences guide how consumers form beliefs. When choosing to rate, it is unclear how raters should expect consumers to update their beliefs. As such, any set of beliefs in which raters and consumers agree is a potential equilibrium. This means there are many potential equilibria.

Of the potential equilibria, two stand out. One where raters rate based on their own preferences, and another where raters rate based on preferences opposite their own. These are the most extreme interpretations of how knowing rater preference may modify the beliefs of consumers.

Claim 2. *When rater preferences are common knowledge, raters may choose to rate as a reflection of a) the attribute they prefer; b) the attribute they less prefer.*

Neither equilibrium is intrinsically more attractive or efficient. However, raters own preferences are likely most focal when rating. Since rater's preference is the only information aside from the rating available to consumers, we conjecture that consumers will interpret ratings in the context of the rater's preference.

Next, we consider the situation where the quality of one dimension of the product is common knowledge. This means consumers have perfect information about this attribute. If ratings take into account this attribute, they do not improve the information consumers have. Therefore, most informative ratings should only take into account the the quality of the dimension of the product that is not common knowledge. Raters, wanting to provide the most informative ratings, should thus only provide ratings based on the dimension of the product that is not common knowledge.

Claim 3. *When one dimension of the product is common knowledge, raters will rate solely based on the unknown dimension.*

We conduct an experiment to test how raters and consumers solve the equilibrium selection problem, and how that solution responds to different information environments.

2.4 Experimental Design

We design a two-part laboratory experiment to examine how raters and consumers resolve the ambiguity arising from ratings in the presence of heterogeneous preferences. The experiment was conducted using the oTree software on Prolific (Chen et al., 2016). A total of 502 subjects were recruited from a sample of the US population. Each subject read a series of instructions, completed a comprehension quiz, and then played 20 rounds of the experiment. Subjects were paid a show up fee of \$3.00 USD (equivalent to \$12.00 USD per hour), and were able to earn bonus payments in tokens, where each token was worth \$0.50 USD. The average subject took just above 13 minutes to complete the experiment and received a total payment (show up fee and bonus) of \$6.37 USD. There were minimal differences in the instructions between treatments, which are highlighted in Appendix B.1.

Each treatment took place in two phases. In the first phase, we recruited subjects to play the role of senders who rated randomly selected products. After all data from the first phase was collected, we conducted the second phase. In the second phase, we recruited a new group of subjects, excluding those who participated in the first phase, to act as receivers who viewed ratings generated in the first phase. In each phase, subjects participated in 20 rounds of decision making. At the beginning of the 20 rounds, they were randomly assigned a preference for X or Y. This preference was fixed for all rounds.

Senders. Subjects in the sender role were then asked to evaluate a two-dimensional product referred to in the instructions as a "prize". The product comprised of two values, X and Y, and both values were drawn from a discrete uniform distribution, $U\{1, 10\}$. Subjects who were assigned a preference for X valued the product as $1 \times x + 0.1 \times y$. Likewise, subjects who were assigned a preference for Y valued the same product as $0.1 \times x + 1 \times y$.

Senders were informed that receivers would also be randomly assigned to have a preference for either X or Y. They were also informed of the receiver's task, and what information receivers would know prior to making their own decisions, which varied by treatment as described below.

In each of the 20 rounds, senders were tasked to rate a prize on a scale of 1 to 5, where 1 was stated to be the worst rating, and 5 the best. Although such a rating scale differs slightly from our theory which uses only positive and negative ratings, we prefer this scale as it is more natural for participants.

After rating the prize, senders were given the choice to share their rating with receivers for a small fee of 0.1 tokens ($0.05 USD). At the end of the 20 rounds, one round was randomly selected and the value of the prize from that round, less any cost of sending a rating in that round, was paid to the sender.

Notice that although our theory is based on a model of rater altruism, we choose not to explicitly impose such a utility function on raters. Our experiment design minimises potential effects from other explanations for why people rate. Hence, where raters do choose to rate, we expect that this is out of altruism. Since we anticipate that altruistic preferences are intrinsic, there is no need to explicitly impose such a utility.

Receivers. Like Senders, each receiver was assigned either a preference for X, valuing the prize at $1 \times x + 0.1 \times y$, or a preference for Y, valuing the same prize at $0.1 \times x + 1 \times y$. Receivers learned about the sender's task and and that senders were randomly assigned a preference for X or Y. Receivers were further split into two groups. One group of receivers were shown only prizes for which senders sent ratings. The other group of receivers were shown only prizes for which ratings were not sent. Receivers knew in advance which group they belonged to, but they were not informed of the existence of the other group.[1]

Each receiver sampled 20 different prizes randomly chosen from the all prizes with sent (or unsent) ratings. Sampling was without replacement for each receiver, but with replacement across receivers. In other words, all receivers independently faced the same probability of drawing any one prize from the prizes with sent (or unsent) ratings.

Each receiver next indicated a willingness to pay for that round's prize via a Becker-DeGroot-Marschak (henceforth BDM) Mechanism (Becker et al., 1964) similar to Healy (2018). In this implementation, receivers were given a series of questions asking them to choose between the prize or an increasing number of tokens and were asked to state the question at which they would begin to choose tokens over the prize. They were allowed to respond in the range of 0.1 to 11.0 tokens, in increments of 0.1.

At the end of each round, receivers were told the actual values of X and Y. At the end of the 20 rounds, one round was randomly selected and their decision on one randomly selected question from that round's BDM was paid to the receiver.

Treatments. The experiment implements a 2×2 design, where information is varied over two dimensions. Our first treatment variable modifies the information provided to the receiver about the sender. In the *Public* treatments receivers learned which of the two values the rater preferred (X or Y) prior to stating their WTP. In the *Private* treatments the receivers never learned the preferences of the sender.

Our second treatment variable modifies the information provided to the receiver about the product. In the *Partial* treatments receivers directly observed the X-value for the product, prior to indicating their willingness to pay. In the *None* treatments receivers did not directly observe any information about the product quality. Note that the revealing only the X-value in the Partial treatment was intentionally asymmetrical, resulting in some consumers (X-types) who had consistently better information than other consumers (Y-types).

All combinations of our treatment variables gives us four treatments: (i) Private-None, where receivers learn nothing about sender preferences or product quality; (ii) Public-None, where receivers learn which of the two dimensions the sender preferred but nothing about product quality; (iii) Private-Partial, where receivers do not learn sender preferences but do learn the value of x; (iv) Public-Partial, where receivers learn both sender preferences and the value x takes. This is summarised in table 2.1.

[1] We decide to over sample sent ratings because we anticipated that ratings were sent with a probability of less than 50%.

		Product Information	
		No	Yes
Rater Information	No	(i) Private-None	(iii) Private-Partial
	Yes	(ii) Public-None	(iv) Public-Partial

Table 2.1: Treatments

2.5 Hypotheses

Our first three hypotheses focus on how raters choose to rate. This is important as it addresses how raters resolve the equilibrium selection problem: either rating based on their own preferences or on the preferences of others. The first hypothesis is that raters will rate based on their own experience when ratings are the only information revealed to consumers. This follows from Claim 1.

Hypothesis 1. *In the baseline Private-None treatment, ratings are more likely to reflect the attribute the rater prefers. Hence, the coefficient of the quality of the attribute the rater prefers will be larger than the coefficient of the attribute they prefer less.*

Our second hypothesis is about how making raters preferences public affects the way they rate. Although Claim 2 suggests multiple equilibria, we anticipate that information about the rater's preferences anchors how raters and consumers interpret ratings. Hence, our second hypothesis is that making rater preferences public causes those preferences to become more focal, resulting in ratings which more closely reflect the rater's preferred attribute.

Hypothesis 2. *Ratings are more sensitive to the rater's preferred attribute if their preferences are public rather than private. Hence, the coefficient of the quality of the attribute the rater prefers is larger in the public than private setting, while the coefficient of the attribute they prefer less will be smaller.*

Claim 3 discusses that when the quality of one attribute is perfectly known to consumers, raters should not take that attribute into account when choosing a rating. This leads to our third hypothesis:

Hypothesis 3. *Ratings are not affected by the revealed attribute in the Partial treatments. Hence, the coefficient of the $X-value$ will be insignificant while the $Y-value$ is positive and significant.*

The next two hypotheses focus on how consumers interpret ratings. They reflect how consumers, with the correct beliefs of how raters choose to rate, will interpret ratings. Our fourth hypothesis focuses on how consumers interpret ratings when raters preferences are public. If consumers anticipate that raters are more likely to rate along the attribute they prefer, consumers who share the same preference are more likely to trust ratings. Conversely, consumers who do not have the same preference as raters will be less willing to use ratings.

Hypothesis 4. *In the Public treatments, willingness to pay will be more sensitive to ratings for consumers who share the rater's preferences than for consumers with the opposite preferences. Hence, when consumers share the same preference as raters, the coefficient on ratings will be larger.*

When the quality of one attribute is revealed, consumers who prefer this attribute have little to gain from information about the unrevealed attribute. Likewise, consumers who prefer the unrevealed attribute are more likely to take the rating into account when reporting their willingness to pay. Our final hypothesis captures this.

Hypothesis 5. *In the Partial treatments, willingness to pay will be more sensitive to ratings for consumers who prefer the unrevealed attribute than for those who prefer the revealed attribute. Hence, when consumers prefer X the coefficient on ratings will be insignificant, while the coefficient on rating will be positive and significant when they prefer Y.*

2.6 Results

Table 2.2 provides summary statistics for the experiment. The means reported in the table are given at the decision level, while our analysis is done at the subject level. We recruited a total of 502 subjects, each of whom played 20 rounds. Raters made 2 decisions, the decision of a rating, and if their rating is sent. Consumers reported their willingness to pay via a BDM.[2]

We provide our analysis in two parts, studying rater behaviour followed by consumer behaviour. Studying their behaviour independently allows us to understand how raters and consumers belief ratings should be interpreted. We may then shed light on the consistency of beliefs between raters and consumers, and discuss if, when and how ratings are informative.

	Private None	**Public None**	**Private Partial**	**Public Partial**
Raters				
Mean rating	3.25	3.16	3.26	3.27
	(1.38)	(1.35)	(1.30)	(1.33)
Ratings sent (%)	23	28	20	35
Mean sent rating	3.87	3.57	3.59	3.84
	(1.21)	(1.41)	(1.45)	(1.27)
Mean sent X value	5.77	6.18	5.56	6.02
	(2.96)	(2.85)	(2.89)	(2.84)
Mean sent Y value	6.53	5.90	6.38	6.04
	(2.81)	(2.90)	(2.97)	(2.80)
Subjects	51	51	50	50
Consumers (with ratings)				
Mean rating	3.44	3.76	4.15	3.99
	(1.57)	(1.49)	(1.11)	(1.21)
Mean WTP	54.65	54.39	60.40	61.00
	(29.54)	(31.77)	(32.54)	(27.96)
Mean X value	5.09	6.84	6.80	6.54
	(3.18)	(2.92)	(2.36)	(2.90)
Mean Y value	6.73	5.76	5.71	6.06
	(2.68)	(2.75)	(2.71)	(2.47)
Subjects	50	50	50	50
Consumers (without ratings)				
Mean rating	3.22	3.77	3.47	3.13
	(1.40)	(1.44)	(1.28)	(1.15)
Mean WTP	52.46	54.05	52.49	49.42
	(28.33)	(27.90)	(32.41)	(30.50)
Mean X value	5.60	6.77	5.81	5.16
	(2.90)	(2.88)	(2.91)	(2.62)
Mean Y value	5.65	5.84	5.50	5.33
	(2.78)	(2.83)	(2.76)	(2.51)
Subjects	25	25	25	25

Table 2.2: Summary Statistics

2.6.1 Rater Behavior

We begin our analysis by examining rater behavior, first in terms of how raters selected their ratings and then how they decided to send those ratings to consumers. Figure 2.1 summarises how raters chose to rate when they shared their ratings with consumers, showing that they generally preferred to share better ratings when they received a product which they valued more. Table 2.3 reports regression results for rater behavior. Columns 1-4 examine how raters selected their ratings, using the choice of rating as the dependent variable. Columns 5 and 6 show the decision to send ratings, with a binary dependent

[2] 51 subjects were recruited in the Private-None and Public-None treatments for raters due to a misunderstanding of how to increase the subject pool size in prolific. 25 subjects were recruited when to act as consumers who received no ratings. This was deliberate as there would be less variation in the information they received.

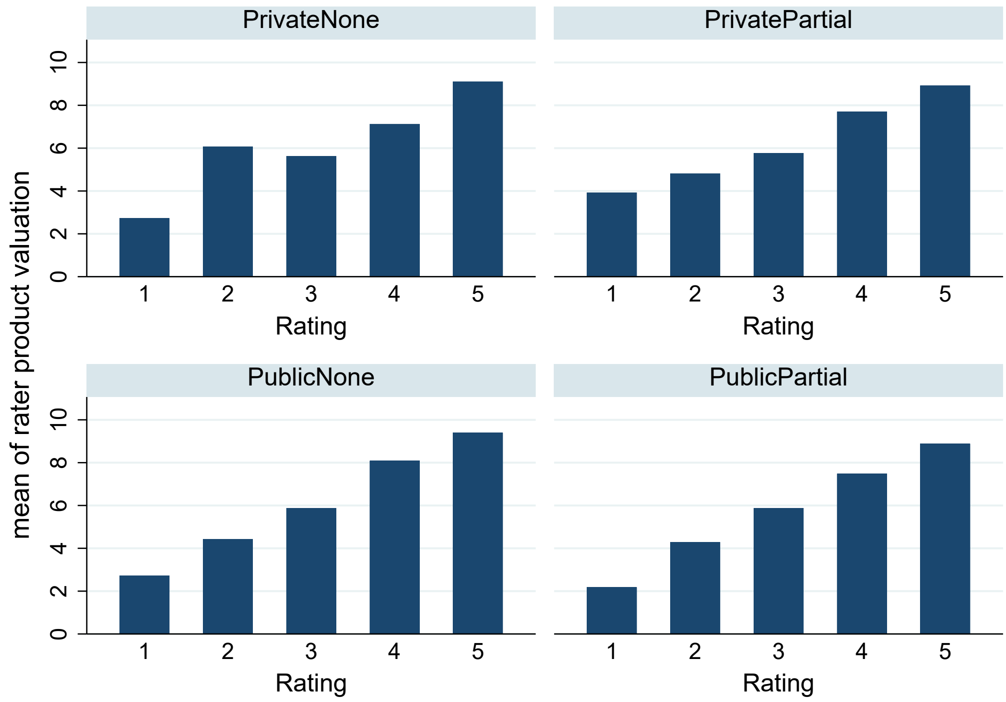

Figure 2.1: Mean of rater's valuation of product for each rating, if they chose to send a rating.

variable describing whether the rating was sent. The variable *ownvalue* describes the quality level for the rater's preferred attribute (X for X-types, Y for Y-types) while the variable *othervalue* captures the rater's less-preferred attribute (Y for X-types, X for Y-types). We include dummies *public* and *partial* for each of our treatment variables.

In all treatments, raters are primarily motivated by their preferred attribute (*ownvalue*), while showing little concern for their less-preferred attribute (*othervalue*). In other words, raters rated almost exclusively based on their own preferences, providing support to hypothesis 1. Interacting *public* with *ownvalue* and *othervalue* shows this effect was stronger in the public treatments, when raters knew that their own preferences were visible to consumers, consistent with hypothesis 2. The differences between all ratings (columns 1 and 3) and sent ratings (columns 2 and 4) can be attributed to the non-uniform decision to send ratings.

The probability of sending a rating is also sensitive to the rater's preferred attribute, and only minimally so for their less-preferred attribute. Raters were more likely to send ratings when their preference is public knowledge. From our theory, this suggests that raters may anticipate ratings to be more likely to change consumers decisions when the rater's preference is known.

Table 2.4 focuses on understanding how making the X-value common knowledge affects the rater's decisions. In particular, the coefficient on the X-value is negative and significant when interacted with *partial* dummy. This shows that the choice of rating is less affected by the X-value when it is common knowledge. However, this effect is small and goes away when using an ordered probit model. Hence, we do not conclude that hypothesis 3 is supported.

	(1) tobit All ratings	(2) tobit Sent ratings	(3) probit All ratings	(4) probit Sent ratings	(5) LPM Send choice	(6) Probit
ownvalue	0.47***	0.44***	0.45***	0.37***	0.022***	0.14***
	(0.012)	(0.031)	(0.034)	(0.067)	(0.0055)	(0.028)
othervalue	0.069***	0.061**	0.065**	0.053	0.0071*	0.042*
	(0.012)	(0.029)	(0.026)	(0.041)	(0.0039)	(0.023)
public	0.12	-1.01***	0.096	-0.88**	0.053	0.62*
	(0.14)	(0.35)	(0.21)	(0.43)	(0.054)	(0.37)
partial	0.24*	0.56	0.20	0.46	0.028	0.35
	(0.14)	(0.35)	(0.21)	(0.43)	(0.054)	(0.37)
public×ownvalue	0.048***	0.18***	0.046	0.15**	0.016**	0.050
	(0.014)	(0.035)	(0.031)	(0.059)	(0.0075)	(0.039)
public×othervalue	-0.074***	-0.012	-0.070**	-0.0089	-0.0073	-0.041
	(0.014)	(0.032)	(0.029)	(0.044)	(0.0051)	(0.029)
partial×ownvalue	-0.021	0.0010	-0.015	0.0063	0.0063	0.017
	(0.014)	(0.035)	(0.031)	(0.058)	(0.0075)	(0.039)
partial×othervalue	-0.0052	-0.076**	-0.0052	-0.067	-0.0066	-0.035
	(0.014)	(0.032)	(0.029)	(0.042)	(0.0051)	(0.029)
round	-0.014***	-0.024***	-0.014***	-0.021**	-0.0018*	-0.010*
	(0.0033)	(0.0077)	(0.0040)	(0.0089)	(0.0010)	(0.0060)
constant	0.46***	1.13***	0.26***	0.58***	0.058	-2.66***
	(0.13)	(0.33)	(0.043)	(0.19)	(0.040)	(0.32)

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.3: Rater decisions. Columns 1-4 are choice of rating, while columns 5 and 6 are the binary decision of whether to send a rating. Columns 1 and 2 are tobits, 3 and 4 are ordered probits. Columns 1 and 3 are for all ratings, columns 2 and 4 are only for ratings that were sent. Column 5 is a linear probability model while column 6 is a probit.

2.6.2 Consumer Behavior

Our results on consumer behavior begin with Figure 2.2, which shows how consumer willingness to pay (WTP) varies with ratings across treatments. Consumers generally reported a higher WTP when they observed better ratings, except in the Private-Partial treatment. This is expected, since we anticipate that consumers who prefer X will be less likely to rely on ratings. We report consumer WTP decisions in Table 2.5. Column 1 shows how ratings influence consumers WTP along the two treatment variables. We find

	(1) tobit-sent	(2) probit-sent
X-value	0.31***	0.23***
	(0.032)	(0.041)
Y-value	0.28***	0.18***
	(0.031)	(0.040)
partial	0.90**	0.55
	(0.45)	(0.39)
partial×X-value	-0.12***	-0.072
	(0.044)	(0.054)
partial×Y-value	-0.014	-0.0087
	(0.043)	(0.057)
round	-0.019*	-0.012*
	(0.010)	(0.0066)
constant	0.53	0.49
	(0.35)	(0.13)
Observations	1071	1071

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.4: Rater decisions, focusing on the effect of X-value and Y-value on sent ratings with partial information (X-value is common knowledge). Column 1 is tobit, and column 2 is ordered probit.

that consumers respond positively to favorable ratings, reporting a higher WTP and, on average, always report a higher WTP when there is more information. In other words, in Public and Partial treatments, consumers have a higher WTP. This may be the case that ratings provide consumers with a sense of certainty, or assurance of product value. Notice in the Private-None case, at all ratings levels, raters interpret ratings as a higher value than what consumers interpret (first panel in Figure 2.1 and Figure 2.2. Moreover, average WTP is less sensitive to ratings in public treatments.

Figure 2.2 also shows that when the rating of 3 is observed, consumers WTP is indistinguishable across all treatments. Considering the rating of 3 as a benchmark, we categorise ratings as favorable (4 and 5) and unfavorable (1 and 2). The dummy variable *binfavrating* represents this. Column 2 in Table 2.5 shows consumer WTP is higher when they observe favorable ratings, and overall higher in Partial treatments. Importantly, we show when more information is available, WTP become less sensitive to ratings. This suggests that when consumers have additional information, they rely less on ratings when evaluating a product.

To further investigate why WTP becomes less sensitive to ratings when more information

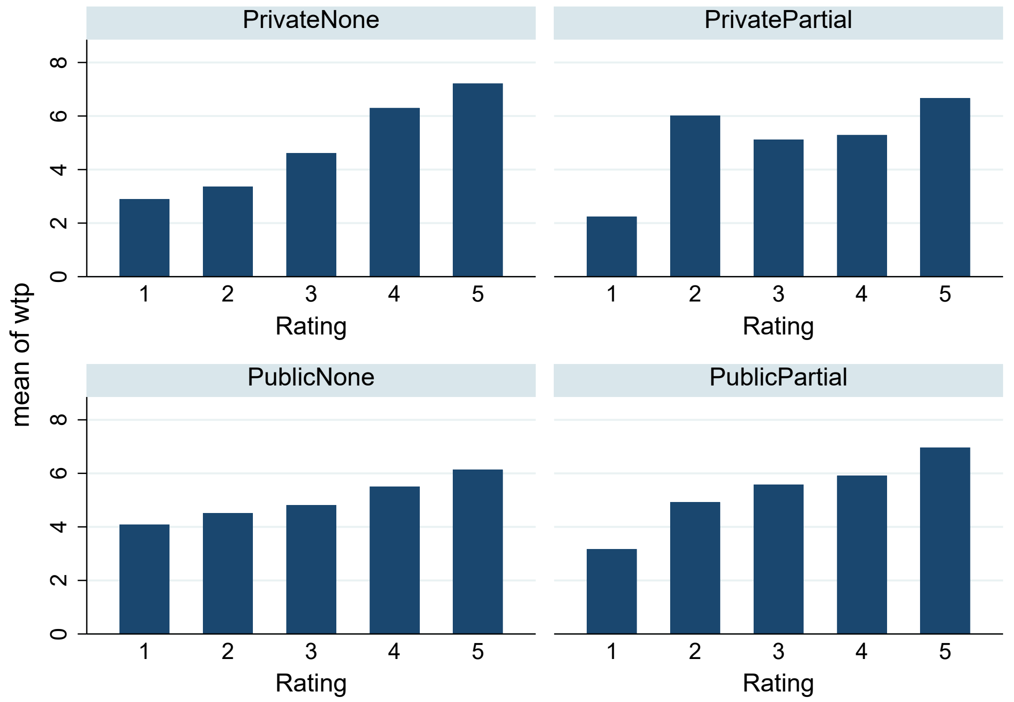

Figure 2.2: Mean WTP for each rating observed, across treatments.

is available, we consider Public and Partial treatments separately. Theory suggests that raters preferences can affect consumer WTP in public treatments. In particular, if the consumer shares the same preference as the rater, their WTP will be more sensitive to the rating. Table 2.6 reports consumer WTP decisions when considering only Public treatments (column 1 and 3) or Partial treatments (column 2 and 4). The variable *rtypex* is a dummy variable for consumers with a preference for X, while the variable *sametype* is a dummy variable for when the consumer had the same preference as the rater whose rating was observed.

In public treatments, columns 1 and 3 show consumers WTP were more sensitive to ratings when they shared the same preferences as the rater. This aligns with the theory, providing support for hypothesis 4. Because we find support for both hypotheses 2 and 4, we show that despite the potential for multiple equilibria (Claim 2), raters and consumers agree that raters should rate based on their own preferences when these preferences are publicly known. This common belief resolves any ambiguity that arises in the interpretation of ratings when rater's preferences can be observed.

In partial treatments, the quality of one dimension of the product is common knowledge. Theory suggests the WTP of consumers who prefer this dimension will be less sensitive to ratings. Columns 2 and 4 in Table 2.6 reports this analysis. We find that consumers who preferred X had, on average, a higher WTP. However, we see no effect from interacting either consumer preferences or X-value with ratings. In other words, consumers who preferred X were willing to pay more than those who preferred Y in partial treatments, but consumers with either preference did not use ratings differently. Hence, we do not find support for hypothesis 5.

Although we do not support both hypotheses 3 and 5, our results can still shed some light on what raters and consumers belief ratings should represent in partial treatments. Raters place less emphasis on the X-value when they rate (column 1 in Table 2.4). However, all consumers do not change the way they use ratings (column 2 in Table 2.6). This suggests that in such settings, raters and consumers may not agree on what a rating should represent.

2.6.3 Efficiency of ratings

We have shed some light on how consumers interpret ratings, and how other information can influence the way they interpret ratings. However, can information help consumers make better consumption decisions?

We use a BDM to elicit the willingness to pay of consumers. This is akin to a second price sealed bid auction, and consumers have an incentive to reveal their true valuation of the product on offer. Using the difference between the actual value of the product and the consumer valuation, we construct a series of variables to describe how effective ratings are at helping consumers make decisions. Table 2.7 shows these results. Columns 1 and 2 shows how ratings and the ratings environment can affect the difference between the actual value of the product and the consumers WTP (*efficiency*). Columns 3 and 4 repeats this analysis on the absolute difference between the actual value of the product and the consumers WTP. Columns 5 and 6 analyses the probability of receiving a positive difference between the actual value of the product and consumers WTP.

Column 1 and 2 show that consumers, on average, underestimate the product value, and ratings have little impact on how effective consumers are at valuing the product. In the public treatments, consumers tend to overestimate the value of product. However, when observing better ratings, consumers are more likely to underestimate product value. In partial treatments, consumers were more likely to overestimate product value when observing better ratings. Columns 5 and 6 report a similar result when considering the probability of consumers receiving a positive outcome.

Our results imply that when consumers know the preferences of raters, they are more likely to purchase products which are of little value to themselves. However, it is important to observe that this negative effect is mitigated when ratings send a positive signal - suggesting that only poorer rated products are overvalued. Moreover, when consumers

	(1) All Ratings	(2) Favorable Ratings
rating	0.93***	
	(0.100)	
public	0.95**	0.43
	(0.43)	(0.33)
partial	0.75*	1.20***
	(0.44)	(0.34)
public×rating	-0.26**	
	(0.12)	
partial×rating	-0.13	
	(0.12)	
round	0.032***	0.033***
	(0.0081)	(0.0089)
binfavrating		3.07***
		(0.34)
binfavrating×public		-0.67*
		(0.38)
binfavrating×partial		-1.21***
		(0.38)
constant	1.80***	3.07***
	(0.34)	(0.25)
Observations	4000	3619

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.5: Consumer WTP, only for consumers with ratings. Column 1 uses all ratings, while column 2 uses a binary variable that is 1 for ratings of 4 or 5 and 0 for ratings of 1 or 2.

can verify the quality of some product attributes, they tend to overvalue such products when observing better ratings. Hence, we find that additional information generally harms consumer decision making, and is only beneficial when they observe a good rating and know the preferences of the rater who sent that rating.

Although consumers continue to misestimate the value of the product in both public and partial treatments, we show that additional information can be useful. Column 3 and 4 show how the degree of misestimation changes in each treatment. Particularly, partial product information can decrease the magnitude of the estimation error. Moreover, information about the raters type may reduce estimation error when ratings are poor, but this

	(1) Public	(2) Partial	(3) Public	(4) Partial
rating	0.21**	0.59***		
	(0.099)	(0.13)		
rtypex	-0.32	1.66***	-0.15	1.95***
	(0.63)	(0.58)	(0.43)	(0.44)
sametype	-2.20***		-1.33***	
	(0.54)		(0.35)	
rtypex×rating	0.11	-0.19		
	(0.16)	(0.15)		
sametype×rating	0.69***			
	(0.14)			
binfavrating			0.48	2.06***
			(0.34)	(0.44)
rtypex×binfavrating			0.26	-1.26***
			(0.48)	(0.44)
sametype×binfavrating			2.43***	
			(0.43)	
X-value		0.39***		0.41***
		(0.080)		(0.062)
rating×X-value		-0.012		
		(0.016)		
binfavrating×X-value				-0.035
				(0.053)
round	0.040***	0.029***	0.043***	0.029**
	(0.011)	(0.011)	(0.012)	(0.012)
constant	4.27***	0.63	4.71***	1.10**
	(0.42)	(0.54)	(0.34)	(0.50)
Observations	2000	2000	1837	1880

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.6: Consumer WTP, only for consumers with ratings. Column 1 and 3 are restricted to the public treatments, while column 2 and 4 are restricted to the partial treatments. Columns 1 and 2 use all ratings, columns 3 and 4 uses a binary variable that is 1 for ratings of 4 or 5 and 0 for ratings of 1 or 2.

effect diminishes with better ratings. This result suggests that even though consumers may make poor consumption decisions, and purchase products that have little value to themselves, additional information can minimise their loss in utility.

	(1)	(2)	(3)	(4)	(5)	(6)
	efficiency		|efficiency|		efficiency	probability
rating	-0.015		-0.0073		-0.092**	
	(0.097)		(0.066)		(0.041)	
public	-2.30***	-0.70***	-0.76***	-0.14	-1.02***	-0.44***
	(0.42)	(0.33)	(0.28)	(0.22)	(0.21)	(0.16)
partial	0.33	0.40	-0.92***	-0.58***	0.24	0.19
	(0.41)	(0.32)	(0.26)	(0.22)	(0.22)	(0.16)
public×rating	0.62***		0.18**		0.26***	
	(0.11)		(0.076)		(0.049)	
partial×rating	-0.19*		0.038		-0.086*	
	(0.12)		(0.072)		(0.051)	
binfavrating		0.011		-0.0049		-0.31**
		(0.34)		(0.21)		(0.15)
public×binfavrating		1.54***		0.37		0.71***
		(0.36)		(0.23)		(0.17)
partial×binfavrating		-0.92***		-0.088		-0.33**
		(0.35)		(0.21)		(0.16)
round	-0.033***	-0.033***	-0.019***	-0.022***	-0.014***	-0.014***
	(0.0087)	(0.0094)	(0.0064)	(0.0069)	(0.0039)	(0.0042)
constant	1.65***	1.30***	3.22***	3.04***	0.88***	0.70***
	(0.35)	(0.30)	(0.23)	(0.18)	(0.17)	(0.15)
Observations	4000	3619	4000	3619	3866	3493

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.7: Efficiency of ratings. Columns 1 and 2 takes the difference between the actual value of the product and consumer willingness to pay. Columns 3 and 4 take the absolute value of this difference. Columns 5 and 6 the dependent variable takes the value of 1 when the difference is positive, and 0 if it is negative. Columns 1, 3 and 5 use all ratings, while 2, 4 and 6 uses a binary variable that is 1 for ratings of 4 or 5 and 0 for ratings of 1 or 2.

2.7 Conclusion

In this paper, we show what raters and consumers think ratings should represent in environments with different types of information. We consider individuals with heterogeneous preferences, preferring one of two dimensions in a two-dimensional product. We propose a theory of rater altruism to future consumers, and design an experiment to show that raters are indeed altruistic to consumers. When the rating is the only information available to consumers, we show that raters and consumers expect better ratings to reflect higher quality products.

We consider 2 types of information: (i) information about rater's preferences, which can be learned by consumers through reading reviews or checking a rater's profile; (ii) verifiable information about product quality, which can describe elements of a product that are fixed, such as the location of a hotel, or a third party verification mark.

When rater's preferences are known, our theory predicts the possibility of multiple equilibria. We show that raters decide to provide ratings that are more sensitive to their own preferences. Importantly, consumers do anticipate how raters choose to rate, and report willingness to pay that are more sensitive to ratings when they have the same preferences as raters. We also show that such information can lead to less spread between product value and consumer willingness to pay, but consumers tend to overvalue products more often.

When there is verifiable product information, we predict that ratings are a reflection of unverifiable product quality. We find that such information may have a small effect on rater behaviour. However, consumers do not change the way they use ratings when verifiable product information is available. This suggests that there might be some disagreement on what a rating should represent. Despite this disagreement, we show that such information can also lead to less spread between product value and consumer willingness to pay, and only cause consumers to overvalue products when ratings are favorable.

Our results hence suggest that the design of rating systems depends on what the goal of the designer is. If designers have the goal of reducing uncertainty of product quality, they should make more information available to consumers, particularly information about the preferences of other consumers. However, if designers aim to help consumers make positive consumption decisions, this can be counter productive, as it will cause consumers to overvalue products with a poor rating, and hence low quality products more often.

Chapter 3

Free and Open-Source Software: Coordination and Competition

Abstract

Free and Open-Source Software (FOSS) are developed by a community of developers led by a coordinator. Coordinators balance the following trade-off: (i) more developers improve FOSS' quality - a positive vertical differentiation effect; (ii) more developers lead to more diverse views, driving FOSS characteristics away from existing developers' preferences - a negative horizontal differentiation effect. Permissive Open-Source licenses can intensify competition when FOSS compete with proprietary software, resulting in lower prices. I study coordinators with different motivations: self-interested, altruistic, and profit-driven. Altruistic and profit-driven coordinators fail to maximize total surplus, while self-interested coordinators generate higher total surplus than altruistic coordinators.

JEL: D21, D26, L14, L17

Keywords: Open-Source Software, Network effects, Licensing

*CORE, Universite Catholique de Louvain, robin@robin-ng.com.

I am especially grateful to Paul Belleflamme, Johannes Johnen, Rachel Tan, and Greg Taylor for many fruitful discussions. I thank Francis Bloch, Julius Goedde, Vili Lehdonvirta, Francois Maniquet, Erika Pini, Manpreet Singh, Jakob Wegmann, and many seminar participants for their comments. This project was realised with funding from the Fédération Wallonie-Bruxelles through the Action de recherche concertée grant 19/24-101 'PROSEco,' and the FNRS-ASP research fellowship.

3.1 Introduction

Free and Open-Source Software (FOSS) are a bedrock of the digital age. They range from prominent household names like Google's Android to niche and independently developed programs like Paint.NET. These software contribute to at least 2% of the global GDP (Blind & Schubert, 2023), and events in the FOSS community can have a wide-ranging impact. For example, the USA's 'Securing Open Source Software Act of 2023' was triggered by a FOSS vulnerability.[1] Although many firms and individuals adopt FOSS, few ever fund it.[2] Even the most prominent FOSS receive funding below the USA poverty line.[3] Developers of FOSS are effectively volunteers and have no obligation to develop any aspect of the software (Franke & Von Hippel, 2003; Hars & Ou, 2002; Shah, 2006). This community-driven nature of FOSS creates difficulties for coordinators in coordinating its development (Michlmayr et al., 2007).[4] Despite the lack of funding, and the challenges of coordinating FOSS development, many still choose to do so.

Prior research has chosen to abstract away from the decision to coordinate the development of FOSS. Instead, they focus on the private motives of developers contributing to a public good (Athey & Ellison, 2014; Casadesus-Masanell & Ghemawat, 2006; Johnson, 2006; Johnson, 2002). I contribute to this body of work by considering how the motives of coordinators influence the development of FOSS. Accounting for the motivations of the key actors involved in FOSS development, coordinators and developers, I ask how FOSS development is coordinated and how FOSS compete with proprietary software.

Coordinators guide the development of FOSS. They act as an interface between users and developers, leading the community of developers by prioritising features and bugs to work on (Klug et al., 2021; Lerner & Tirole, 2002).[5,6,7] Hence, higher levels of coordination can improve the quality of the FOSS as efforts by developers are more likely complementary. Conversely, less coordination leads to a lower quality FOSS as developments are more substitutable. I investigate three common motivations of coordinators (Lerner & Tirole,

[1] Other significant events include: the Heartbleed bug (John Biggs, 2014), malicious corruption by Marak Squires (Emma Roth, 2022), un-publication of leftpad (Chris Williams, 2016) to name a few.

[2] Firm adoption of FOSS is documented and encouraged by the Boston Consulting Group (2021). Denis Pushkarev (2023) and Marak Squires describe the funding challenges faced by developers (Emma Roth, 2022).

[3] André Staltz (2019) documents on the 58 most popular GitHub projects. The Apache Software Foundation (2023b) runs almost 400 FOSS projects on annual funding of USD 1 to 3 million.

[4] Chris Stokel-Walker (2014) and Denis Pushkarev (2023) document the problems they face in coordinating development.

[5] Raymond (1999) discusses the centrality of the coordinator's role and how he had to establish his competency to lead the FOSS Fetchmail, building relationships with users and developers.

[6] The Apache Software Foundation (2023a), The Document Foundation (2023), and The Linux Foundation (2023) are examples of FOSS foundations which have a hierarchical structure and are governed by volunteers.

[7] Brian Proffitt (2016) describes how governance of FOSS is decided in a meritocratic manner on Red Hat's Blog.

2002): (i) **self-interested Founders** who are unsatisfied with existing software, looking to maximise their own private utility (Klug et al., 2021); (ii) **volunteering Altruistic** from the community who care about the outcomes of others (Raymond, 1999)[6]; (iii) **profit-driven Managers** who may provide paid features and services adjacent to the FOSS, maximising their own profits (Andersen-Gott et al., 2012)[2,8]. To model this, I introduce coordinators who decide on the level of coordination the FOSS receives, focusing on self-interested Founders in the main model.

Developers choose when, what, and how they contribute to FOSS. Evidence shows developers contribute to FOSS because they are attracted to their high levels of customisability, only contributing to the development of FOSS features that directly benefit themselves (Franke & Von Hippel, 2003; Hars & Ou, 2002; Lerner & Tirole, 2002; Mustonen, 2003; Raymond, 1999; Shah, 2006). Selective contributions by developers mean the characteristics of a FOSS depend on who decides to contribute to it. However, the ability of a developer to alter the characteristics of a FOSS depends on the governance of the FOSS. I consider the software license adopted by the FOSS as a proxy of its governance structure.[6] More restrictive licenses, such as copyleft licenses, provide the coordinator with more control over the FOSS characteristics. This means the characteristics are less sensitive to the preferences of developers.[9] Conversely, permissive licenses, such as the MIT license, allow the characteristics of the FOSS to be more sensitive to developers' preferences.[10] I capture these effects in a model of product differentiation, allowing product characteristics to depend on which developers contribute to the FOSS and the software license adopted by the FOSS.

My main results show that by considering developers' preference, FOSS induce a more competitive environment than previously understood. I discuss how the Founder's trade-off leads to an equilibrium level of coordination which depends on the FOSS license. When comparing coordinators with different incentives, I show that self-interested Founders are more effective than Altruists or profit-driven Managers at coordinating FOSS, generating a higher total surplus.

Formally, I adopt a simple model of product differentiation à la Hotelling, modelling users with preferences over the characteristics of a software (d'Aspremont et al., 1979; Hotelling, 1929). I focus on the situation where a duopoly exists, and users choose between paying a fee and using a proprietary software or paying nothing to use the FOSS. I assume that

[8]Prominent examples of firms with a profit-driven motive that actively manage their own FOSS include Microsoft, IBM, and Google.

[9]A FOSS with a copyleft license would require any modifications or extensions to the software to be made a FOSS. Meaning proprietary software cannot reuse code from copyleft software.

[10]Permissive licenses impose few restrictions on how the FOSS' source code is used or distributed. The MIT license is listed among the most popular open-source licenses by Open Source Initiative (2022b) and GitHub (Balter (2015) blogs about this on GitHub.com).

users are also developers, and where they choose to use the FOSS, they contribute to its development.[11] They face some mismatch cost if their preferred characteristics differ from the characteristics of the chosen software.

A proprietary software is located at one end of the Hotelling line, and has a fixed quality. Its price is set by a profit-maximising firm.

The Founder's location is fixed at the opposite end of the Hotelling line. It decides on the level of coordination of the FOSS, facing the following trade-off. On the one hand, higher levels of coordination ensure that developers' efforts complement each other, leading to distinct contributions to the FOSS. Thus, a higher level of coordination improves the quality of the FOSS, which depends on the number of developers and the extent to which they complement each other. On the other hand, the characteristics of a FOSS are endogenous and depend on who develops it. Hence, the preferences of developers affect the location of the FOSS. Because the Founder's location is fixed at one end of the Hotelling line, every additional developer shifts the location of the FOSS away from the Founder's preference. This reduces the horizontal differentiation between the FOSS and proprietary firm, increasing the mismatch cost for the Founder.

First, by considering how developers' incentives may influence the location of a FOSS, I show that FOSS generate more competitive effects than previously understood. When more developers contribute to the FOSS, the characteristics of the FOSS become closer to the proprietary software, leading to fiercer horizontal competition. This effect is in addition to previously studied scenarios where more users can improve product quality (Fainmesser & Galeotti, 2020; Shy & Thisse, 1999). Because more developers benefit the FOSS' competitive position in two dimensions, proprietary firms have a stronger incentive to lower prices, increasing competition in the market. I discuss how government policies supporting the development of FOSS - even where the FOSS remains underutilised - can be beneficial as they drive competition (Madiega, 2020).[12]

Second, I show how the equilibrium level of coordination depends on the FOSS license. Restrictive licenses mean the FOSS location is not sensitive to the marginal developer. As developers are less able to influence the final software design, it is harder to encourage them to participate in the development of the FOSS. Fewer developers result in a FOSS of lower quality. Hence, when licenses are restrictive, Founders prefer distinct contributions to boost the quality of the FOSS. When a FOSS license is permissive, the location of the FOSS is more sensitive to the marginal developer, and additional developers increase the Founder's mismatch cost. As such, Founders may be deterred from managing the

[11] One way to think of this is that users are also contributing to the task of debugging the software (Raymond, 1999).

[12] See Sean Fleming (2021) and Benjamin Cedric Larsen (2022) for discussions on digital sovereignty.

FOSS, leading to a lower level of coordination. My model suggests that the prominence of permissive licenses, such as the MIT license, may explain why Founders decide to reduce, or even cease, their coordination efforts.[13]

Finally, I study the surplus generated by coordinators with different incentives. I show that self-interested Founders prefer a higher level of coordination and generate the highest total surplus. This result connects with anecdotal evidence on how many influential FOSS are coordinated by their Founders.[14] It also highlights the importance of directing FOSS funding to the correct coordinators, and questions how funding with stipulations of affiliation or requirements to adopt specific FOSS licenses can negatively affect the market.

This paper studies the coordination of Free and Open-Source Software. Previous theoretical work has focused on the eccentricities of developers (Athey & Ellison, 2014; Casadesus-Masanell & Ghemawat, 2006; Johnson, 2006; Johnson, 2002). My paper is perhaps most similar to Johnson (2002), as we consider how the community-driven nature of FOSS can influence the characteristics of the FOSS. I build upon this, taking into account the motives of coordinators, and study three types of coordinators: (i) the founder (Klug et al., 2021), (ii) members of the community who are recognised and appointed (Raymond, 1999), (iii) firms.[15] See Lerner and Tirole (2002) for a substantive survey.

I extend the workhorse model of product differentiation provided in Hotelling (1929). This model has been used to study vertical network effects, where more users can lead to better quality products (Fainmesser & Galeotti, 2020; Shy & Thisse, 1999). I introduce a horizontal network effect, one that affects location. Using this model, I study the situation where a product's compatibility with users can depend on who the users are. Doing so allows me to capture empirical evidence on the motives of developers who choose when, what, and how they contribute to FOSS (Fielding, 1999; Franke & Von Hippel, 2003; Hars & Ou, 2002; Shah, 2006).

By capturing the motives of developers, this paper relates to the literature on crowd-sourced innovation and competition. I show that a FOSS with a small market share can have large competitive effects, and this is driven by the proprietary firm preventing the development of the FOSS by lowering prices. My result is similar to Boone (2001). Through a different mechanism, Boone (2001) shows that market concentration can drive innovation by the incumbent, and does not imply a lack of competition.

[13] For example, the malicious corruption of color.js and faker.js by Marak Squires (Emma Roth, 2022) and the un-publication of leftpad by Chris Stokel-Walker (2014).

[14] Examples include core-js (Denis Pushkarev, 2023) and leftpad (Chris Williams, 2016). Also, see the xkcd comic on dependency (https://xkcd.com/2347/).

[15] Prominent examples of firms that actively manage their own FOSS include Microsoft, IBM, and Google.

There is little theoretical work discussing open-source licenses. Most papers avoid this discussion or discuss their results in the context of a specific license. Lerner and Tirole (2005b) and Gaudeul (2005) discuss simple models of license choice. Gaudeul (2008) assumes the selection of a given binary license, either permissive or restrictive, and discuss the implications of using either license. My model also takes the license choice as given but allows for a potential continuum of license designs. I show how the FOSS license can affect its existence and continued coordination.

In Section 3.2, I describe the setting in detail, discussing my main results in Section 3.3, and surplus implications in Section 3.4. In Section 3.5, I provide a series of extensions, and Section 3.6 concludes. I delegate proofs to Appendix C.1.

3.2 Model

I begin with the construction of a standard Hotelling model. I then introduce the concept of a FOSS, which accounts for network effects that simultaneously influence both the quality and location of the FOSS.

Users. There exists a unit mass of users. Users have heterogeneous preferences for product characteristics, x, and these are uniformly distributed, $x \sim U(0,1)$, along a Hotelling line. Users may choose between consuming the FOSS or an existing proprietary software, $i \in \{o, p\}$ respectively. Let $\bar{x}$ be the location of the user indifferent between consuming the FOSS or the proprietary software. The utility users receive from each software is evaluated as follows: $u_i = v_i - p_i - t|L_i - x|$. v_i represents the quality of the software, p_i the price, L_i the location of the software, and $t > 0$ the cost associated with a poor match ('mismatch cost').

Firm. The firm produces an existing proprietary software. This software is located at one end of the Hotelling line, fix $L_p = 0$. The firm is endowed with a software of quality v_p, which is large enough to ensure that the market is covered in equilibrium. The firm sets a price, p_p, that allows it to maximise its profits, $\pi_p = p_p \cdot D_p$, where D_p is demand for the firm's software. In this setting, $D_p = max\{\bar{x}, 1\}$, $\bar{x}$ being the indifferent user.

By fixing the proprietary software's location at 0, I study the situation where the firm is unable to change the characteristics of its software. Seemingly a strict assumption, this is rooted in the notion that the firm is an incumbent which has already decided what features its software is going to possess.[16]

Thus far, I have described a standard Hotelling model with a fixed location, where firms

[16]For example, it is public knowledge when Microsoft releases a feature, and what features are under development (Microsoft, 2023a, 2023b).

only make a pricing decision. To this model, I introduce a Free and Open-Source Software whose quality and location are influenced by the network of developers contributing to it.

Developers. There exists a unit mass of developers. Developers have heterogeneous preferences for what product characteristics they wish to develop, o, and these are uniformly distributed, $o \sim U(0,1)$, along the same Hotelling line as users. They receive a utility of $w_o = s_o - k|L_o - o|$. s_o captures the benefits to a developer and this depends on the number of users of the FOSS, $s_o = \beta \cdot D_o$, where D_o is the mass of users using the FOSS and β captures the prominence each additional user brings to the software. $k > 0$ captures the cost associated with developing features for a software that does not perfectly align with the developers preferences. Developers may either choose to contribute to the software, or some outside option.

Modelling developers this way captures that the popularity of a software allows the skills of individual developers to be recognised. This benefits the developers if they were motivated intrinsically to build something useful, or extrinsically, translating to job opportunities. Moreover, the mismatch cost here captures the preferences to work on features of a software that may interest or benefit developers directly, and much less on other features.

Free and Open-Source Software. To capture the characteristics of a FOSS, I introduce a model of endogenous (vertical) quality and (horizontal) location, both of which depend on the number of users. I dub these effects 'quality' and 'location' network effects respectively.[17] By definition, a FOSS is free to use, and hence p_o, the price of the FOSS, is 0.[18]

The quality of the FOSS depends on the number of developers, $v_o = v_c + \gamma \cdot D_o$. Here, $\gamma \in [0, 1]$ captures the probability that each contribution is distinct. Because developers are free to select the type of contributions they make, there is always a possibility that multiple contributions overlap.[19] v_c is the quality of the FOSS when the coordinator is its sole developer. This captures that some FOSS may exist without additional developers, and thus do not benefit from network effect. I normalise $v_c = 0$ to reduce parameters.[20]

[17] Network effects arise when each user's utility is affected by the number of other users. My main model captures 'direct' network effects. In the general setting where users and developers are mutually exclusive, network effects are 'indirect'. See Belleflamme and Peitz (2015) for details.

[18] FOSS are free to use because anyone is able to download it source code to build and run the software.

[19] And $(1 - \gamma)$ the probability that the contributions overlap.

[20] More generally, the reader may choose to interpret $v_p > 0$ as the difference between the quality of the proprietary software and the FOSS without network effects. Notice that $v_p \leq 0$ implies that the proprietary firm is inactive. See Etzion and Pang (2014) for details.

The location of the FOSS depends on the preferences of its developers, $L_o = 1 - l \cdot D_o$, $l \in (0, 1)$ is a parameter that regulates the FOSS' location, and $D_o = 1 - \bar{x}$ is the demand the FOSS faces, $\bar{x}$ being the marginal developer. In particular, $l \to 0$ suggests that the location of the software is insensitive to the marginal developer, and $l \to 1$ implies that the location is very sensitive to the marginal developer. Mathematically, this implies that when more developers contribute to the development of the FOSS, the location of the FOSS shifts closer to the proprietary software. I interpret this as the permissiveness of the license adopted by the FOSS, and suppose that l is fixed ex-ante.[21]

A permissive license guarantees the right to use modify, and redistribute the underlying source code. More restrictive licenses may grant limited rights, disallowing the software to be used in specific situations, for example for profit or in programs that may harm others. The most restrictive of licenses would be a closed source software where the underlying source code is public but there is no right of use, or a proprietary license where no code is made public. Hence, permissive licenses encourages users to suggest modifications to the coordinator, and make a coordinator more likely to accept changes.[22] This way, I interpret l as the permissiveness of a license, which relates to the ability of developers to influence the characteristics and location of the FOSS.

Coordinator. Free and Open-Source Software that are actively developed are always coordinated by someone. The coordinator chooses the probability that each contribution is distinct, $\gamma \in [0, 1]$. This proxies the coordination level of the FOSS, such that a larger γ reflects how the efforts of developers complement each other. In the main model, I discuss a FOSS which is coordinated by its **Founder**. Founders are motivated to develop a FOSS because they face a huge mismatch cost from using the proprietary software. Hence, I assume that the Founder is located at 1. Moreover, Founders develop FOSS to fix problems of their own. I model Founders as self-interested, maximising their utility, $\pi_o = \gamma \cdot D_o - t(l \cdot D_o)$. This utility function is identical to the utility of function of users, but is specific to the user located at 1.

For simplicity, I shall assume that all users are developers, collapsing both groups into 'users', and setting $\beta = \gamma$ and $k = t$. This reduces the number of actors, and variables, in the main model and allows me to focus on the decision of the coordinator.

The sequence of events is as follows. The coordinator chooses the probability that each contribution is distinct, γ, firms then set the price, p_p, and users finally decide between the proprietary software or the FOSS. This captures the notion that it takes time to

[21] The reader may also choose to interpret l as the hierarchical structure of the FOSS, or the inverse of the degree of control that coordinators have over the direction of the FOSS.

[22] When the FOSS has a permissive license and the coordinator chooses not to accept a change, other developers can simply 'fork' or spin-off the project. Denis Pushkarev (2023), the founder of core-js, a standard library for JavaScript, describes this in his regular update to the community.

develop a software and the fluidity of software prices.

This concludes the description of the model environment.

Restrictions. I look for a subgame perfect equilibrium. To characterise this equilibrium, I introduce the following restrictions as equilibrium selection assumptions. These restrictions mirror the 'standard' restrictions in Hotelling models.

Restriction 1. *The market is covered.*

Restriction 2. *Both the firm and the FOSS are active in the market.*

Combined, these restrictions allow me to focus on the scenario where the development of a FOSS creates competitive pressures for a proprietary firm. This is the scenario of interest as the evidence suggests that many FOSS are developed as a response to the needs of users which cannot be met by existing software - meaning that in many cases, there exists a proprietary software that serves as a less than perfect substitute. This allows me to study a rich environment, and explain many realistic competitive environments.

3.3 Coordinating Free and Open-Source Software

Before considering the decisions of coordinators, I determine the choices made by users and firms. Their decisions are the best response functions to the decision of other users and the coordinator, and do not depend on the motives of the coordinator. While the quality of the proprietary software and the mismatch cost affect the equilibrium, these are economically understood. Instead, I focus my discussion on the novel effects that arise from the coordination of FOSS and license of FOSS.

First, I determine the users' choice of software, the FOSS or the proprietary software. In this setting, there exists a user who is indifferent between the FOSS and the proprietary software,

$$\bar{x} = \frac{v_p - p_p + p_o - \gamma + t(1-l)}{t(2-l) - \gamma} \tag{3.1}$$

Intuitively, higher levels of coordination increases the marginal network effect ($\gamma \uparrow$), improving the quality of the FOSS. Assuming prices to be 0, and hence ignoring the strategic pricing decision of firms, this makes the FOSS more attractive to users.

To understand the role FOSS licenses play, consider that more permissive licenses ($l \uparrow$) give more rights and freedom to other developers to modify and use the FOSS code. Conversely, restrictive licenses prevent the use of the underlying code, the most extreme example being closed proprietary software, where only the binary file is shared, removing the ability of others to view (and hence use) the underlying code. Therefore permissive software licenses mean the characteristics of the FOSS are more sensitive to the marginal user. When a FOSS license is permissive, the marginal user faces a lower mismatch cost when choosing the FOSS. Hence, by ignoring the strategic pricing decisions of firms and assuming prices to be 0, permissive licenses also increase the FOSS' appeal to users.

Second, I consider the decision of the proprietary firm. The firm chooses to lower prices when the FOSS has a higher level of coordination or a more permissive license. First, a higher level of coordination increases the quality of the FOSS, which leads to a virtuous cycle of network effects where the higher quality FOSS attracts even more users. The firm hence faces stronger quality-based (vertical) competition. Second, more permissive licenses mean the characteristics of the FOSS are more sensitive to the preferences of the marginal user. This means as more users use the FOSS, the location of the FOSS becomes closer to the firm. This decreases horizontal differentiation, which increases characteristics-based (horizontal) competition.

Lemma 1. *The equilibrium price set by the proprietary firm increases in its quality and the price of the FOSS, and decreases in the level of coordination of the FOSS and permissiveness of the FOSS license.*

$$p_p^* = \frac{v_p + p_o - \gamma + t(1-l)}{2} \tag{3.2}$$

See C.1.1 for proof.

Because FOSS are free to use, I impose that $p_o = 0$. In equilibrium, when accounting for the price of the proprietary software, (3.1) becomes

$$\bar{x} = \frac{v_p - \gamma + t(1-l)}{2(t(2-l) - \gamma)}$$

and the first derivative with respect to γ and l are positive if and only if $v_p > t$.[23] Therefore, a proprietary firm with a sufficiently high quality software is able to capture a larger share of the market by setting lower prices even when FOSS have a higher level of coordination or more permissive licenses.

Interestingly this means that setting a lower price allows the proprietary firm to recapture demand and counter competitive forces generated by the FOSS in two ways. First, by decreasing the size of the FOSS network, the quality of the FOSS decreases, reducing vertical competition. Second, lower prices allow the proprietary firm to recapture users with less extreme preferences. Because only users with more extreme preferences (closer

[23]Consider the incentives of a FOSS coordinator. In the case of Founders or Altruists, they are providing a service for free, and many have argued that no money should change hands. See Footnote 12 and 18. When considering a profit-driven Manager, negative pricing cannot be profitable.

to 1) remain in the FOSS network, this increases the differentiation between the two software, reducing horizontal competition.

The dual incentive for proprietary firms to compete on prices suggests that the existence of FOSS is able to generate a larger competitive effect than previously understood. Because standard models of competition do not account for how FOSS' characteristics depend on developers' preferences, they do not account for the additional horizontal competitive effect. Lemma 1 and $\bar{x}$ together show that proprietary firms have an additional incentive to lower prices when licenses are more permissive. Hence, by accounting for this horizontal competitive effect, I show that more permissive FOSS licenses can reinforce competition.

Despite gaining a larger share of the market, the proprietary software is less profitable when competing against a FOSS with a higher level of coordination or a more permissive license. The negative effect on profit from the fall in prices dominates any gains in profit from a larger market share.

Corollary 1. *The proprietary firm's profits decreases in the level of coordination of the FOSS and permissiveness of the FOSS license.*

$$\pi_p = \frac{[v_p + p_o - \gamma + t(1-l)]^2}{4(t(2-l) - \gamma)}$$

See C.1.2 for proof.

Corollary 1 and $\bar{x}$ hence suggest that even niche FOSS that serve a small user base can contribute to a more competitive market. This result has similar implications as Boone (2001), suggesting that it is not always appropriate to use market concentration as a measure of its competitiveness.

Having discussed the decisions of users and the proprietary firm, I now turn my attention to the self-interested Founder.

3.3.1 Self-Interested Founder

As a baseline, I consider the incentives of a Founder of FOSS. As suggested by many, Founders of FOSS are usually motivated by some degree of self-interest, looking to resolve a problems that they face (Franke & Von Hippel, 2003; Hars & Ou, 2002; Klug et al., 2021; Lerner & Tirole, 2002). I take an extreme view where Founders are completely self-interested, seeking to maximise their own utility.

Recall that the Founder's utility is

$$\pi_o = \gamma \cdot D_o - t(l \cdot D_o)$$
$$= \frac{(\gamma - tl)(t(3-l) - v_p - \gamma)}{2(t(2-l) - \gamma)}$$

and the FOSS is free to use, $p_o = 0$. Founders decide on the probability that contributions are distinct, γ, which proxies the level of coordination of the FOSS. This provides the following first-order condition:

$$\frac{\partial \pi_o}{\partial \gamma} = \underbrace{\frac{(\gamma - tl)(t(3-l) - v_p - \gamma)}{2(t(2-l) - \gamma)^2} + \frac{t(3-l) - v_p - \gamma}{2(t(2-l) - \gamma)}}_{>0} \underbrace{- \frac{\gamma - tl}{2(t(2-l) - \gamma)}}_{<0}.$$

The first-order condition shows the trade-off of the Founder. On the one hand, the Founder has the incentive to grow the network of users. By leveraging the contributions of additional users, the Founder can improve the quality of the FOSS, and by extension, their own utility. Hence, the Founder prefers that contributions are somewhat distinct, $\gamma > 0$. However, as more users join the network, the characteristics of the FOSS change. This change results in a software that is less compatible with the Founder's needs and preferences, shifting the location of the software away from the Founder. Under these circumstances, the Founder has less incentive to ensure that contributions to the development of the FOSS are distinct, as doing so incurs a higher mismatch cost, harming itself.

Using this trade-off, I find there exists a unique subgame perfect Nash equilibrium that satisfies Restrictions 1 and 2.

Proposition 1. *Founders are active only if the proprietary software is of a sufficiently low quality ($v_p < t(3 - 2l)$). Their level of coordination is weakly decreasing in the permissiveness of FOSS license ($l \uparrow$), and the quality of the proprietary software. And strictly decreasing if contributions are not perfectly distinct ($\gamma^* < 1$).*

$$\gamma^* = min\{t(2-l) - \sqrt{2t(v_p - t)(1-l)}, 1\}$$

with $\gamma^* = 1 \iff 3t + \frac{tl^2}{2(1-l)} + \frac{1}{2t(1-l)} - \frac{2-l}{1-l} \geq v_p$ *and* $t > \frac{1}{2-l}$.

See C.1.3 for proof.

Proposition 1 reiterates how more permissive licenses can make the Founder worse off. Because more permissive licenses allow other developers to have a greater influence over the characteristics of the FOSS, the location of the software is more sensitive to the preferences of the marginal user. Hence, more permissive licenses increase the mismatch cost for the Founder. Founders of permissive FOSS hence choose lower levels of coordination, allowing overlapping contributions to the FOSS. This serves to increase the mismatch cost of the marginal user, discouraging changes to the characteristics of the FOSS, and reducing the Founder's own mismatch cost. However, doing so has the additional effect of reducing the FOSS quality. The lower quality dominates the recovered mismatch cost, making the Founder worse off.

One way of understanding such licenses is in the ability to control the future of the underlying code. This is not dissimilar to our understanding of Intellectual Property Rights (IPR), where a larger l captures limited IPRs for the rights holder. This way, Proposition 1 connects well to prior research on how IPRs can lead to more investment and innovation (Besley, 1995; Bloxam, 1957; Chen & Puttitanun, 2005; Moser, 2005; Qian, 2007).

Software licenses are usually selected at its onset (Vendome et al., 2015b).[24] Because of the collaborative nature of FOSS it is often not possible to modify a license once it has been adopted (Gamblin, 2021).[25] This can lead to a license "lock-in" effect. Anecdotal evidence suggests that Founders are encouraged to, and often select permissive FOSS licenses - such as the MIT license.[10] However, this has been the subject of a long-standing debate.

Supporters of permissive licenses argue that permissive licenses eases the rules of participation in the FOSS community: (i) coordinators minimise their legal liability; (ii) users face no legal risk for using the FOSS; (iii) developers find it easier to contribute to the FOSS (Gamblin, 2021).[26] Permissive licenses thus encourage more users to join the FOSS community, allowing the Founder to depend on network effects to grow the FOSS (Gamblin, 2021). They also argue that permissive licenses are in the spirit of FOSS, allowing anyone (and everyone) access (and permission) to use and modify the source code.

On the other hand, detractors claim that permissive licenses go against the spirit of liberty and digital sovereignty - the ideological view that the inability to pay for a proprietary software should not exclude individuals from participating in the digital age.[12,26] They worry that overly permissive licenses allow corporations to take advantage of FOSS developers, repacking FOSS into proprietary software and profiteering from a public good.[27] Still others warn that permissive licenses make it difficult for developers and coordinators to monetise their work, which can lead to unsustainable software development.

[24] Vendome et al. (2015a), Vendome et al. (2015b) also comment that where a license change occurs, a more permissive license is subsequently adopted and this is often attributed to developers or funding requests, or to allow for commercialisation.

[25] Stackexchange (2015) provides a detailed discussion regarding the problems of subsequently changing FOSS licenses.

[26] Nicolas Suzor (2013) describes on the website opensource.com.

[27] Paint.NET was released in 2004 with the MIT license, developer Rick Brewster decided to adopt a stricter license to prevent others from profiting from the software (Brewster (2009) explains in his community update).

Corollary 2 contributes to this discussion by showing how permissive licenses can lead to the underdevelopment of FOSS. Licenses need to provide Founders with sufficient control over the use and development of the FOSS. This level of control depends on v_p and t. The differences in market conditions across each software market can thus explain the variety of open-source licenses in active use, and the continued challenges in designing and choosing boilerplate open-source licenses (Lerner & Tirole, 2002; Subramaniam et al., 2009).[28,29]

Corollary 2. *Founders only coordinate FOSS when the FOSS license is not too permissive. The upper bound on the permissiveness of the FOSS license decreases in the quality of the proprietary software, and increases in the strength of user preferences* ($t \uparrow$).

Mathematically, $\gamma^* > 0 \iff l < \bar{l} = \frac{3t - v_p}{2t}$.

See C.1.4 for proof.

The intuition follows: when users' preferences are weaker (small t), they only face a small mismatch cost when using products with characteristics further from their preferences. This means the FOSS is less able to use location to improve the utility of users. As a result, the FOSS' advantage from location network effects diminishes. Hence, Founders are only active when licenses are more restrictive.

While permissive licenses may help avoid legal issues, they diminish the willingness of Founders to manage and continue the development of a FOSS. This is not easy to observe in the data, but many prominent anecdotes where Founders stopped developing FOSS exist. For instance, leftpad, colors.js and faker.js.[1] Others, such as Paint.NET, have adopted more restrictive licenses when market conditions evolved.[27] Corollary 2 supports those who believe permissive software licenses may lead to unsustainable software development. My result questions the default recommendation of permissive licenses, such as the MIT license. Instead, FOSS licenses should be tailored for each software, not unlike IPRs, and no license should be promoted generically above another without consideration of the specific software market.

Corollary 3. *FOSS serve a smaller share of a market when licenses are more permissive.*

Proof.
In equilibrium, (3.1) becomes

[28] Almeida et al. (2017) shows developers have a hard time comparing licenses, turning to tools such as choosealicense.com and license.md.
[29] See also Janis Lesinski (2020) for problems with the MIT license.

$$\bar{x} = \frac{v_p - \gamma^* + t(1-l)}{2(t(2-l) - \gamma^*)}$$
$$\frac{\partial \bar{x}}{\partial l} = \frac{t(v_p - t)^2}{2(2t(v_p - t)(1-l))^{\frac{3}{2}}}$$

Notice that $v_p > t$ for a real solution to γ^*. Therefore $\frac{\partial \bar{x}}{\partial l} > 0$. □

The negative impact that permissive licenses have on the level of coordination (Proposition 1) and the price of the proprietary software (Lemma 1) can explain why many FOSS serve only niche markets and never become mainstream (Corollary 3) (Lakhani & Hippel, 2004; Lerner & Tirole, 2005a; Robles et al., 2019). First, it is not beneficial for a Founder to over-invest in the development of the FOSS. Despite the virtuous cycle of network effects, where quality leads to more users, and more users lead to a software with mass market appeal, the Founder suffers from a growing mismatch of preferences. Second, firms become more price competitive, reducing the incentive for users to use the FOSS. My model also provides an explanation for the 'pet project' mentality found in anecdotal evidence.[30] That is to say, self-interested Founders may limit the quality and scope of the software despite the potential societal benefits, leading to the many FOSS which serves only niche markets or just a few individuals.

To summarise, the main model studies how the design of FOSS licenses influence the development and competitiveness of the FOSS. I show that self-interested Founders only manage a FOSS if they retain sufficient control of the software. When a FOSS license is more restrictive, the Founder prefers that each contribution to the FOSS is distinct. However, if a more permissive license is adopted, Founders may prefer overlapping contributions, causing the FOSS to be of a lower quality, and attract fewer users. When licenses are permissive, despite being of a lower quality, the FOSS is still able to create a more competitive environment.

3.3.2 Altruist

The literature suggests that social preferences have been noted as a motivation for contributing to and managing a FOSS. Hars and Ou (2002) suggest that developers are often motivated by building a sense of community. Whereas Lerner and Tirole (2002) state that altruism is one motivation often reported. Moreover, many large FOSS communities are user managed: through a meritocratic process, users are recognised and rise in the hierarchy (Raymond, 1999).[6] Arguably, this implies that coordinators (as an individual or a group) are motivated by some sense of altruism. Altruistic coordinators are also more likely to be open to changing the characteristics of a FOSS to fit the needs of the users of the software. To model such an Altruist, I consider that a coordinator might be

[30] Adem Ait Fonollà (2022) documents that most projects on GitHub are abandoned within 5 years.

motivated by a sense of community, working towards maximising the value of the FOSS for its users.

This model follows directly from the base setting, with the following modification: instead of a self-interested Founder, the coordinator is an Altruist and seeks to maximise the value of the FOSS network, and the FOSS is free to use, $p_o = 0$. This means they are concerned with generating the largest possible benefit for their members,

$$\begin{aligned}\pi_o^A &= \int_{\bar{x}}^1 u_o(x)dx \\ &= \int_{\bar{x}}^1 v_o - t|L_o - x|dx \\ &= \int_{\bar{x}}^1 \gamma(1-\bar{x}) - t|1 - l(1-\bar{x}) - x|dx.\end{aligned}$$

Taking into account the equilibrium decisions of users, (3.1), and the proprietary firm, (3.2),

$$\pi_o^A = \frac{(2\gamma - t(1-2l+2l^2))(v_p + \gamma - t(3-l))^2}{8(t(2-l)-\gamma)^2}.$$

Where coordinators are Altruists, I find that there exists a unique subgame perfect Nash equilibrium that satisfies Restrictions 1 and 2.

Proposition 2. *An Altruist always selects a positive level of coordination. Their level of coordination is weakly decreasing in the quality of the proprietary software, and strictly decreasing if contributions are not perfectly distinct* ($\gamma^A < 1$).

$$\gamma^A = min\left\{\frac{t(3-2l) + v_p - \sqrt{(v_p - t)(11t - 8tl^2 + v_p)}}{2}, 1\right\}$$

with $\gamma^A = 1 \iff v_p \in (t, \frac{1+t^2(5-3l-l^2)-t(3-2l)}{1+t(1+l-2l^2)}]$*, and* $t > \frac{1}{2-l}$.

See C.1.5 for proof.

To understand why Altruists are always active, consider that Altruists are motivated by the surplus of users of the FOSS, and only receive a positive utility when such a software exists. Hence, as long as there exists some user located arbitrarily close to 1 such that their mismatch cost is negligible, then the Altruist will coordinate the development of the FOSS. This result may explain why FOSS foundations, such as The Apache Software Foundation, support projects covering a diverse range of topics.[31]

[31] The Apache Software Foundation (2022) describes the diversity of topics covered by Apache supported projects in 2021.

Corollary 4. *Founders prefer contributions which are more distinct than what Altruists prefer. Hence, $\gamma^* > \gamma^A$ when $\gamma^* \neq 1$ and $\gamma^A \neq 1$.*

See C.1.6 for proof.

To understand Corollary 4, notice that the key difference between FOSS coordinated by Founders and Altruists is the objective function of the coordinator. Founders are concerned with their own mismatch cost, whereas Altruists consider the mismatch cost of all FOSS users. This means Altruists are less willing to select a higher level of coordination, because this increases the costs of many others, while Founders only consider their personal utility.

For intuition, suppose that a coordinator selects a lower level of coordination ($\gamma \downarrow$). This decreases the quality of the FOSS. By extension, reducing the number of FOSS users, and shifting the location of the FOSS closer to 1. As a result, the mismatch cost of the extreme user, Founders located at 1, decreases. Additionally, the mismatch cost of the marginal user decreases. Because the Altruist concerns itself with all users of the FOSS, they consider this additional effect. Thus, Altruists have a larger benefit from selecting a lower γ than Founders.

3.3.3 Profit-Driven Manager

Many prominent coordinators - such as Google and IBM - are driven by a profit motive. Lerner and Tirole (2002) and Andersen-Gott et al. (2012) also discuss how firms may adopt a code-release strategy, allowing the community to further develop the software, while providing features and services adjacent to the FOSS for profit. This setting is analogous to one where some of the software's code is open-source and other (non-employee) developers may create plug-ins or extensions to the software. Although a firm allows this, it may still choose to moderate the type of plug-ins and extensions readily available. The firm then earns profit from the sale of data or other auxiliary services. A prime example of this is Google's Chrome web browser.

To model profit-driven Managers, I deviate from the main model and allow the Manager to set prices and the level of coordination with profit in mind. The profit-driven Manager maximises the following $\pi_o^M = p_o \cdot D_o$, where $D_o = 1 - \bar{x}$ is the demand of the FOSS, and p_o the price set by the Manager. The timing of the game follows: (i) the Manager selects the probability that contributions are distinct, γ; (ii) the Manager and the proprietary firm simultaneously select prices, p_o and p_p respectively; (iii) users decide between the FOSS or the proprietary software.

Taking into account the equilibrium decisions of users, (3.1),

$$\pi_o^M = \frac{p_o(p_p + t - v_p - p_o)}{t(2-l) - \gamma}$$

and it is immediate that profits are concave in price, and the optimal pricing strategy is $p_o^M = \frac{p_p+t-v_p}{2}$. Accounting for the pricing strategy of the proprietary firm, (3.2),

$$p_o^M = \frac{t(3-l) - \gamma - v_p}{3} \tag{3.3}$$

$$p_p^M = \frac{v_p - 2\gamma + t(3-2l)}{3} \tag{3.4}$$

Higher levels of coordination should lead to higher quality products. Despite this, managers choose to lower prices when selecting a higher level of coordination. This is because lower prices have the additional effect of attracting more users. Hence, lower prices complement the higher level of coordination, building upon the virtuous cycle of network effects, to further improve the FOSS' quality. Together, equations (3.3) and (3.4) show how higher levels of coordination can lead to a more competitive environment, as both firms engage in fiercer price competition, lowering prices.

Taking into account equilibrium prices, (3.1) becomes

$$\bar{x} = \frac{t(3-2l) + v_p - 2\gamma}{3(t(2-l) - \gamma)}$$

$$\frac{\partial \bar{x}}{\partial \gamma} = \frac{v_p - t}{3(t(2-l) - \gamma)^2} \tag{3.5}$$

$$\frac{\partial \bar{x}}{\partial l} = \frac{t(v_p - t)}{3(t(2-l) - \gamma)^2} \tag{3.6}$$

When the quality of the proprietary software is sufficiently high ($v_p > t$), the proprietary firm has a dual incentive to compete on prices. First, setting a lower price allows it to capture a larger share of the market. Second, increasing the market share of the proprietary software lowers the quality of the FOSS as less developers contribute to it. Hence, in response to a larger γ, the proprietary firm decreases its prices to capture a larger share of the market.

However, when the quality of the proprietary software is low ($t > v_p$), the mismatch cost faced by users is too large; this diminishes the proprietary firm's ability to use prices to gain market share. Hence, when the FOSS' level of coordination is higher, the proprietary firm is unable to retain demand and loses users to the FOSS instead.

Proposition 3. *Profit-driven Managers are active only if the FOSS has a sufficiently permissive license ($l \geq \frac{3t-2}{2t}$). When they are active, profit-driven Managers choose between binary levels of coordination. They prefer distinct contributions ($\gamma^M = 1$) when the proprietary firm is of a sufficiently low quality ($v_p < t$). When the proprietary firm has a high quality ($v_p \geq t$), Managers choose the lowest level of coordination ($\gamma^M =$*

$max\{t, t(1.5 - l)\})$ for which users continue to use the FOSS.

See C.1.7 for proof.

Proposition 3 shows how Managers competing with higher quality proprietary software can take advantage of the permissiveness of FOSS to lower their level of coordination. Notice that when the FOSS license is restrictive, then $\gamma^M = t(1-5-l)$, but as l increases, or FOSS adopt more permissive licenses, γ^M becomes smaller until $\gamma^M = t$. This result suggests that Managers adopting an open-source approach to software development can choose lower levels of coordination, which may make it easier for them to coordinate a FOSS. Moreover, despite a lower level of coordination, the coordination of the FOSS is still profitable.

Proposition 3 also shows how Managers of FOSS prefer distinct contributions when they face low quality proprietary competitors. This is driven by the gains in revenue from a larger user base (3.5), which dominates any loss that arises from a lower per-unit price (3.4).

These results connect with recommendations that firms should adopt open-source business models.[32] Anecdotal evidence shows how Microsoft, which once called open-source 'a cancer', has recanted on this position, and is now the largest contributor to the development of FOSS.[33]

To further illustrate Proposition 3, I use the case study of competition between internet browsers.

Browser War 2.0 In 2008, the leading internet browser was Microsoft's Internet Explorer (IE). This was thought to be a frustrating and outdated browser ($v_p < t$).[34] Google developed an open-source browser, Chromium, upon which Chrome is built.[35] Additionally, they chose to adopt a code-release strategy, allowing them to benefit from the contributions of external developers, and better understand what features users desired. My model suggests that a code-release strategy can be justified by a specific market environment: where the proprietary software has a low quality (small v_p) and the mismatch cost for users is very high (large t). When this is the case, code-releasing allows the FOSS to capture a large share of the market (3.6). Therefore, my model captures closely how Chrome managed to become a market leader in under 3 years.[36]

[32] Consulting companies such as Boston Consulting Group (2021) promote that firms rely on the open-source community for software development.

[33] Thomas Greene (2001) reports on Microsoft's stance on open-source as a cancer, and Tom Warren (2020) reports of their retraction on this position.

[34] This is documented by Luke Little (2021) and Tom Warren (2018) on Andriod Authority and The Verge respectively.

[35] Releasing Chromium's code corresponds to a permissive licensing strategy, $l \geq \frac{1}{2}$.

[36] As tracked by Statcounter (2022).

3.4 Surplus Analysis

In Corollary 3, I show that permissive licenses can lead to a FOSS that serves a smaller market, and in Proposition 3, I discuss how permissive licenses can encourage firms to adopt code-release strategies. Therefore, it is unclear if permissive licenses are beneficial to the development of FOSS. Here, I investigate how the design of FOSS licenses can influence consumer surplus.

I begin by discussing the self-interested Founder, followed by the Altruist, and the profit-driven Manager, focusing on the situation where they do not prefer perfectly distinct contributions. Finally, I study how total surplus changes with the level of coordination of the FOSS.

3.4.1 Self-Interested Founder

Corollary 5. *When $\gamma^* \neq 1$ and FOSS coordinated by Founders have more permissive licenses, the consumer surplus generated by the FOSS is lower.*

See C.1.8 for proof.

This result arises for two reasons. First, the decrease in consumer surplus is directly driven by the fall in the number of FOSS users (Corollary 3). Second, because there are less FOSS users, this leads to an indirect decrease in the quality of the FOSS.

Since a more permissive license means that more users use the proprietary software, we naturally expect that the consumer surplus contributed by the consumption of the proprietary software increases in l. Moreover, because firms lower prices when facing intensified location based competition (Lemma 1), lower prices also increases consumer surplus from the proprietary software.

Corollary 6. *When $\gamma^* \neq 1$ and FOSS coordinated by Founders have more permissive licenses, the consumer surplus generated by the proprietary software is higher.*

See C.1.9 for proof.

3.4.2 Altruist

Shifting attention to Altruistic coordinators, I show that consumer surplus of FOSS users is weakly decreasing when the FOSS adopts a more permissive license. Moreover, users of the proprietary software always benefit when the FOSS license is more permissive.

Corollary 7. *When $\gamma^A \neq 1$ and FOSS coordinated by Altruists have more permissive li-*

censes, the consumer surplus of FOSS users (weakly) decrease, while the consumer surplus of proprietary software users increase.

See C.1.10 for proof.

Using Corollary 8, I show that Founders generate more overall consumer surplus than Altruists. This is because Founders prefer a higher level of coordination than Altruists (Corollary 4). In doing so, Founders generate larger competitive effects, causing the proprietary firm to lower prices and extract less surplus from users. Hence, even though the FOSS consumer surplus may fall, the overall consumer surplus is higher when Founders coordinate FOSS.

Moreover, Founders generate a higher total surplus than Altruists. This is because the FOSS benefits from network effects, which expands the total surplus available to both users and firms. Such a comparison between the Founder and Altruist is possible because they face identical market conditions, and the equilibrium only differs in their level of coordination.

Corollary 8. *In equilibrium, when Altruists prefer contributions that are not distinct, both consumer surplus and total surplus is increasing in the level of coordination of the FOSS.*

See C.1.11 for proof.

Although this result suggests that Altruistic coordinators generate less surplus than Founders, they still play an important role in the FOSS community. Since Founders may decide to stop coordinating a FOSS when facing unfavourable market conditions, they may choose to pass the role of coordinator on. Altruists are always willing to coordinate FOSS, and ensure that these software continue to be supported.[37] Hence, despite generating a lower level of surplus, this is preferable to a situation where no FOSS exists, and a proprietary firm forms a monopoly.

3.4.3 Profit-Driven Manager

When studying the surplus of FOSS coordinated by profit-driven Managers, there is an additional complexity stemming from the pricing strategy employed by the Manager. The combination of strategic decisions by the Manager means that the changes in surplus generated by each software is non-monotone.

Corollary 9. *When Managers coordinate FOSS, more permissive licenses weakly decrease*

[37] This argument is made by Denis Pushkarev (2023), and exemplified by how members of the community took over the management of Fetchmail, colors.js, and faker.js (Emma Roth, 2022; Raymond, 1999).

profits of both the FOSS and proprietary firm when $v_p > t$ and Managers do not prefer distinct contributions. If licenses are sufficiently permissive, $l \geq 0.5$, profits of both the FOSS and proprietary firm strictly decrease.

Moreover, when $v_p < t$ and Managers prefer distinct contributions, the profits of the FOSS increases when licenses are more permissive.

See C.1.12 for proof.

Corollary 9 reiterates how more permissive FOSS licenses can generate location based competition. More importantly, the second half of the Corollary points out how potential conflicts of interest can lead to the promotion of permissive licenses in the Open-Source community. When profit-driven Managers can influence the definition of "Open-Source", they may prefer that this definition applies only to the most permissive of licenses when the proprietary software has a low quality.

Additionally, Corollary 10 highlights how more permissive licenses always harms consumers when Managers coordinate the FOSS. Together, these results question the motives of organisations such as the Open Source Initiative (2022a) which provides an 'independent' stamp of approval on what licenses are deemed Open Source. Because organisations that participate in FOSS development are able to directly propose members to its governing body, the incentives of such members may be skewed to benefit firms that participate in FOSS development.

Corollary 10. *When Managers coordinate FOSS, more permissive licenses unambiguously reduce consumer surplus.*

See C.1.13 for proof.

Finally, I show that whenever it is possible for the Manager to costlessly improve their level of coordination, this improves total surplus. This is because a higher level of coordination can expand the quality in the market, improving the total surplus available to both firms and consumers. However, when the proprietary software has a sufficiently high quality ($v_p > t$), Managers prefer lower levels of coordination. This allows them to occupy a niche segment of the market, reducing competition with the proprietary firm, setting higher prices, and maximising profits. This is not unlike a traditional Hotelling result of maximal differentiation between firms.

Corollary 11. *In equilibrium, when Managers prefer contributions that are not distinct, total surplus is increasing in the level of coordination of the FOSS.*

See C.1.14 for proof.

Together, Corollary 8 and 11 show that it is always possible to improve total surplus from the Altruist and Manager equilibrium. This suggests that when they coordinate FOSS, they fall short of the socially optimal level of coordination.

I have shown that permissive licenses may benefit or harm consumers, depending on who coordinates the FOSS. When Founders coordinate the FOSS, it is ambiguous if consumers benefit from more permissive licenses. However, Managers coordinate FOSS, consumer surplus unambiguously decreases. Hence, although permissive licenses can create a more competitive environment, decreasing profits of the firm, they may not benefit consumers.

3.5 Extensions

I provide a series of extensions, considering the situations where: (i) Founders decide on a license choice instead of level of coordination; (ii) developers have heterogeneous skill levels; (iii) users and developers are mutually exclusive groups.

3.5.1 Founder's License Choice

In my setting, I have assumed that coordinators are unable to influence the license adopted by a FOSS, taking it as a given. I consider this a realistic setting because many developers are told to adopt 'standard' FOSS licenses such as the MIT license. Unlike coordination effort, once a license has been selected, it is much more sticky. Hence, as the market environment changes, or as coordination changes hands, licenses are already entrenched, and it is more interesting to focus on how the level of coordination evolves. However, it is true that Founder's do make an initial license choice at the beginning of a FOSS. To capture the initial license decision, I consider the situation where a Founder dos not select the level of coordination, instead selecting the license adopted by the FOSS.

Proposition 4. *Founders choice of FOSS license is binary. They prefer the most restrictive license* ($l = 0$) *whenever the proprietary software has a sufficiently low quality and the mismatch cost is sufficiently small.*[38] *In all other cases, Founders prefer the most permissive license* ($l = 1$).

See C.1.15 for proof.

Proposition 4 shows that self-interested Founders prefer permissive licenses under most conditions. They only select restrictive licenses when the quality of the proprietary software is sufficiently small and the marginal gains in quality due to network effect are

[38]Mathematically, $v_p < \frac{t^2(6-4l+l^2)+\gamma^2-2t\gamma(3-l)}{2(t-\gamma)}$ and $t < \gamma$.

larger than the mismatch cost.

This result has two interesting features. First, notice that the degree of permissiveness is affected by $\frac{\gamma}{t}$. This suggests that when the marginal vertical network effect increases, coordinators are willing to select a less permissive license. This could reflect that once a large enough market share is captured, the Founder may prefer to limit the permissiveness of the license to reduce its mismatch cost.[39] Conversely, when mismatch costs are higher, then the Founder selects a less permissive license to reduce its mismatch cost.

This result connects with anecdotal evidence from Paint.NET and the 2022 FOSS scandal involving Marak Squires.[27,40] Highlighting how the initial selection of the permissive MIT license could come back to haunt the Founder when the market evolves.

Importantly, this result suggest that the one size fits all recommendation of the MIT license might be an equilibrium. This is because selecting a permissive license best allows a Founder to take advantage of the virtuous cycle of network effects. Hence, despite its critics, it might be easier for Founders to converge upon the MIT license when trying to start a FOSS.

3.5.2 Skilled Users

An interesting observation made by Lerner and Tirole (2002) suggests that self-interested Founders tend to be highly skilled, and developers to FOSS tend to be more skilled than the average programmer. To reflect this observation, I propose the following modification: Developers are heterogeneous in skill level and this is uniformly distributed according to their location along the Hotelling line. Those located closer to 1 have a higher ability, normalised to $\alpha > 0$ at 1, and those further away lower ability, where the user located at 0 has 0 ability. I assume that if a developers chooses to contribute to the FOSS, their ability limits the level of their contribution.

This brings about two effects. First, the value of the software is dependent on the total skill of developers, with developers located closer to 1 generating more value, while those further away generate less value. This means the value of the FOSS is $v_o = \gamma \int_{\alpha\bar{x}}^{\alpha} \frac{x}{\alpha} dx = \gamma \frac{\alpha(1-\bar{x}^2)}{2}$. Second, because developers with a higher skill level are more likely to contribute to more components/critical components of the software, the location of the software is more affected by their output. This means that the location of the FOSS is at $L_o = 1 - l \int_{\alpha\bar{x}}^{\alpha} \frac{x}{\alpha} dx = 1 - \frac{\alpha(1-\bar{x}^2)}{2}$.

[39] Admissibly this may be a result of the features of a Hotelling model.

[40] Marak Squires is the Founder of two popular Open-Source libraries, Faker.js and Colors.js, both using the MIT license, decided to upload malicious code in place of the original out of frustration (Emma Roth, 2022).

Proposition 5. *When there exist heterogeneous skill levels, a self-interested Founder prefers contributions to be unique, $\gamma = 1$.*

See C.1.16 for proof.

Proposition 5 suggests that when there are heterogeneous skill levels, with more skilled developers located further away from the proprietary software, a self-interested Founder prefers unique contributions. This is because the location of the software shifts less for each additional developer than in the base model. As such, the mismatch cost to the Founder increases at a lower rate than the base model, allowing the Founder to prefer more distinct contributions.

3.5.3 Non-User Developers

I begun with a simplified model where users are also developers. Here, I consider the opposite extreme, where users and developers are mutually exclusive groups. In reality, developers are often a subset of users, and the groups do not perfectly overlap nor are they mutually exclusive. In showing these extremes, I expect that when developers are a subset of users, such a model will sit in-between the results shown in the main model and here.

Recall from Section 3.2 that users utility is $u_i = v_i - p_i - t|L_i - x|$. v_i captures the quality of the software, in particular $v_o = \gamma \cdot D_D$ where γ is the coordinator's effort to manage developers and D_D is the mass of developers. This is virtually identical to the main model, except that $D_D \neq D_o$.

Also recall that developers exhibit preferences $o \sim U(0, 1)$ along the same Hotelling line as users. Their utility is $w_o = s_o - k|L_o - o|$, s_o captures benefit to the developer, $k > 0$ the cost associated with working on a FOSS with features they do not prefer. As suggested in the literature, developers may be motivated by skill development, and recognition both within the community and outside the community (Hars & Ou, 2002; Li et al., 2012; Mustonen, 2003). The opportunities to work on complex tasks, or to gain prominence demand on the number of users of a FOSS. Hence, $s_o = \beta \cdot D_o$, where β proxies the prominence and recognition each additional FOSS user brings.

The timing of the game is identical to the main model, a self-interested Founder chooses the level of coordination, γ, the proprietary firm sets the price, p_p, users and developers simultaneously decide between consuming the proprietary software or the FOSS and contributing to the FOSS respectively.

Proposition 6. *Founders only coordinate FOSS in one of the following two situations:*

- *Proprietary software is low quality,* $v_p < t$ *and the FOSS license is restrictive* $l < \frac{1}{2}$*;*
- *Proprietary software is high quality,* $v_p > t$ *and the FOSS license is permissive* $l > \frac{1}{2}$*,*

preferring distinct contributions when they do so, $\gamma = 1$*.*

Proposition 6 highlights how the ability to attract users, and hence developers, can influence when a Founder is active. When $v_p < t$, the proprietary software is less competitive, making it easier for the FOSS to attract users. In such circumstances, Founders, maximising their own utility, prefer restrictive licenses as this minimises their mismatch cost. Conversely, when $v_p > t$, and it is more difficult to attract users, Founders only coordinate FOSS if they are able to attract more users to join the FOSS. This is only possible when the FOSS characteristics appeal more towards the median user, and hence the Founder is only active if the FOSS license is sufficiently permissive.

3.6 Conclusion

This paper provides a fresh look at the development of Free and Open-Source Software. I account for the different motivations of each actor involved in the development process, using a stylised model with quality and location network effects. In addition to the traditional view of quality network effects, I argue that permissive licenses induce location network effects, and this promotes competition. I discuss the extent to which a coordinator controls the complementarity of each developer's contribution to the FOSS, showing that this depends on the quality of proprietary software in the market, and the type of license adopted by the FOSS. I discuss how coordinators incentives changes the way FOSS are developed, studying 3 different motivations: (i) self-interested Founders, (ii) volunteering Altruists, (iii) profit-driven Managers.

Self-interested Founders are only active if the FOSS has a sufficiently restrictive license. When licenses are more restrictive, they induce developers to provide distinct contributions to the FOSS. Altruistic coordinators are always active, and prefer distinct contributions when they face low-quality proprietary software. Profit-driven Managers are only active when licenses are both sufficiently permissive and restrictive, allowing them to attract more developers, but minimising horizontal competition. They also prefer distinct contributions when competing with a low quality proprietary software.

This model captures the proliferation of niche FOSS that exist, and serve very small communities, despite prominent proprietary firms. Moreover, findings about profit-driven Managers closely relate to the growth of Google's Chrome, and their decision to adopt a code release strategy.

I study the surplus of users, showing that more permissive licenses can both benefit or harm users of FOSS and proprietary software when the Founder coordinates development. However, when the coordinator is an Altruist, more permissive licenses harm users of the FOSS but drive competition with firms, which benefits users of the proprietary software. Moreover, when coordinators are profit-driven Managers, permissive licenses drive competition and decrease the profitability of the FOSS and proprietary software.

I show that both Altruists and Managers fail to maximise total surplus, while Founders generate a higher level of total surplus than Altruists. I also study a series of extensions where users have heterogeneous skill levels, users and developers are mutually exclusive groups, and consider the Founder's initial choice of license.

My stylised model provides a new view of network effects in the context of a Hotelling model, showing this can generate interesting competitive effects. This model is then able to better capture developments of FOSS. Because of its stylistic nature, my model does not consider the situation where choice of license and degree of contributions to a FOSS can both influence the quality of the proprietary software.

At its heart, this paper is about social movements, organised efforts by a group of people to achieve a common goal. Such movements often begin with an extremely dissatisfied individual. Facing sufficient unhappiness, they may find it beneficial to coordinate a group of like-minded individuals to address their shared dissatisfaction (Breton & Breton, 1969). As groups grow, they become more valuable to their members. However, differences of opinions are more likely to manifest with larger groups. This friction can diminish the value of the group to an individual. Therefore, in addition to Free and Open-Source Software, location network effects can also be found in other social movements such as crowd-sourcing and resource sharing, and social and political activism. I leave such applications for future research.

Appendix A
Ratings and Reciprocity

A.1 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^j_{t,R_t} 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:* $p^H_{1,1} = min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p^H_{1,1}]\}$
- *Low-quality firm:*

$$p^L_{1,1} = \kappa q^L - \frac{e}{\Delta}$$

$$p^L_{1,0} = E[q_1|R_0, p^H_{1,1}]$$

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 $p^H_{1,1} = \min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p^H_{1,1}]\}$ with probability 1.

Given Restriction 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,

$$[\kappa q^j - p_{t,1}^j]\Delta - e \geq 0 \Leftrightarrow p_{t,1}^j \leq \overline{p_t^j} \equiv \kappa q^j - \frac{e}{\Delta}.$$

Therefore, the upper bound on prices such that the high-quality firm receives a positive rating is $\overline{p_1^H} = \kappa q^H - \frac{e}{\Delta}$.

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 Restriction 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 $p_{1,1}^H = \min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p_{1,1}^H]\}$ with probability 1. To see this, note that $\overline{p_1^L} = \kappa q^L - \frac{e}{\Delta}$ 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 Restriction 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 $p_{1,1}^H = \min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p_{1,1}^H]\}$.

We conclude that high-quality firms set a unique price $p_{1,1}^H = \min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p_{1,1}^H]\}$ 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 $p^L_{1,1} = \kappa q^L - \frac{e}{\Delta}$.

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 $\overline{p^L_1}$, $\overline{p^L_1} = \kappa q^L - \frac{e}{\Delta}$. 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 $\overline{p^L_1}$. Since demand and ratings are the same for all prices weakly below $\overline{p^L_1}$, a low-quality firm that obtains a rating must optimally set $\overline{p^L_1}$ with probability 1. We conclude that conditional on obtaining a rating—the low-quality firm sets $\overline{p^L_1}$ 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 $E[q_1|R_0, p^H_{1,1}]$. We have shown in step 2 that all prices above $\overline{p^L_1}$ induce the same beliefs $E[q_1|R_0, p^H_{1,1}]$. Thus, low-quality firms who obtain a no rating optimally set the highest possible price, $E[q_1|R_0, p^H_{1,1}]$ 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*

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} & \text{if } p_1 > \overline{p^L_1} \\ q^L & \text{if } p_1 \leq \overline{p^L_1}, \end{cases}$$

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 $p^H_{1,1}$ and obtain a rating with probability 1, and low-quality firms obtain a rating if and only if they charge a low price $p^L_{1,1}$. We denote the probability δ as the probability with which low-quality firms charge $p^L_{1,0}$. 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 Restriction 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) $p^H_{1,1} = \min\{\kappa q^H - \frac{e}{\Delta}, E[q_1|R_0, p^H_{1,1}]\} = E[q_1|R_0, p^H_{1,1}]$ and (ii) $p^H_{1,1} = \min\{\kappa q^H - \frac{e}{\Delta}, E[q_1|R_0, p^H_{1,1}]\} = \kappa q^H - \frac{e}{\Delta}$.

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

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

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

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

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

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

Because $p^H_{1,1} > \overline{p^L_1}$, $p^L_{1,0} > \overline{p^L_1}$, and $p^L_{1,1} \leq \overline{p^L_1}$, this concludes the proof. □

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

Proposition 6. *All perfect Bayesian equilibria satisfy the following. In period 1:*

1. *High-quality firms receive a good rating with probability 1 and charge* $\overline{p} \equiv E[q_1|R_0, p^H_{1,1}]$.
2. *Low-quality firms randomize their strategy.*
    a. *They charge* $\overline{p}$ *and obtain no rating with probability* δ^*.
    b. *They charge* $\underline{p} \equiv \kappa q^L - \frac{e}{\Delta}$ *and obtain a good rating with probability* $1-\delta^*$.
3. *Consumers beliefs of equilibrium prices are given by* $E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma+\delta^*(1-\gamma)} & \text{if } p_1 > \kappa q^L - \frac{e}{\Delta} \\ q^L & \text{if } p_1 \leq \kappa q^L - \frac{e}{\Delta} \end{cases}.$

In period 2:

4. *Prices are equal to expected quality conditional on ratings.*
5. *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* $\kappa q^L - \frac{e}{\Delta} + q^H \leq \gamma q^H + (1-\gamma)q^L + q^L$, *and* $\delta^* \in (\frac{1}{2}, 1)$ *otherwise. The equilibrium exists if* $\kappa q^H - \frac{e}{\Delta} \geq \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma+\delta^*(1-\gamma)}$.

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

Proof of Proposition 6.
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 $p^H_{1,1} = \min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p^H_{1,1}]\}$, and it remains to show that $\min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p^H_{1,1}]\} = E[q_1|R_0, p^H_{1,1}]$. Suppose towards a contradiction that $\min\{\kappa q^H - \frac{e}{\Delta},\ E[q_1|R_0, p^H_{1,1}]\} = \kappa q^H - \frac{e}{\Delta}$. As we argued in the proof of Lemma 3, we then have $p^H_1 \neq p^L_{1,0}$, and $E[q_1|R_0, p^H_{1,1}] = q^H$, and $p^L_{1,0} = E[q_1|R_0, p^H_{1,1}] = q^H$ is played with probability $\delta^* = 0$. Because the low-quality firm charges $p^L_{1,1} = \kappa q^L - \frac{e}{\Delta}$ with probability 1, all firms get a rating with probability 1. By Lemma 3, low-quality firms earn up to $\kappa q^L - \frac{e}{\Delta} + \gamma q^H + (1-\gamma)q^L$. Low-quality firms can strictly increase profits by charging a first period price of q^H. By Restriction 2, consumers believe $E[q_1|R_0, p_1] = q^H$ for all prices $p_1 \geq \kappa q^L - \frac{e}{\Delta}$, 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 $q^L > \kappa q^L - \frac{e}{\Delta}$, this deviation profitable for low-quality firms, a contradiction.

We conclude that $p^H_{1,1} = \min\{\kappa q^H - \frac{e}{\Delta}, E[q_1|R_0, p^H_{1,1}]\} = E[q_1|R_0, p^H_{1,1}]$, which proves statement 1. Because we have shown that $p^H_{1,1} = E[q_1|R_0, p^H_{1,1}]$, and this is the same as $p^L_{1,0}$, to simplify notation, we state that $\overline{p} = p^H_{1,1} = p^L_{1,0} = E[q_1|R_0, \overline{p}]$. To further simplify notation, we label $\underline{p} = p^L_{1,1} = \kappa q^L - \frac{e}{\Delta}$.

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 proofs 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 $\kappa q^H - \frac{e}{\Delta} > E[q_1|p^H_{1,1}]$, implying that $\overline{p} = \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$.

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

$$\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)}. \tag{A.1}$$

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

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

$$\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} + q^L, \tag{A.2}$$

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 (A.1) and (A.2) become $\kappa q^L - \frac{e}{\Delta} + q^H$ and $\gamma q^H + (1 - \gamma)q^L + q^L$, respectively. For $\delta^* = 1$ to be optimal, it must be the case that

$$\kappa q^L - \frac{e}{\Delta} + q^H \leq \gamma q^H + (1-\gamma)q^L + q^L. \tag{A.3}$$

Since (A.1) strictly increases in δ^* and (A.2) 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 (A.1) and (A.2) become $\kappa q^L - \frac{e}{\Delta} + \gamma q^H + (1-\gamma)q^L$ and $q^H + q^L$, respectively. But since $q^H > q^H + (1-\gamma)q^L$ and $q^L > \kappa q^L - \frac{e}{\Delta}$, low-quality firms optimally set $\overline{p} = q^H$ 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 (A.1) = (A.2) and low-quality firms must play some δ^* such that these beliefs are correct. Thus, in a mixed-strategy equilibrium, we have

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

We have two candidates that solve this equation:

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

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

Consider the scenario where the last term is subtracted.

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

.

Therefore, we conclude that

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

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

We see from (A.1) and (A.2) that $\delta^* < 1$ if $\kappa q^L - \frac{e}{\Delta} + q^H > \gamma q^H + (1-\gamma)q^L + q^L$. Also note that

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

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 (A.4) holds, it must be unique, because (A.1) strictly increases, and (A.2) strictly decreases in δ^*. We conclude that if mixed-strategy equlibria 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

$$(1-\gamma)(q^H - q^L) > q^L(1-\kappa) + \frac{e}{\Delta}. \tag{A.5}$$

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.

$$\delta^* = \begin{cases} \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1-\gamma)(q^L(1-\kappa)+\frac{e}{\Delta})} + \frac{(4\gamma^2(q^H-q^L)^2+(1+\gamma)^2(q^L(1-\kappa)+\frac{e}{\Delta})^2)^{\frac{1}{2}}}{2(1-\gamma)(q^L(1-\kappa)+\frac{e}{\Delta})} & \text{if (A.5) holds} \\ 1 & \text{otherwise} \end{cases} \tag{A.6}$$

We now show that these equilibria exist.

To start, consider the case where (A.3) holds. In the candidate equilibrium in period 1, the high-quality firm sets $\overline{p} = \gamma q^H + (1-\gamma)q^L$ with probability 1 and obtains a rating, and the low-quality firm charges $\overline{p}$ 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

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

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 restrictions. Because high-quality firms obtain a rating with probability 1, the candidate equilibrium is consistent with Restriction 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 Restriction 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 $\kappa q^L - \frac{e}{\Delta}$ does not improve the rating and only reduces margins without raising demand, this is not a profitable deviation. Deviations to lower prices below $\kappa q^L - \frac{e}{\Delta}$ leads to profits weakly below $\kappa q^L - \frac{e}{\Delta} + q^H$, which is not a profitable deviation since (A.3) holds.

We conclude that if (A.3) holds, no profitable deviations exist for either type of firm.

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

We conclude that if $(1-\gamma)(q^H - q^L) \leq q^L(1-\kappa) + \frac{e}{\Delta}$ and $\gamma q^H + (1-\gamma)q^L \leq \kappa q^H - \frac{e}{\Delta}$, 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 (A.5) 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 $\overline{p} = \frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$ with probability 1 and obtains a rating. The low-quality firm charges $\overline{p}$ with probability δ^* and gets no rating, and sets $\underline{p} = \kappa q^L - \frac{e}{\Delta}$ 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

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

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 restrictions. Because high-quality firms obtain a rating with probability 1, the candidate equilibrium is consistent with Restriction 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 Restriction 2.

We now show that firms have no profitable deviations.

We start with low-quality firms, who earn total profits $\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta^*)(1-\gamma) q^L}{\gamma + (1-\delta^*)(1-\gamma)}$. The firm is indifferent between charging $\overline{p}$ and $\underline{p}$, where $\overline{p} > \underline{p}$. If the firm deviates to a price above $\overline{p}$, demand drops to zero and total profits are weakly below $0 + q^L$, this is not a profitable deviation. Deviations to a price $p_1 \in (\underline{p}, \overline{p})$, for which the firm gets the same rating as charging $\overline{p}$ and therefore earns the same profit in period 2, but the firm earns a lower margin than when it charges $\overline{p}$ 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 $\overline{p} = \min\{\kappa q^H - \frac{e}{\Delta}, E[q_1|\overline{p}]\} = E[q_1|\overline{p}]$, which requires
$\frac{\gamma q^H + \delta^*(1-\gamma) q^L}{\gamma + \delta^*(1-\gamma)} \leq \kappa q^H - \frac{e}{\Delta}$.

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

This concludes the proof. $\square$

Proof of Proposition 2

Proof of Proposition 2.
In this proof, we show that $\frac{\partial\delta^*}{\partial\Delta} < 0$ if $(1-\gamma)(q^H - q^L) - (1-\kappa)q^L > \frac{e}{\Delta}$ and argue it is 0 otherwise.

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

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

When $(1-\gamma)(q^H - q^L) - (1-\kappa)q^L \leq \frac{e}{\Delta}$, $\delta^* = 1$, and further increases to Δ cannot increase δ^* as it is is a probability and is weakly bound by 0 and 1.

Thus we have shown that when low-quality firms would play a mixed-strategy, as consumers kindness increases, low-quality firms are more likely to participate in ratings harvesting. Otherwise, low-quality firms play a pure strategy of price mimicking.

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 $(1-\gamma)(q^H - q^L) - (1-\kappa)q^L > \frac{e}{\Delta}$ holds.

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

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

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 3

Proof of Proposition 3.
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, $\overline{\gamma}$, and show its uniqueness and existence.

To start, note from (A.4), δ^* is such that low-quality firms are indifferent between charging low prices and high prices,

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

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

$$\frac{\partial \delta^*}{\partial \gamma} = \frac{(2\delta^* - 1)(\gamma^2 + \delta^{*2}(1-\gamma)^2 - \delta^*(1-\gamma)^2)}{\gamma(1-\gamma)((\gamma + (1-\delta^*)(1-\gamma))^2 + (\gamma + (1-\gamma)\delta^*)^2)} \tag{A.7}$$

In particular, we show that the sign of $\frac{\partial \delta^*}{\partial \gamma}$ switches at some point $\overline{\gamma}$. 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.

$$\overline{\gamma} = \frac{(q^H - q^L)^2 - (q^L(1-\kappa) + \frac{e}{\Delta})^2}{3(q^H - q^L)^2 + (q^L(1-\kappa) + \frac{e}{\Delta})^2} \tag{A.8}$$

We show that $\overline{\gamma} \in (0, 1)$. The denominator is strictly positive. From (1.1), the numerator is strictly positive. Hence, $\overline{\gamma} > 0$. Moreover, because the denominator is always strictly larger than the numerator, $\overline{\gamma} < 1$.

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

Notice that when $\gamma \to 0$, and from (A.6) $\delta^* \to 1$. Thus $\delta^*(1-\delta^*) \to 0$ and $\frac{\gamma^2}{(1-\gamma)^2} \to 0$. And when $\gamma \to 1$, $\delta^* = 1$. 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 $\gamma \to 0$, $\frac{\partial \frac{\gamma^2}{(1-\gamma)^2}}{\partial \gamma} = \frac{2\gamma}{(1-\gamma)^3} \to 0$. Further, consider that $\frac{\partial \delta^*(1-\delta^*)}{\partial \gamma} = \frac{\partial \delta^*}{\partial \gamma}(1-2\delta^*)$. Evaluated at $\gamma \to 0$, $\delta^* \to 1$ 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 $\gamma \to 1$, $\delta^* \to 1$ and $\frac{\gamma^2}{(1-\gamma)^2} \to \infty$. 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 $\gamma \to 1$, $\frac{\partial \delta^*(1-\delta^*)}{\partial \gamma} = \frac{\partial \delta^*}{\partial \gamma}(1-2\delta^*) < 0$. 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 $\overline{\gamma} = \gamma'' = \gamma'$.

We conclude that there is a unique switching point $\overline{\gamma}$, 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 $\gamma = \overline{\gamma}$, and $\overline{\gamma} \in (0,1)$. □

Proof of Proposition 4

Proof of Proposition 4.
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.

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

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

$$CS = (1-\delta^*)(1-\gamma)[q^L - \kappa q^L + \frac{e}{\Delta}]. \tag{A.9}$$

This is (1.2). From (1.2) we can evaluate the effects of changes to effort cost.

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

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

This is positive if and only if

$$(1-\gamma)^2\gamma(q^H-q^L)^2 > (1+\gamma)^2(q^L(1-\kappa)+\frac{e}{\Delta})^2$$

, and $\frac{\partial CS}{\partial e} = 0$ with equality. This condition is (1.3).

We now look at the second derivative.

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

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

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

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

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

where e^{cs} is the level of effort cost that maximises consumer surplus. This e^{cs} is indeed positive when (1.3) holds. Which is the restriction required for $\frac{\partial CS}{\partial e} > 0$ at $e = 0$.

When (1.3) 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,
$e^{cs} = -(1-\kappa)q^L\Delta + \frac{\Delta(q^H-q^L)(1-\gamma)\sqrt{\gamma}}{1+\gamma}$ when (1.3) 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,

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

Taking the derivative to e,

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

and this is negative when (1.3) 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 (1.1) 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 (1.3) is a stricter condition than (1.1).

From (1.3),

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

And from (1.1),

$$(1-\gamma)(q^H - q^L) - (1-\kappa)q^L > \frac{e}{\Delta}$$
$$(1-\gamma)(q^H - q^L) > q^L(1-\kappa) + \frac{e}{\Delta} \text{ since both sides are positive,}$$
$$(1-\gamma)^2(q^H - q^L)^2 > (q^L(1-\kappa) + \frac{e}{\Delta})^2$$
$$1 > \frac{(q^L(1-\kappa) + \frac{e}{\Delta})^2}{(1-\gamma)^2(q^H - q^L)^2}$$

Finally, notice that $\gamma \in (0,1)$ and thus $\frac{\gamma}{(1+\gamma)^2} < 1$. Therefore, (1.3) is a stricter condition than (1.1).

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. □

Proof of Proposition 5

Proof of Proposition 5.
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 4 that total consumer surplus reduces to (1.2), i.e. $CS = (1-\delta^*)(1-\gamma)[q^L - \kappa q^L + \frac{e}{\Delta}]$. We take the derivative of (1.2) with respect to γ and show that there is some $\gamma^{cs} < \overline{\gamma}$ that maximises consumer surplus.

$$\frac{\partial CS}{\partial \gamma} = -(1-\delta^* + (1-\gamma)\frac{\partial \delta^*}{\partial \gamma})(q^L(1-\kappa) + \frac{e}{\Delta}) \tag{A.10}$$

Since $q^L(1-\kappa) + \frac{e}{\Delta} > 0$, therefore, 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.

$$\gamma^{cs} = \frac{-(q^H - q^L)(q^L(1-\kappa) + \frac{e}{\Delta})^2}{(q^H - q^L)(4(q^H - q^L)^2 + (q^L(1-\kappa) + \frac{e}{\Delta})^2)} \pm \frac{((q^H - q^L)^3(q^L(1-\kappa) + \frac{e}{\Delta})(2(q^H - q^L) - (q^L(1-\kappa) + \frac{e}{\Delta}))^2)^{\frac{1}{2}}}{(q^H - q^L)(4(q^H - q^L)^2 + (q^L(1-\kappa) + \frac{e}{\Delta})^2)}$$

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 $(q^H - q^L)(2(q^H - q^L) - (q^L(1-\kappa) + \frac{e}{\Delta}))^2 \geq (q^L(1-\kappa) + \frac{e}{\Delta})^3$.

We now show that $(q^H - q^L)(2(q^H - q^L) - (q^L(1-\kappa) + \frac{e}{\Delta}))^2 \geq (q^L(1-\kappa) + \frac{e}{\Delta})^3$ always holds. By applying $(q^H - q^L) > \gamma q^H + (1-\gamma)q^L - \kappa q^L + \frac{e}{\Delta} > q^L(1-\kappa) + \frac{e}{\Delta}$,

$$(q^H - q^L)(2(q^H - q^L) - (q^L(1-\kappa) + \frac{e}{\Delta}))^2 \geq (q^L(1-\kappa) + \frac{e}{\Delta})^3$$
$$(q^H - q^L)(q^L(1-\kappa) + \frac{e}{\Delta})^2 > (q^L(1-\kappa) + \frac{e}{\Delta})^3$$
$$(q^H - q^L) > (q^L(1-\kappa) + \frac{e}{\Delta})$$

This is a weaker condition than (1.1).

Thus 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 $\gamma^{cs} < \overline{\gamma}$. 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 $\overline{\gamma}$ solves $\frac{\partial \delta^*}{\partial \gamma}$ implies $\frac{\partial \delta^*}{\partial \gamma} = 0$ at $\overline{\gamma}$. Further, $\delta^* \in (\frac{1}{2}, 1)$ for $\gamma > 0$, implying that $(1-\delta^*) > 0$. Therefore, from (A.10), at $\overline{\gamma}$, $\frac{\partial CS}{\partial \gamma} < 0$. This implies that $\gamma^{cs} < \overline{\gamma}$.

Second, we show that when $\gamma > \gamma^{cs}$, $\frac{\partial CS}{\partial \gamma} < 0$. When $\gamma > \overline{\gamma}$, we know that $(1-\delta^*) > 0$, $(1-\gamma) > 0$, and from Proposition 3 $\frac{\partial \delta^*}{\partial \gamma} > 0$. 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 $(1 - \delta^* + (1-\gamma)\frac{\partial \delta^*}{\partial \gamma}) \to 1 - \frac{q^H - q^L}{q^L(1-\kappa) + \frac{e}{\Delta}}$ when evaluated at $\gamma \to 0$. From (1.1) we know that this is strictly negative. Therefore, $\frac{\partial CS}{\partial \gamma} > 0$ at $\gamma \to 0$. 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 $\gamma^{cs} \in (0, \overline{\gamma})$ that maximises consumer surplus,

$$\begin{aligned}\gamma^{cs} = &\frac{-(q^H - q^L)(q^L(1-\kappa) + \frac{e}{\Delta})^2}{(q^H - q^L)(4(q^H - q^L)^2 + (q^L(1-\kappa) + \frac{e}{\Delta})^2)} + \\ &\frac{((q^H - q^L)^3(q^L(1-\kappa) + \frac{e}{\Delta})(2(q^H - q^L) - (q^L(1-\kappa) + \frac{e}{\Delta}))^2)^{\frac{1}{2}}}{(q^H - q^L)(4(q^H - q^L)^2 + (q^L(1-\kappa) + \frac{e}{\Delta})^2)}\end{aligned}$$

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.

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

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

Therefore,

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

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

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

$$\pi^H = \frac{\gamma q^H+\delta^*(1-\gamma)q^L}{\gamma+\delta^*(1-\gamma)}+\frac{\gamma q^H+(1-\delta^*)(1-\gamma)q^L}{\gamma+(1-\delta^*)(1-\gamma)}$$
$$\frac{\partial \pi^H}{\partial \gamma} = \frac{32\gamma^2(q^H-q^L)^3(q^L(1-\kappa)+\frac{e}{\Delta})^2\sqrt{4\gamma^2(q^H-q^L)^2+(1+\gamma)^2(q^L(1-\kappa)+\frac{e}{\Delta})^2}}{(\gamma+\delta^*(1-\gamma))^2(\gamma+(1-\delta^*)(1-\gamma))^2} - \frac{32\gamma^2(q^H-q^L)^3(q^L(1-\kappa)+\frac{e}{\Delta})^2(2\gamma(q^H-q^L))}{(\gamma+\delta^*(1-\gamma))^2(\gamma+(1-\delta^*)(1-\gamma))^2}$$

Because $\sqrt{4\gamma^2(q^H-q^L)^2+(1+\gamma)^2(q^L(1-\kappa)+\frac{e}{\Delta})^2}-2\gamma(q^H-q^L)>0$, 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. □

A.2 Web Appendix

A.2.1 Extensions and Robustness

Designing a profit maximizing ratings system

So far we studied how different features of a marketplace environment affects the informativeness of ratings and surplus. In practice, rating systems are typically designed by online platforms, and many platforms constantly tweak their ratings system. We now investigate how a platform would actually design its ratings system. To do so, we introduce some modifications to our base model.

Suppose there is a unit mass of consumers who are heterogeneous only in their outside option to joining the platform. This outside option is uniformly distributed between

$[0,1]$. We denote the mass of consumers who join the platform as n_b. Further, we call the per transaction consumer benefit u_b. Consumers choose to join the platform and are then randomly assigned (with equal probability) to the first period or the second period. This captures that consumers use the platform both when there are more or less ratings available to guide their choices.

We suppose that there exists a unit mass of sellers that are monopolists in a product category. Sellers are ex-ante homogeneous and face an outside option of $\bar{v}_s$. We denote the mass of sellers who join the platform by n_s and call π_s the per transaction revenue of the seller. We assume that sellers face some additional platform specific marginal cost of selling on the platform, t.[1] To simplify illustration, we suppose that sellers join the platform prior to period 1 of our main model and therefore do not have information over their true quality when deciding to join the platform.[2]

To maximize profits, the platform makes two choices. First, the platform sets a royalty, r, which it charges sellers. This gives the platform a share r of the sellers' revenue. Second, the platform chooses how easy it is for buyers to leave a rating, i.e. the platform sets e. We assume that the platform can choose any $e \in [0, \bar{e}]$. To simplify exposition, we assume $\bar{e} \leq e^{cs}$. This implies that a larger effort always increases u_b and reduces π_s. The platform makes these decisions prior to period 1 in the base model.

The effort choice e captures that the platform designs a ratings system to make it easier or more difficult for consumers to leave a rating. For example, to facilitate ratings, a platform can introduce a one-click ratings system or provide users a link to go directly to the ratings page, automated ratings[3], etc. Conversely, a platform can make rating more costly for consumers if it introduces additional authentication and verification steps—such as proof of identity, proof of purchase[4], multiple ratings components—such as the use of multi-dimensional ratings (Schneider, Weinmann, Mohr, & vom Brocke, 2021a), or requiring a written review along with the rating[5].

Proposition 7. *Platforms design a ratings environment that favours sellers ($e^* = e^s = 0$) if and only if $\pi_s - t < u_b$ at $e = 0$.*

[1]This cost captures the difference in cost of selling on a platform rather than direct to consumers. Reflecting costs in addition to the platforms ad valorem fees. For instance, on Amazon, in addition to the ad valorem fees, there are additional charges for fulfilled by Amazon and other transaction fees.

[2]This assumption captures the importance of relative quality of products on a marketplace, and sellers learning their true relative quality after joining the marketplace. This assumption also allows us to abstract away from seller selection by the platform and focus on the role that the platform plays in influencing ratings.

[3]Documented by a website which guides non-mandarin speakers the use of Taobao (TaobaoTranslate, 2021).

[4]As explained by Amazon (Amazon, 2021).

[5]As is required by Steam ("Introducing Steam Reviews", 2021).

The proposition characterizes when the platform wants to make it easier to rate. It captures the common finding that platforms want to balance surplus between buyers and sellers to generate value on the platform (Armstrong, 2006; Caillaud & Jullien, 2003; Rochet & Tirole, 2003). Combining this result with Proposition 4, we find that when $\pi_s - t < u_b$, the platform encourages ratings to raise $\pi_s - t$. Intuitively, a lower e raises profits per seller but also lowers the number of buyers who join the platform. $\pi_s - t < u_b$ implies that the former effect dominates. Thus, the platform facilitates ratings to lower reciprocity rents and shift surplus to sellers, but thereby undermines the informativeness of ratings.

These results support concerns of regulators that platforms design insufficiently informative ratings environments for consumers, and that minimum standards of ratings may help to protect consumers (Competition and Markets Authority (UK), 2017).

We made the simplifying assumptions that sellers are ex-ante identical and do not yet know their relative quality level when entering the platform. Corollary 2 suggests that relaxing this assumption could lead to a novel selection effect. Platforms may still prefer to lower e in order to raise average seller surplus. But this discourages high-quality firms from joining the platform and leads to lower average quality on the platform. This way, facilitating ratings can undermine a key purpose of a ratings system.

These results suggest that a platform may change its strategy over time. For example, a new platform may choose a relatively large e to get informative ratings that favor buyers in order to grow and reinforce network effects, but later lower e to extract more profits from sellers.

Competitive environment

Our basic model focuses on monopoly sellers. We now introduce competition to the model. To do so in a simple way, we assume that there exists a competitive fringe of non-strategic firms that offers a product of quality q at a price equal to their marginal cost c, where $q \geq 0$. Intuitively, we capture competition by changing the outside option of purchasing in the market to $q - c \geq 0$. This could capture (i) established firms for which consumers do not need ratings to evaluate the quality of their products; (ii) brick-and-mortar stores; or (iii) the expected utility a consumer gets when participating in another ratings environment such as another marketplace.

We now study how the presence of the competitive fringe affects the equilibrium strategy of our previous strategic sellers.

To start, suppose $q - c = 0$. Then the competitive fringe offers the same value as our previous outside option and the equilibrium is the same as in Proposition 1. Now suppose c decreases marginally and the competitive fringe offers a small but strictly positive surplus

$q - c > 0$. This encourages sellers to harvest more ratings and reduces the informativeness of ratings. Intuitively, in period 1 firms that charge $\overline{p}$ extract consumers' conditional expected value. As c decreases, firms can no longer extract all surplus and reduce $\overline{p}$ to remain more beneficial to consumers than the fringe. The low-quality firms that charge $\underline{p}$, however, already leave a rent to consumers, which is why $\underline{p}$ does not decrease. Thus, the fringe puts more pressure on $\overline{p}$ than on $\underline{p}$, which is why low-quality firms harvest more ratings, leading to less-informative ratings in equilibrium (lower δ^*).

As c decreases further, $q - c$ becomes so large that it puts equal pressure on $\overline{p}$ and $\underline{p}$, so that the competitive fringe no longer affects the incentives to harvest ratings and the informativeness of ratings, δ^*.

The following proposition summarizes this result.

Proposition 8. $\frac{\partial \delta^*}{\partial c} > 0$ *if* $q^L - (q - c) \geq \kappa q^L - \frac{e}{\Delta}$. *Otherwise,* $\frac{\partial \delta^*}{\partial c} = 0$

The proposition shows that competition puts disproportionate pressure on firms that extract more surplus from consumers, i.e. firms that do not harvest ratings. This is why competition encourages firms to harvest ratings and makes ratings less-informative. Crucially, however, ratings will not become uninformative: as the fringe becomes more competitive ($q - c$ increases), it will put equal pressure on all prices and no longer affect incentives to harvest ratings, leaving the informativeness of ratings constant at $\delta^* = \frac{1}{2}$.

Even though competition makes ratings less-informative, it unambiguously benefits consumers through the following two channels. First, competition exerts pressure on the higher price and lowers $\overline{p}$. Second, competition encourages low-quality firms to harvest ratings and charge $\underline{p}$, for which consumers receive a higher surplus. Thus, our results suggest that making consumers aware of alternative sellers can make ratings less-informative, but benefits consumers nonetheless.

Negative Ratings

For exposition we made the simplifying assumption of only good and no ratings. In this section, we relax this assumption to allow for more than 2 ratings, allowing consumers to rate positively, negatively, or not at all. In doing so, we show that if consumers rate value for money, only extreme ratings are played in equilibrium.

Formally, we setup a model that is identical to the base model, with the exception that $R_t \in \{1, 0, -1\}$, representing good, bad or no ratings respectively. Consumers provide a negative rating when they receive what they perceive to be unfair, and want to respond with negative reciprocity, potentially harming the firm. Hence, instead of obtaining a positive warm glow of Δ, we say consumers exhibit a negative warm glow of $-\Delta$ when leaving a bad rating. Finally, consumers only rate (bad or good) if they have a weakly positive rating utility.

In this setting, we show that there exist a range of prices for which each rating is possible: (i) when consumers receive a sufficiently high value-for-money, they choose to leave a good rating; (ii) when consumers receive a sufficiently low value-for-money, they choose to leave a bad rating; (iii) and for middle levels of value-for-money, consumers choose not to leave a rating.

However, there exist equilibria where low-quality firms choose between setting the prices of $\overline{p}$ and $\underline{p}$. This is because with Restriction 1, we search for equilibria where consumers' beliefs are such that both $R_t \in \{0, -1\}$ are a reflection of a low-quality firm. This means that, conditional on receiving no or a bad rating, firms receive the same pay off in the second period. Hence, a profit maximising firm sets the highest possible price of $\overline{p}$ if it chooses not to receive a good rating.

Proposition 9 summarises the implication of this.

Proposition 9. *There exists a range of prices for which good, no and bad ratings may occur. In equilibrium, it is always beneficial for low-quality firms to obtain a bad rating over no rating, if consumers continue to buy at the price that induces a bad rating.*

For completeness, we conduct a minor extension to show that our model connects well with suggestions in the empirical literature that consumers find it more difficult to leave a bad rating than they do a good rating (Cabral & Hortaçsu, 2010; Dellarocas & Wood, 2008; Filippas & Horton, 2022; Filippas et al., 2022). To do so, we modify the model with negative ratings to study the situation where consumers face different costs to leaving good and bad ratings. We say that consumers face a cost of e_b when leaving a bad rating, and a cost of e_g when leaving a good rating, where $e_b > e_g$.

Lemma 4. *When it is more difficult to leave a bad rating than a good rating, firms are more likely to obtain good ratings, and less likely to obtain bad ratings.*

Lemma 4 shows that when e_b is larger, firms are able to set a much higher price before they obtain a bad rating. This result implies that when e_b is sufficiently large, low-quality firms may not receive any bad ratings in equilibrium. Moreover, because e_g is relatively smaller, combined with Corollary 1, low-quality firms are likely to participate in ratings harvesting and receive a good rating.

We show that extreme ratings by the low-quality firm is an equilibrium result in a model of more than 2 ratings, and our model does not lose any information by considering a binary ratings system. Further, when considering the difference in cost needed to provide a rating, we provide theoretical foundation for the J-shape distribution of ratings found in the empirical literature, suggesting that it is an equilibrium result of reciprocity-based ratings (Dellarocas & Wood, 2008; Filippas et al., 2022; Hu et al., 2009).

Continuum of quality

For exposition, we discussed a binary setting of firms with only high or low quality. In this section, We show that our results are robust to firms with a continuum of quality types.

Formally, we setup a model that is identical to the base model, with the following modifications: firm's quality is uniformly and continuously distributed on the interval $[0, 1]$; for simplicity there are no ratings prior to period 1 and consumers do not leave a rating in period 2; our solution concept is a pure-strategy perfect Bayesian equilibrium.

We show there exists an equilibrium where firms can be divided into 3 groups: (i) higher quality firms which set the highest possible price and receive a good rating in equilibrium; (ii) middle quality firms which choose to set lower prices, depending on their quality, to participate in ratings harvesting; (iii) lower quality firms that choose to set the highest possible price and forgo ratings.

Proposition 10. *There exists an equilibrium where: Firms with quality $q \in [\hat{q}, 1]$ obtain a good rating in equilibrium. Firms with quality $q \in [\hat{\hat{q}}, \hat{q})$ choose to participate in ratings harvesting in equilibrium. Firms with quality $q \in [0, \hat{\hat{q}})$ do not obtain a rating in equilibrium.*

In equilibrium, firms of low-quality obtain no rating and high- and middle- quality a good rating. Our result that easier ratings lead to less-informative ratings is robust: as it becomes easier to rate, firms of more types receive a good rating, reducing the expected quality of firms with a good rating. As in the main text, this occurs because as e decreases, firms harvest more ratings, causing $\hat{\hat{q}}$ to fall. This suggests that ratings become less informative because a good rating represents a wider range of quality types. However, despite ratings becoming less informative, the lowest quality firms still receive a good rating, suggesting that ratings are probably still somewhat useful for signalling quality.

Multi-period model

We check for robustness to time horizon by looking at a three period model. In this model, we find that equilibria similar to those described in our main model exist.

Notice that in a three period model, ratings take the following possible histories:

- $R_0 = \{0\}$

- $R_1 = \{01, 00\}$
- $R_2 = \{011, 010, 001, 000\}$

To save on notation, we omit the 0 in the initial period, capturing that all sellers start with the same reputation. We also extend Restriction 1 to reflect the additional period.

Restriction 3. *We focus on the equilibria where high-quality firms obtain a rating of $R_t = 1$, $t \in \{1, 2\}$ with probability 1.*

Further, in each period low-quality firms may decide to randomize with different probabilities, δ_t. Aside from the addition of a third period, the model remains the same as that of the main body of this paper.

Our intention is to uncover equilibrium strategies similar to the main body of the paper and not to consider all possible equilibria. We show that our findings continue to exist in a 3 period model. That is, we show that low-quality firms may choose to play a mixed strategy in every period. The result differs from our base setting in that a stricter set of restrictions is required for a unique mixed strategy to exist in every period.

We show that there exist equilibria where the low-quality firm chooses to play a mixed strategy in period 1 and in period 2 play a pure strategy: they harvest ratings if they received a good rating in the past or charge q^L otherwise. This means that low-quality firms choose to obtain good ratings in the first period even if they are unable to sustain this ratings outcome. As a result, we believe that our results are not driven by the end game effect of the final period. Instead, this effect is continuous and trading off pay offs in earlier periods for a big pay off in latter periods is a consideration made independent of the terminal period.

This is summarised in the following proposition

Proposition 11. *A mixed strategy equilibrium with properties similar to the base model exists in a 3-period model.*

Therefore, we conclude that the main results of our paper are robust to multiple periods and limiting our analysis to 2 periods in the main text helps to focus the discussion on the direct implications of reciprocity-based ratings.

A.2.2 Proofs for Extensions Robustness (Intended For Online Appendix)

Proof of Proposition 7

Proof of Proposition 7.
To proof Proposition 7, we show that $\frac{\partial \pi_p}{\partial e} < 0$.

To begin we characterize the actions of users on either side of the platform. Afterwards, we look at the strategy of the profit-maximizing platform.

To begin we characterize when consumers join the platform. The consumers' value from joining the platform is $n_s u_b$, i.e. the consumer gets expected surplus u_b from each interaction with a seller, and there are n_s sellers in total. Consumers' outside option is uniformly distributed on $[0, 1]$, which is why buyer demand is given by $n_b^* = n_s^* u_b$.

Next, we consider the firms. Since firms are ex-ante homogeneous, they have the same ex-ante expected revenue per transaction, π_s. Firms also face the same commission fee, r, which is set by the platform. Additionally, all firms face a marginal cost t of selling on the platform. Therefore the per-transaction profit for firms is $\pi_s(1 - r) - t$. Since all firms face the same cost of entry $\bar{v_s}$, then any firm whose total profits are weakly above the outside option joins the platform. This means that $n_b^*(\pi_s(1 - r) - t) \geq \bar{v_s}$. Because all firms face the same decision, the number of firms to join the platform is either $n_s^* = 0$ or $n_s^* = 1$. If $n_s^* = 0$, there is no activity on the platform and the platform earns zero profits, which is not optimal. We conclude that $n_s^* = 1$. Note that this implies $n_b^* = u_b$.

We now turn our attention to the platform. Because sellers are homogeneous and $n_s^* = 1$, the profit-maximizing platform extracts the highest possible benefit from the royalty fee subject to sellers participating. This implies that the platform sets the optimal r^* such that

$$n_b^*(\pi_s(1 - r^*) - t) = \bar{v_s} \Leftrightarrow r^* = 1 - \frac{\bar{v_s}}{n_b^* \pi_s} - \frac{t}{\pi_s}.$$

Using this, we can simplify the platform's profits to

$$\pi_p = n_b^* \pi_s r^* = n_b^* \pi_s - \bar{v_s} - t n_b^* = n_b^*(\pi_s - t) - \bar{v_s}.$$

Now, we consider the effects of the ratings environment on the profits of the platform. To see this, we need to understand how the platform's profits are affected by changes to effort cost,

$$\frac{\partial \pi_p}{\partial e} = \frac{\partial u_b}{\partial e}(\pi_s - t) + \frac{\partial \pi_s}{\partial e} u_b. \tag{A.11}$$

To understand the platform's strategy, we evaluate $\frac{\partial \pi_p}{\partial e}$. To do so, we first show that

$\frac{\partial u_b}{\partial e} = -\frac{\partial \pi_s}{\partial e}$.

Consider first u_b. This is the transaction benefit of each consumer. Because consumers purchase first or second with equal probability, their ex-ante expected benefit per transaction is $u_b = \frac{1}{2}CS$, where we know from (1.2) that $CS = (1-\delta^*)(1-\gamma)[(1-\kappa)q^L + \frac{e}{\Delta}]$. Thus,

$$\frac{\partial u_b}{\partial e} = \frac{1}{2}\frac{(1-\gamma)(1-\delta^*)}{\Delta} - \frac{1}{2}\frac{\partial \delta^*}{\partial e}[(1-\kappa)q^L + \frac{e}{\Delta}].$$

Next, consider π_s. This is the per transaction profit of firms before taking into account the commission fee of the platform. Ex-ante, this is equivalent to the expected revenue that the firms receives per consumer, i.e.

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

Taking the derivative,

$$\frac{\partial \pi_s}{\partial e} = \frac{1}{2}\frac{\partial \delta^*}{\partial e}[(1-\kappa)q^L + \frac{e}{\Delta}] - \frac{1}{2}\frac{(1-\gamma)(1-\delta^*)}{\Delta}.$$

Thus we conclude that $\frac{\partial u_b}{\partial e} = -\frac{\partial \pi_s}{\partial e}$.

Returning to equation (A.11), if $\frac{\partial \pi_p}{\partial e} < 0$, the platform designs a ratings system such that it minimizes the effort cost associated with rating.

$$\frac{\partial \pi_p}{\partial e} < 0 \Leftrightarrow \frac{\partial u_b}{\partial e}(\pi_s - t) + \frac{\partial \pi_s}{\partial e}u_b < 0 \Leftrightarrow \frac{\frac{\partial u_b}{\partial e}}{u_b} < -\frac{\frac{\partial \pi_s}{\partial e}}{(\pi_s - t)}.$$

From this formulation, we see that the platform's design decision depends on the relative elasticity of consumer transaction surplus and firms transaction revenue. In particular, since $\frac{\partial u_b}{\partial e} = -\frac{\partial \pi_s}{\partial e}$,

$$\frac{\frac{\partial u_b}{\partial e}}{u_b} < \frac{\frac{\partial u_b}{\partial e}}{(\pi_s - t)} \iff (\pi_s - t) < u_b$$

implies that $\frac{\partial \pi_p}{\partial e} < 0$.

We now show when this condition can be satisfied. First, note that $e < e^{cs}$, so $\frac{\partial u_b}{\partial e} > 0$.

We now provide the conditions for which $\frac{\frac{\partial u_b}{\partial e}}{u_b} < \frac{\frac{\partial u_b}{\partial e}}{(\pi_s - t)}$ is satisfied.

$$\pi_s - t < u_b \Leftrightarrow \frac{\gamma q^H + (1-\gamma)q^L - t}{(1-\kappa)q^L + \frac{e}{\Delta}} < (1-\delta^*)(1-\gamma). \tag{A.12}$$

Therefore, when (A.12) holds, we show that platforms favour sellers and minimise the effort cost required to leave a rating.

We conclude that when (A.12) holds platforms will minimize the effort costs required to leave a rating.

Specifically, when $t \geq \frac{q^L((1+\gamma)+\kappa(1-\gamma))\sqrt{4\gamma^2(q^H-q^L)^2+(1+\gamma)^2(q^L(1-\kappa))^2}}{2}$, $e^* = 0$. Otherwise,

$$e =\Delta \frac{q^L(-1+\gamma(2\kappa-1)) + t(1-\gamma) \pm \sqrt{(q^L-t)^2(1+2\gamma+\gamma^2) - 4\gamma^3(q^H-q^L)^2}}{2\gamma}$$

Since $e^* > 0$, we reject the negative and

$$e^* = \Delta \frac{q^L(-1+\gamma(2\kappa-1)) + t(1-\gamma) + \sqrt{(q^L-t)^2(1+2\gamma+\gamma^2) - 4\gamma^3(q^H-q^L)^2}}{2\gamma}$$

We have provided the condition for which platforms are incentivised to minimize effort costs associated with rating. And we have also shown more generally the level of effort that maximises platform's profit when (A.12) holds.

This concludes the proof. □

Proof of Proposition 8

To proof Proposition 8, we first need to show some lemmas. First, Lemma 5 guides when the competitive fringe may be considered to be active in the market. Second, Lemma 6 shows the adjusted pricing strategies of strategic firms. Finally, Lemma 7 characterizes the adjusted consumer's beliefs in the presence of the competitive fringe. We then use these results to prove Proposition 8.

Lemma 5. *The competitive fringe is only relevant whenever $c < q$.*

Proof of Lemma 5.

Whenever $c > q$, consumers receive negative utility from the competitive fringe. Whenever $c = q$, consumers receive 0 utility. Since the strategic firms provide at least non-negative utility, by assumption they would receive the full consumer demand. Therefore, the fringe only captures consumers if $c < q$.

This concludes the proof. □

Lemma 6. *Pricing strategy of the high-quality firm:*

- *Period 1:* $p_1'^H = E[q_1|R_0, p_1^H] - (q - c)$

- *Period 2:* $p_2'^H = E[q_2|R_1] - (q-c)$

Pricing strategy of the low-quality firm:

- *Period 1:*
  - $R_1 = 1$, $p_{1,1}'^L = min\{q^L - (q-c), \kappa q^L - \frac{e}{\Delta}\}$
  - $R_1 = 0$, $p_{1,0}'^L = p_1'^H$
- *Period 2:* $p_2'^L = E[q_2|R_1] - (q-c)$

Proof of Lemma 6.
To prove this, we first discuss the tie breaker rule. Second, we consider the second period as this is straight forward. Then we consider the high-quality firm in the first period, followed by the low-quality firm in the first period. The proofs are similar to that of Lemma 2.

Our tie breaker rule is such that whenever both the strategic firm and the fringe provide the same level of consumer surplus, consumers choose to purchase from the strategic firm. Here, we shall argue why this tie breaker rule is a simplification that yields virtually identical results to a tie breaker rule where consumers randomise between the two firms.

By construction, the fringe sets some fixed price c, while the strategic firm selects prices. Therefore, for any situation where the strategic firm and fringe provide the same level of consumer surplus, the strategic firm can always choose to set prices some small $\epsilon > 0$ lower, such that they obtain the full demand. Therefore, by assuming this tie breaker rule, we are able to simplify our discussion without considering the need for some small ϵ deviation.

We now turn our attention to the discussion of equilibrium prices in the second period. In the second period, consumers are aware that consumption from the fringe will leave them a surplus of $q - c$. This forms the outside option for consumers. Strategic firms therefore have to provide consumers with a surplus of at least $q - c$, doing so will shift all the demand towards the strategic firm. We conclude that the firm will not set prices lower than $E[q_2|R_1] - (q^L - c)$.

We now turn our attention to the high-quality firm in the first period. As in Lemma 2, a high-quality firm wishing to obtain a positive rating must set prices no higher than $\kappa q^H - \frac{e}{\Delta}$. As in the proof of Proposition 6, consumers are aware that consumption of the outside good will provide $q - c$ level of utility. Therefore, firms must set prices no higher than $E[q_1|R_0, p_1] - (q - c)$. Because we show in the proof of Proposition 6 that we must have $min\{\kappa q^H - \frac{e}{\Delta}, E[q_1|R_0, p_1^H]\} = E[q_1|R_0, p_1^H]$ in equilibrium, high-quality firms will continue to receive a good rating at $E[q_1|R_0, p_1] - (q-c)$. It follows that the equilibrium price is $p_1'^H = E[q_1|R_0, p_1^H] - (q-c)$.

We now look at the low-quality firm in period 1. When the low-quality firm prefers to obtain a good rating, it has to set a price no higher than $\kappa q^L - \frac{e}{\Delta}$. With the presence of the competitive fringe, consumers are aware that they would be able to receive at least $q-c$ of surplus. Therefore, in order for the low-quality firm to capture demand and get a good rating, they are unable to command prices higher than $q^L - (q-c)$. Thus, the equilibrium price is $p_{1,1}'^L = min\{c, \kappa q^L - \frac{e}{\Delta}\}$.

For the low-quality firm receiving no rating, it sets prices above $\kappa q^L - \frac{e}{\Delta}$. The profit maximizing firm will set the highest possible price at which consumers buy. For the same argument as in the proof of Lemma 2, low-type firms do not set prices higher than the high-quality firm. Therefore, when low-quality firms receive no rating, it sets prices equal to that of the high-quality firm.

This concludes the proof. □

Lemma 7. *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 7.
The proof is identical to that of Lemma 3. □

Proof of Proposition 8.
We now proof Proposition 8.

First, we discuss the possible cases for how the competitive fringe might affect the behavior of strategic firms. Then we consider these cases one at a time and find the effect that changes in c have on the mixed-strategy of the low-quality firm.

Since we know that the second period prices are equally affected by the competitive fringe in all cases, we turn our attention to the first period. Notice that in period 1, the high-quality firm sets a single price, and the low-quality firm sets two out of three possible prices, i.e. $q^L - (q-c)$, $\kappa q^L - \frac{e}{\Delta}$, or $p_1'^H$. Further, notice that $\overline{p_1^H} > \overline{p_1^L}$ and $E[q_1|p_1'^H] - (q-c) > q^L - (q-c)$.

We know from the proof of Proposition 6 that in equilibrium we have $\overline{p_1^H} > E[q_1|p_1'^H]$, this leaves us with the following possible scenarios:

$$
\begin{aligned}
&1.\overline{p_1^H} > E[q_1|p_1'^H] - (q-c) > \overline{p_1^L} > q^L - (q-c) \\
&2.\overline{p_1^H} > E[q_1|p_1'^H] - (q-c) > q^L - (q-c) \geq \overline{p_1^L}
\end{aligned}
$$

We consider the scenarios individually.

1. $\overline{p_1^H} > E[q_1|p_1'^H] - (q-c) > \overline{p_1^L} > q^L - (q-c)$

Because $\overline{p_1^L} = \kappa q^L - \frac{e}{\Delta} > q^L - (q-c)$, we know from Lemma 6 that the low-quality firm is indifferent between setting prices $p_1'^H$ and $q^L - (q-c)$. When charging $p_1'^H$, the firm gets no rating and charges $q^L - (q-c)$ in period 2. When charging $q^L - (q-c)$, the firm gets a good rating and earns the expected quality of a firm with a good rating in period 2 minus $(q-c)$. This leads to the following condition.

$$
\begin{aligned}
&E[q_1|p_1'^H] - (q-c) + q^L - (q-c) = q^L - (q-c) + E[q_2|R_1] - (q-c) \\
\Leftrightarrow &\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} = \frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)} \\
\Leftrightarrow &\delta^* = 0.5,
\end{aligned}
$$

which is independent of c.

This concludes case 1.

2. $\overline{p_1^H} > E[q_1|p_1'^H] - (q-c) > q^L - (q-c) \geq \overline{p_1^L}$

Because $q^L - (q-c) \geq \overline{p_1^L}$, and we know from Lemma 6 that the low-quality firm is indifferent between setting prices $p_1'^H$ and $\overline{p_1^L} = \kappa q^L - \frac{e}{\Delta}$. To be indifferent between these prices, the following condition must hold. The left-hand side is as in the previous case. On the right-hand side, the firm charges $\kappa q^L - \frac{e}{\Delta}$ in period 1 and obtains a good rating. In period 2, it earns the expected quality of a firm with good rating in period 2 minus $(q-c)$. This leads to the following condition.

$$
\begin{aligned}
&E[q_1|p_1'^H] - (q-c) + q^L - (q-c) = \overline{p_1^L} + E[q_2|R_1] - (q-c) \\
\Leftrightarrow &\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} - q + c = \kappa q^L - \frac{e}{\Delta} + \frac{\gamma(q^H - q^L)}{\gamma + (1-\delta^*)(1-\gamma)}.
\end{aligned}
$$

Using the implicit-function theorem then leads to

$$\frac{\partial \delta^*}{\partial c} = \frac{(\gamma + (1-\delta^*)(1-\gamma))^2(\gamma + \delta^*(1-\gamma))^2}{\gamma(1-\gamma)(q^H - q^L)[(\gamma + (1-\delta^*)(1-\gamma))^2 + (\gamma + \delta^*(1-\gamma))^2]} > 0.$$

This concludes case 2.

We conclude that $\frac{\partial \delta^*}{\partial c} > 0$ if $q^L - (q - c) \geq \kappa q^L - \frac{e}{\Delta}$ and 0 otherwise. □

Proof of Proposition 9

We can proof this proposition using the following two Lemmas

Lemma 8. *Consumers obtain the same signal from a negative rating and no rating.*

Proof of Lemma 8.

This proof holds directly from Restriction 1. Consumers' belief is such that high-quality firms always get a good rating. Hence, on observing no or bad ratings, they would believe this to be obtained by low-quality firms.

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

where δ^* continues to represent the probability that low-quality firms get no rating.

This concludes the proof. □

Lemma 9. *When low-quality firms choose not to receive a good rating, it prefers to receive a bad rating over no rating, if consumers continue to buy at the price level that induces a bad rating.*

Proof of Lemma 9.

First, we discuss the prices set in the second period. Then we look at consumer's beliefs in the first period and the prices set in the first period. We then show that in equilibrium, there exist some price for which low-quality firms receive no rating, and another price where they receive bad ratings.

When looking at the prices set in the second period, notice first that firms of all types will set the highest possible price in the second period. This is equivalent to consumer's expectation in the second period. Since this is only dependent on ratings, on receiving a good rating in the first period, firms set $\frac{\gamma q^H + (1-\delta^*)(1-\gamma)q^L}{\gamma + (1-\delta^*)(1-\gamma)}$; on receiving no rating, they set a price of q^L; on receiving a bad rating, they also set a price of q^L.

Now, consider the possible prices in the first period. In the first period, high-quality firms will set a price $p_1^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_1|R_0, p_1^H]\}$ and receives a good rating. The proof of this is identical to the proof of high-quality firm's prices in Lemma 2.

Low-quality firm's do one of the following: set a price such that it receives a good rating, no rating or a bad rating.

When a low-quality firm receives a good rating, it sets a unique price and this is $p^L_{1,1} = \kappa q^L - \frac{e}{\Delta}$. This proof is identical to that of Lemma 2 when low-quality firms receive a good rating.

We turn our attention to the situation when low-quality firms receive a bad rating. When a low-quality firm receives a bad rating, this only occurs if $[\kappa q^L - p^L_1]\Delta' - e \geq 0$. Further, bad ratings imply that consumers are exhibiting negative reciprocity. Therefore, we define $\Delta' = -\Delta$. And consumer's leave a bad rating whenever $p^L_1 \geq \kappa q^L + \frac{e}{\Delta}$, gaining a positive rating utility from doing so. From Restriction 2, we know that since this price is larger than $\overline{p^L_1}$, consumers expectations are fixed and from Lemma 2, we know that the firm set the highest possible price. Therefore, when obtaining a bad rating, the low-quality firm sets the price $p^L_1 = E[q_1|R_0, p^H_1]$.

Now we turn our attention to the situation when low-quality firms receive no rating. This occurs when $[\kappa q^L - p^L_1]\Delta' - e < 0$. That is the rating utility from leaving a good or bad rating is negative. Recall that the utility from giving no rating is 0.

Suppose that the consumer considers between leaving a good rating and no rating. By reciprocity, the consumer acts with a kindness of their own, implying that $\Delta' = \Delta > 0$. Therefore, no rating only occurs if $[\kappa q^L - p^L_1] < \frac{e}{\Delta}$. In other words, $p^L_1 > \kappa q^L - \frac{e}{\Delta}$.

Now suppose that the consumer considers between leaving a bad rating and no rating. By reciprocity, the consumer acts negatively, implying that $\Delta' = -\Delta < 0$. Therefore, no rating only occurs if $p^L_1 < \kappa q^L + \frac{e}{\Delta}$ and $p^L_1 > \kappa q^L - \frac{e}{\Delta}$ hold together.

Notice that $\kappa q^L - \frac{e}{\Delta} < \kappa q^L + \frac{e}{\Delta}$. This implies that there exist a range of prices, $p^L_1 \in (\kappa q^L - \frac{e}{\Delta}, \kappa q^L + \frac{e}{\Delta})$ such that consumers maximising their utility provide no ratings to low-quality firms.

In equilibrium, the total profit of a low-quality firm obtaining a bad rating is $\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)} + q^L$, while the total profit of obtaining no rating is strictly below $\kappa q^L + \frac{e}{\Delta} + q^L$.

We can now show that low-quality firms set some unique price in equilibrium, when they choose not to obtain a good rating. In equilibrium, when consumers observe price of p^H_1, they anticipate a utility of $\frac{\gamma q^H + \delta^*(1-\gamma)q^L}{\gamma + \delta^*(1-\gamma)}$. If this is above $\kappa q^L + \frac{e}{\Delta}$, low-quality firms maximise their profits by extracting the full surplus from consumers, and obtaining a bad rating. Alternatively, if the anticipated utility is below $\kappa q^L + \frac{e}{\Delta}$, it is not profitable for firms to set $\kappa q^L + \frac{e}{\Delta}$ as this would lead to zero demand. Instead, the low-quality firm sets p_1^H and obtains no rating. Therefore, when choosing not to obtain a good rating, low-quality firms always sets the unique price mimicking price of p_1^H.

We conclude that when we allow for negative ratings, there exists a range of prices for which good, bad, and no ratings may occur. However, in equilibrium, a profit maximising low-quality firm mixes between setting a price that allows it to harvest ratings and price mimicking. This allows it to obtain the extreme ratings of no, or negative rating depending on the cost of leaving a rating.

This concludes the proof. □

Together, Lemma 8 and Lemma 9 show that in the equilibrium, negative ratings will replace no ratings when going from a system of 2 ratings options to 3 ratings options, and consumers continue to buy at the price level that induces a negative rating. This concludes the Proof of Proposition 9.

Proof of Lemma 4.
To complete this proof, first consider the possible prices played by the high-quality firm. This is unchanged and follows directly from the Proof of Lemma 9. Second, we now consider the possible prices played by the low-quality firm.

When receiving a good rating, the low-quality firm sets a unique price and this is given by $p_{1,1}^L = \overline{p_1^L} = \kappa q^L - \frac{e_g}{\Delta}$. This proof is identical to that of Lemma 2 when the low-quality firm receives a good rating.

Next, we turn our attention to the low-quality firm receiving a bad rating. This only occurs if $[\kappa q^L - p_1^L](-\Delta) - e_b > 0$. As in Lemma 9 consumers leaving a bad rating exhibit negative reciprocity, this is why the warm glow is negative, and negative ratings occur if $p_1^L > \kappa q^L + \frac{e_b}{\Delta}$.

Notice that applying the same logic as in Lemma A.2.1, there exists a range of prices such that no rating occurs, and this is $p_1^L \in (\kappa q^L - \frac{e_g}{\Delta}, \kappa q^l + \frac{e_b}{\Delta})$.

This has two implications. First, as e_g is smaller, then from Corollary 1, low-quality firms harvest ratings more.

Second, as e_b is larger, low-quality firms only receive a bad rating if they set a higher price. However, as this runs up against the upper bound of prices that they may set, prices must be less than consumers willingness to pay, firms are unlikely to receive bad ratings.

This concludes the proof. □

Proof of Proposition 10

Proof of Proposition 10.

To begin the proof, we first define the game. This is a two period game comprising of a firm active in both periods, and a different consumer in each period. Before the first period, the firm draws quality q from a uniform distribution on the interval $[0, 1]$, learning its true quality.

In the first period, firms choose any price p_1, $\sigma_1^f : q \in [0, 1] \to p_1 \in \mathbb{R}$. First period consumers learn of this price, forming beliefs over the quality of the firm, $E[q|p_1] : p_1 \in \mathbb{R} \to [0, 1]$. They then make a consumption decision, $\sigma_{1,B}^c : p_1 \in \mathbb{R} \to \{0, 1\}$, where 1 represents the decision to purchase and 0 not. Upon consumption, and payment, consumers learn of the true quality of the product. They may then choose to leave a rating, $\sigma_{1,R}^c : q \times p_1 \in [0, 1] \times \mathbb{R} \to R_1 \in \{0, 1\}$, where 1 represents the decision to rate, and 0 not.

At the end of the first period, first-period consumers leave the market and second-period consumers arrive. The rating R_1 is revealed to both the firm and the second-period consumer. In the second period, the firm sets a price p_2, $\sigma_2^f : q \times R_1 \in [0, 1] \times \{0, 1\} \to p_2 \in \mathbb{R}$. Second-period consumers learn this price, forming beliefs over the quality of the firm, $E[q|p_2, R_1] : p_2 \times R_1 \in \mathbb{R} \times \{0, 1\} \to [0, 1]$. Consumers then make a consumption decision, $\sigma_2^c : p_2 \times R_1 \in \mathbb{R} \times \{0, 1\} \to \{0, 1\}$.

Candidate equilibrium: We now describe the candidate equilibrium: Firms with quality $q \in [\hat{q}, 1]$ obtain a good rating in equilibrium. Firms with quality $q \in [\hat{\hat{q}}, \hat{q})$ choose to participate in ratings harvesting in equilibrium. Firms with quality $q \in [0, \hat{\hat{q}})$ do not obtain a rating in equilibrium. The cutoffs are $\hat{q} = \frac{(2\kappa-1)[4\kappa\frac{e}{\Delta}+1+2\kappa]+2}{4\kappa(1+2\kappa^2-\kappa)}$ and $\hat{\hat{q}} = \frac{(2\kappa-1)(1+4\kappa\frac{e}{\Delta})}{4\kappa(1+2\kappa^2-\kappa)}$.

We first state consumers' equilibrium beliefs in period 1.

$$E[q|p_1] = \begin{cases} \frac{\int_{\hat{q}}^1 q dq + \int_0^{\hat{\hat{q}}} q dq}{\int_{\hat{q}}^1 1 dq + \int_0^{\hat{\hat{q}}} 1 dq} = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})} & \text{if } p_1 \geq \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})} \\ q & \text{if } p_1 = \kappa q - \frac{e}{\Delta} < \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})} \\ 0 & \text{otherwise} \end{cases}$$

Note that $\hat{q}$ is such that $p_1(\hat{q}) = \kappa\hat{q} - \frac{e}{\Delta} = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$, and the expectation is well defined.

Next, we state consumer's strategies in period 1.

$$\sigma^c_{1,B} = \begin{cases} 1 & \text{if } p_1 \leq E[q|p_1] \\ 0 & \text{if } p_1 > E[q|p_1] \end{cases}$$

$$\sigma^c_{1,R} = \begin{cases} 1 & \text{if } p_1 \leq \kappa q - \frac{e}{\Delta}, \ \forall q \\ 0 & \text{if } p_1 > \kappa q - \frac{e}{\Delta}, \ \forall q \end{cases}$$

Next, we state the firms' strategies in period 1.

$$\sigma^f_1 = \begin{cases} \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})} & \text{if } q \in [\hat{q}, 1] \\ \kappa q - \frac{e}{\Delta} & \text{if } q \in [\hat{\hat{q}}, \hat{q}) \\ \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})} & \text{if } q \in [0, \hat{\hat{q}}) \end{cases}$$

We now turn our attention to period 2, stating consumers' equilibrium beliefs, their strategies, and the firms' strategies.

$$E[q|p_2, R_1] = \begin{cases} \frac{\int_{\hat{\hat{q}}}^1 q dq}{\int_{\hat{\hat{q}}}^1 1 dq} = \frac{1+\hat{\hat{q}}}{2} & \text{if } R_1 = 1, \, \forall p_2 \\ \frac{\int_0^{\hat{\hat{q}}} q dq}{\int_0^{\hat{\hat{q}}} 1 dq} = \frac{\hat{\hat{q}}}{2} & \text{if } R_1 = 0, \, \forall p_2 \end{cases}$$

$$\sigma^c_{2,B} = \begin{cases} 1 & \text{if } p_2 \leq E[q|p_2, R_1] \\ 0 & \text{if } p_2 > E[q|p_2, R_1] \end{cases}$$

$$\sigma^f_1 = \begin{cases} \frac{\int_{\hat{\hat{q}}}^1 q dq}{\int_{\hat{\hat{q}}}^1 1 dq} & \text{if } R_1 = 1, \forall q \\ \frac{\int_0^{\hat{\hat{q}}} q dq}{\int_0^{\hat{\hat{q}}} 1 dq} & \text{if } R_1 = 0, \forall q \end{cases}$$

Note that $E[q|p_2, R_1 = 1] = E[q|p_2, R_1 = 0] + \frac{1}{2}$. This implies that obtaining a good rating allows firms to set a higher price in the second period.

Optimality: We now show that consumers make rational choices given specified beliefs. In the first period, given their information set, i.e. p_1, consumers only choose to consume if their expected utility is weakly positive. Moreover, they only choose to leave ratings if their utility from doing so, $v_t(q)$ is weakly positive. Period 2 consumers observe p_2 and R_1, making a consumption decision only if their expected utility conditional on observing p_2 and R_1 is weakly positive. Therefore, we conclude that the candidate equilibrium specifies optimal strategies for consumers given their beliefs.

We now show that firms make rational choices given specified beliefs. To do so, we first solve for $\hat{q}$ and $\hat{\hat{q}}$, providing the conditions over parameters κ that characterises the equilibrium. Next, we show that there is no profitable deviation from the firm's specified strategy.

To solve for $\hat{q}$ and $\hat{\hat{q}}$, we use two conditions. First, consumers who observe $\hat{q}$ are indifferent between rating or not, i.e. $\kappa\hat{q} - \frac{e}{\Delta} = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$. Second, firm $\hat{\hat{q}}$ is indifferent between rating harvesting and price mimicking, that is $\kappa\hat{\hat{q}} - \frac{e}{\Delta} + \frac{1}{2} = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$. Combining these two equations, $\hat{q} - \hat{\hat{q}} = \frac{1}{2\kappa}$. This implies that the range of firms participating in ratings harvesting is dependent on only κ and decreases in κ. And we have $\hat{q} = \frac{(2\kappa-1)[4\kappa\frac{e}{\Delta}+1+2\kappa]+2}{4\kappa(1+2\kappa^2-\kappa)}$, $\hat{\hat{q}} = \frac{(2\kappa-1)(1+4\kappa\frac{e}{\Delta})}{4\kappa(1+2\kappa^2-\kappa)}$. Because $\hat{\hat{q}} \geq 0$, it is sufficient to have $\kappa \geq \frac{1}{2}$.

We now show that firms have no profitable deviation. In the second period, firms charge prices equal to their expected quality conditional on their rating and earns strictly positive profits. Thus, deviating to a larger price reduces demand to zero and is therefore not profitable. Deviating to a lower price reduces margins without increasing demand, and is therefore also not profitable. We conclude that firms have no profitable deviation in the second period.

Next, we show that firms have no profitable deviation in the first period. Note first that any deviation to a price $p_1 > \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$ induces zero demand and profits and is therefore not a profitable deviation.

Next, consider firms of quality $q \in [\hat{q}, 1]$. These firms get a rating and therefore charge the large price in period 2. We have shown above that charging a larger price above $\frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$ is not a profitable deviation. Additionally, deviating to a strictly lower price only reduces margins in period 1 without increasing demand in period 1 or changing the rating in period 2. Thus, deviating to a lower price in period 1 is not a profitable deviation. We conclude that firms of $q \in [\hat{q}, 1]$ do not have a profitable deviation in period 1.

Next, consider firms of quality $q \in [\hat{\hat{q}}, \hat{q})$. In the candidate equilibrium, they set a price $p_1 = \kappa q - \frac{e}{\Delta}$ and receive a good rating. To start, consider an upward deviation. We have shown above that the firm does not want to deviate to a price strictly above $\frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$. Because any deviation to a price above $\kappa q - \frac{e}{\Delta}$ leads consumers to no longer rate the firm, the most profitable deviation to a larger price is therefore to the price $\frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$, i.e. the firm would deviate to price mimicking. Because $\hat{q}$ is such that $\kappa\hat{q} - \frac{e}{\Delta} = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$, and because $q < \hat{q}$, this is indeed a deviation to a larger price. But because $\hat{\hat{q}}$ is such that $\hat{\hat{q}}$ is indifferent between rating harvesting and price mimicking, i.e. $\kappa\hat{\hat{q}} - \frac{e}{\Delta} + \frac{1}{2} = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{1-\hat{q}+\hat{\hat{q}}}$, and because $\kappa q - \frac{e}{\Delta} + \frac{1}{2} > \kappa\hat{\hat{q}} - \frac{e}{\Delta} + \frac{1}{2}$, firms $q \in [\hat{\hat{q}}, \hat{q})$ earn strictly larger profits from rating harvesting than price mimicking. Thus, firms $q \in [\hat{\hat{q}}, \hat{q})$ do not have a profitable deviation to a larger price. Additionally, if any firm $q \in [\hat{\hat{q}}, \hat{q})$ deviates to a strictly lower first-period price that $\kappa q - \frac{e}{\Delta}$, the firm earns a strictly lower margin in period 1 without earning a larger demand in period 1 or a larger profit in period 2. We therefore conclude that no firm $q \in [\hat{\hat{q}}, \hat{q})$ has a profitable deviation in period 1.

Next, consider firms $q \in [0, \hat{\hat{q}})$, recall that in equilibrium they set a first-period price

$\frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$ and receive no rating. We have therefore shown above that they have no profitable upward deviation. To see that there is no downward deviation, notice that the firms currently receive no rating. If they lower prices slightly, above $\kappa q - \frac{e}{\Delta}$, they will still receive no rating, thus period 2 profits are unaffected and period 1 profits decrease, which does not constitute a profitable deviation. However, if they reduce prices to weakly below $\kappa q - \frac{e}{\Delta}$, then these firms receive a good rating, and gain a profit of $\frac{1}{2}$ in the second period. But since $q < \hat{\hat{q}}$, we know that $\kappa q - \frac{e}{\Delta} + \frac{1}{2} < \kappa \hat{\hat{q}} - \frac{e}{\Delta} + \frac{1}{2} = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$ and these deviations are also not profitable. We conclude that no firm $q \in [0, \hat{\hat{q}})$ has a profitable deviation in period 1.

Overall, we conclude that not firm has a profitable deviation and that the firms' strategies in the candidate equilibrium are optimal given their beliefs.

Correct beliefs on the equilibrium path: We now show that the beliefs in the candidate equilibrium follow Bayes rule on the path of play.

In the second period, consumers beliefs are such that $E[q|p_2, R_1] = E[q|R_1]$. Consumers correctly believe that only firms of quality $q \geq \hat{\hat{q}}$ receive a good rating in period 1. Similarly, consumers correctly believe that firms of quality $q < \hat{\hat{q}}$. Therefore their beliefs are consistent with Bayes rule.

In the first period, consumer beliefs $[q|p_1]$ only depend on the first-period price. Consumers who observe a price $p_1 = \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$ correctly take into account that only firms $q \in [\hat{q}, 1]$ and $q \in [0, \hat{\hat{q}})$ set this price on the path of play, and these beliefs are correct. Next, note that the remaining firms $q \in [\hat{\hat{q}}, \hat{q})$ charge a price $p_1 = \kappa q - \frac{e}{\Delta} < \frac{1-\hat{q}^2+\hat{\hat{q}}^2}{2(1-\hat{q}+\hat{\hat{q}})}$. Thus, each of these firms charges a unique price that no other firm charges in equilibrium, and consumers have correct beliefs. We therefore conclude that consumers' beliefs are consistent with Bayes rule on the path of play. Since firms have full information, also their beliefs are consistent with Bayes rule. We conclude that all players' beliefs are consistent with Bayes rule on the path of play.

Finally, we show that there indeed exist parameter ranges for which $\hat{q} \in (0, 1)$ and $\hat{\ } \in \hat{(}0, 1q)$. Because $\hat{\hat{q}} < \hat{q}$, it suffices to show that $0 < \hat{\hat{q}}$ and $\hat{q} < 1$. Reformulating $\hat{\hat{q}}$ to $\hat{\hat{q}} = \frac{(2\kappa-1)(1+4\kappa\frac{e}{\Delta})}{4\kappa(1+\kappa(2\kappa-1))}$, we can see that $\kappa \geq \frac{1}{2}$ implies $0 < \hat{\hat{q}}$. Simplifying $\hat{q} < 1$ using $\kappa \geq \frac{1}{2}$ leads to $0 \leq 4\kappa^2 - (2\kappa - 1) - 4\kappa\frac{e}{\Delta}$. Note that for $\kappa \to \frac{1}{2}$, this conditions becomes $0 \leq 2(\frac{1}{2} - \frac{e}{\Delta})$, which holds if $\frac{e}{\Delta}$ is sufficiently small. We therefore conclude that both conditions can be satisfied at the same time. Thus, we conclude that the equilibrium indeed exist.

A larger e makes ratings more informative. Finally, we can show that the informativeness of ratings increases in e. Note that when e increases, more firms participate in price mimicking, $\frac{\partial \hat{q}}{\partial e} = \frac{2\kappa-1}{\Delta(1-\kappa+2\kappa^2)} > 0$. This implies that conditional on observing a good rating, consumers believe that the rating comes from firms of higher quality, as the cut off for firms receiving a good rating has increased. Thus ratings are more informative of quality.

We have shown that there exists a pure-strategy perfect-Bayesian equilibrium with 3 groups of firms, $q \in [\hat{q}, 1]$ obtaining a good rating, without influencing ratings, $q \in [\hat{\hat{q}}, \hat{q})$ choosing to participate in ratings harvesting, and $q \in [0, \hat{\hat{q}})$ which do not obtain a rating in equilibrium. And also in this equilibrium, the informativeness of ratings increases in e.

□

Proof of Proposition 11

To show Proposition 11, we prove a more general proposition, Proposition 12. Further, recall that to save on notation we omit the 0 rising from $t = 0$.

Proposition 12. *All perfect Bayesian equilibria satisfy the following.*

1. *High-quality firms receive a good rating with probability 1 and charge* $p_t^H = E[q_t|R_{t-1}, p_t^H]$, $t \in \{1, 2\}$.
2. *Low-quality firms randomize their strategy in period* $t \in \{1, 2\}$.
   a. *They charge* $\overline{p_t^L} = \kappa q^L - \frac{e}{\Delta}$ *and obtain a good rating with probability* $1 - \delta_t^*$.
   b. *They charge* p_t^H *and obtain no rating with probability* δ_t^*.
3. *Firms set prices equal to expected quality conditional on ratings in the last period.*
4. *Consumer beliefs given equilibrium prices are given by:*
   a. *In period 1,* $E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)} & \textit{if } p_1 > \overline{p_1^L} \\ q^L & \textit{if } p_1 \le \overline{p_1^L} \end{cases}$
   b. *In period 2,* $E[q_2|R_1, p_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} & \textit{if } p_2 > \overline{p_2^L} \textit{ and } R_1 = 1 \\ q^L & \textit{for any other } p_2, R_1 \textit{ combination} \end{cases}$
   c. *In period 3,* $E[q_3|R_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)} & \textit{if } R_2 = \{11\} \\ q^L & \textit{for any other } R_2 \end{cases}$

Furthermore, we show that $\delta_1^* \in (0, 1)$ *and* $\delta_2^* \in (0, 1)$ *is an equilibrium if* $(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\delta_1^*)(1-\gamma)(q^H-q^L)}{\gamma+(1-\delta_1^*)(1-\gamma)}$*, and* $\delta_1^* \in (0, 1)$ *and* $\delta_2^* = 1$ *is an equilibrium if* $(1-\kappa)q^L + \frac{e}{\Delta} \ge \frac{(1-\delta_1^*)(1-\gamma)(q^H-q^L)}{\gamma+(1-\delta_1^*)(1-\gamma)}$*. These equilibria exist if* $(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\gamma)(q^H-q^L)}{2-\delta_2^*}$*,* $\gamma > \frac{1}{3}$*, and*

$\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}, \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}\}$.

We proceed as follows. Before proving Proposition 12, we show two Lemmas. First, we use our restrictions to characterize firms' pricing strategies. Second, we pin down equilibrium beliefs. Finally, we use these results to show Proposition 12. Most of the arguments used in proof are similar to the ones from Proposition 1, which is why we briefly sketch them here.

We begin by pinning down the equilibrium prices in period 1 and 2 in Lemma 10.

Lemma 10. *In equilibrium, firms play the following prices in period 1 and 2 with weakly positive probability.*

- *High-quality firm:* $p_t^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\}$
- *Low-quality firm:*

$$p_{1,1}^L = p_{2,11}^L = \kappa q^L - \frac{e}{\Delta}$$

$$p_{1,0}^L = p_{2,10}^L = E[q_t|R_{t-1}, p_t^H]$$

$$p_{2,00}^L = p_{2,01}^L = q^L$$

Proof of Lemma 10.
The proof of Lemma 10 is similar to the proof of Lemma 2.

The key difference lies in the possible rating histories. Given the extension to 3 periods, the possible rating histories are now $R_0 \in \{0\}, R_1 \in \{0,1\}, R_2 \in \{00, 01, 10, 11\}$. As in the proof of Lemma 2, we will consider the pricing strategy of the high-quality firm followed by low-quality firm.

From Restriction 1, a high-quality firm wants to set prices which allows it to get a good rating and therefore in equilibrium her rating's history are $R_0 = \{0\}, R_1 = \{1\}$ and $R_2 = \{11\}$. We show that on the equilibrium path $p_t^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\}$ for $t \in \{1,2\}$.

To show that the high-quality firm sets a unique price in each period $t \in \{1,2\}$ with probability 1, we follow the proof set out in Lemma 2. Suppose towards a contradiction that the high-quality firm sets more than one price with positive probability. Without loss of generality, suppose that the high-quality firm sets a continuous distribution of prices in either period, i.e. it charges prices in some interval $p_t \in [p_t', p_t'']$ such that $p_t'' > p_t'$, and the high-quality firm receives a good rating with probability 1 for all $p_t \in [p_t', p_t'']$. Note that when the high-quality firm gets a rating for all $p_t \in [p_t', p_t'']$, consumers purchase products with probability 1. Notice that for any price $\hat{p}_t > p_t$ such that $\hat{p}_t \in [p_t', p_t'']$, we have $\pi_t^H(\hat{p}_t) > \pi_t^H(p_t)$. The reason is that both prices induce the same demand in period t, and the same rating and therefore the same continuation profits. Thus, the firm can strictly increase profits by shifting all probability mass from $[p_t', p_t'']$ to p_t''. This contradicts the assumption that the high-quality firm sets a uncountably infinite prices with positive probability. Similarly, the firm will not set countably finite or infinite prices that induce the same rating. Therefore, we conclude that the high-quality firm sets a unique price in each period t on the equilibrium path with probability 1. We denote this as p_t^H.

Next, we show that there exist an upper bound on prices, $\overline{p_t^j}$ for $j \in \{L, H\}$ such that j-quality firm receives a positive rating. In order for a firm to induce a positive rating, the rating utility must be positive. Therefore,

$$[\kappa q^j - p_t^j]\Delta - e \geq 0 \iff p_t^j \leq \overline{p_t^j} \equiv \kappa q^j - \frac{e}{\Delta}$$

Finally, consider that prices are bound by consumer's beliefs, $E[q_t|R_{t-1}, p_t^H] < \overline{p_t^H}$. Under such scenarios, by Restriction 1, high-quality firms prefer obtaining a good rating. This can only be achieved if consumers buy. Therefore, p_t^H is bound by $E[q_t|R_{t-1}, p_t^H]$.

We now show that $p_t^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\}$ with probability 1. As with Lemma 2, note that $\overline{p_t^H} > \overline{p_t^L}$ and $q^L > \overline{p_t^L}$. Thus, because for equilibrium expectations we have $E[q_t|R_{t-1}, p_t^H] > q^L$, the high-quality firm sets equilibrium prices strictly larger than $\overline{p_t^L}$. Applying Restriction, 2, consumers have the same beliefs for all prices strictly above $\overline{p_t^L}$ in each period t. Since $p_t^H > \overline{p_t^L}$, these beliefs are given by $E[q_t|R_{t-1}, p_t^H]$, the correct equilibrium beliefs. Further, since consumers have the same beliefs for all prices above $\overline{p_t^L}$ in each period, the high-quality firm optimally sets the highest possible price at which consumers purchase and rate with probability 1. Hence, $p_t^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\}$. We conclude that high-quality firms set this unique price, $p_t^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\}$, with probability 1.

We turn our attention to low-quality firms. First, we show that the price which the low-quality firm sets when it receive a good rating in any period, $t \in \{1, 2\}$, is unique. And this price is $p_{1,1}^L = p_{1,11}^L = \kappa q^L - \frac{e}{\Delta}$. The argument here is the same as the argument used for when high-quality firms receive a good rating. The low-quality firm receives a good rating at any price weakly below $\overline{p_t^L} = \kappa q^L - \frac{e}{\Delta}$, and $\overline{p_t^L} < q^L$. As consumers beliefs are weakly above q^L, they buy at any price weakly below $\overline{p_t^L}$. Since demand and ratings are the same for all prices weakly below $\overline{p_t^L}$, a low-quality firm obtaining a positive rating optimally sets $\overline{p_t^L}$ with probability 1. We can conclude that low-quality firms who obtain a rating in period t set $\overline{p_t^L}$ with probability 1.

Next, we consider the situation where the low-quality firm receives no rating for the first time. The related ratings history are $R_1 = \{0\}, R_2 = \{10\}$. We show that $p^L_{1,0} = E[q_1|R_0, p^H_1]$ and $p^L_{2,10} = E[q_2|R_1, p^H_2]$. We have shown above that all prices above $\overline{p^L_t}$ induces the belief $E[q_t|R_{t-1}, p^H_t]$ - as a result of Restriction 2. Therefore, a low-quality firm receiving no rating optimally sets the highest possible price, $E[q_t|R_{t-1}, p^H_t]$ with probability 1.

Finally, consider the situation where the low-quality firm already has a history of receiving no rating. The related ratings history is $R_2 = \{00, 01\}$. Since by Restriction 1 only low-quality firms receive no rating on the equilibrium path, consumer's belief on observing any history of no rating is that the firm is of a low-quality. Hence $p^L_{2,00} = p^L_{2,01} = E[q_2|R_1 = \{0\}] = q^L$.

This concludes the proof.

□

Lemma 11. *In the first period, consumer's beliefs for each equilibrium price p_1 is given by*

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)} & \text{if } p_1 > \overline{p^L_1} \\ q^L & \text{if } p_1 \leq \overline{p^L_1} \end{cases},$$

in the second period,

$$E[q_2|R_1, p_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} & \text{if } p_2 > \overline{p^L_2} \text{ and } R_1 = 1 \\ q^L & \text{for any other } p_2,\ R_1 \text{ combination} \end{cases},$$

and in the third period,

$$E[q_3|R_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)} & \text{if } R_2 = \{11\} \\ q^L & \text{for any other } R_2 \end{cases}.$$

Proof of Lemma 11.
As with Lemma 3, we prove this lemma by constructing expected quality using Bayes rule. We start by considering the third period, followed by the second then the first.

We begin with the third period. In the third period, consumers are aware of historical ratings, R_2 and current prices p_3. Given that this is the final period of the game, new ratings are not useful for firms - as there is no future period to signal to. By Restriction 2, this implies that the expected quality in period 3 is independent of prices. Thus, the expected quality in period 3 is independent of prices and only depends on the rating history R_2. As a result, firms set the highest possible price and extract the full consumer surplus.

Note that the possible ratings history are $R_2 \in \{00, 01, 10, 11\}$. When consumers observe a ratings history of $\{00, 01, 10\}$, they expect that the firm is a low-quality firm (Restriction 1).

Using Bayes rule, we pin down consumers' expectations on observing $R_2 = \{11\}$ in period 3. In equilibrium, low-quality firms receive no rating with some probability δ_t^* in period $t \in \{1, 2\}$. Hence, consumers observe a good rating with probability $\gamma + (1 - \delta_1^*)(1 - \delta_2^*)(1 - \gamma)$ - and know that the probability of a high-quality firm is γ. Hence, $E[q_3|R_2 = 11] = \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$.

We conclude that

$$E[q_3|R_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)} & \text{if } R_2 = \{11\} \\ q^L & \text{for any other } R_2 \end{cases}.$$

Next, we consider the second period. In the second period, consumers observe $R_1 \in \{0, 1\}$ and the price p_2. As with the third period, on observing $R_1 = 0$ Restriction 1 implies that the firm is of a low-quality.

Using Bayes rule, we pin down consumers' expectations in equilibrium in period 2, conditional on $R_1 = 1$. We distinguish between two cases: the low-quality firm choosing between a good or no rating. When the low-quality firm receives no rating, it sets $p_2^L = E[q_2|R_1 = 1, p_2^H]$ and when it receives a good rating, $p_2^L = \overline{p_2^L} = \kappa q^L - \frac{e}{\Delta}$. Note that $p_2^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_2|R_1, p_2^H]\}$.

We now start with the first case, i.e. the low-quality firm sets $p_2^L = \overline{p_2^L} = \kappa q^L - \frac{e}{\Delta}$ and gets a good rating. Since $\kappa q^H - \frac{e}{\Delta} > \kappa q^L - \frac{e}{\Delta}$ and $E[q_2|R_1 = 1, p_2^H] \geq q^L > \kappa q^L - \frac{e}{\Delta}$, in the equilibrium, consumers who observe $\kappa q^L - \frac{e}{\Delta}$ believe the firm is a low-quality firm, i.e. $E[q_2|R_1 = 1, p_2 = \kappa q^L - \frac{e}{\Delta}] = q^L$.

Next, consider the second case where the low-quality firm sets, by Lemma 10, $p_2^L = E[q_2|R_1 = 1, p_2^H]$ and gets no rating. We distinguish two scenarios. First, suppose $p_2^H = min\{E[q_2|R_1 = 1, p_2^H], \kappa q^H - \frac{e}{\Delta}\} = E[q_2|R_1 = 1, p_2^H]$. By Lemma 10, $p_{2,10}^L = E[q_2|R_1 = 1, p_2^H] = p_2^H$. On observing this price level, Bayes rule implies $E[q_2|R_1 = 1, p_2^H] = \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}$ in equilibrium. Second, consider the scenario where $p_2^H = min\{E[q_2|R_1 = 1, p_2^H], \kappa q^H - \frac{e}{\Delta}\} = \kappa q^H - \frac{e}{\Delta}$. Since $p_2^L = E[q_2|R_1, p_2^H] \neq \kappa q^H - \frac{e}{\Delta} = p_2^H$, consumers who observe p_2^H believe $E[q_2|R_1, p_2^H] = q^H$. Because $p_2^L = E[q_2|R_1, p_2^H] = q_H > \overline{p_2^L}$ and $p_2^H = \kappa q^H - \frac{e}{\Delta} > \overline{p_2^L}$, Restriction 2 implies that both p_2^L and p_2^H induce the same equilibrium beliefs. But then we must have $\delta_2^* = 1$ in equilibrium. Note that beliefs $E[q_2|R_1, p_2^H] = q_H$ are then a special case of $E[q_2|R_1 = 1, p_2^H] = \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}$ for

$\delta_2^* = 1$.

This concludes the second case.

We conclude that $E[q_2|R_1 = 1, p_2^H] = \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}$.

The final stage is to consider period 1. The proof for this is identical to proof in Lemma 11.

This concludes the proof.

□

With Lemma 10 and 11, we can now prove Proposition 12.

Proof of Proposition 12.
We show that all equilibria satisfying our equilibrium restrictions exhibits similar characteristics as those in the main section of the paper. To do so, we show that a perfect Bayesian equilibrium exists, and that all equilibria satisfying our equilibrium restrictions satisfy similar properties as those in the main body. Specifically,

1. High-quality firms receive a good rating with probability 1 and charge $p_t^H = E[q_t|R_{t-1}, p_t^H]$, $t \in \{1, 2\}$.
2. Low-quality firms randomize their strategy in period $t \in \{1, 2\}$.
   a. They charge $\overline{p_t^L} = \kappa q^L - \frac{e}{\Delta}$ and obtain a good rating with probability $1 - \delta_t^*$.
   b. They charge p_t^H and obtain no rating with probability δ_t^*.
3. Firms set prices equal to expected quality conditional on ratings in the last period.
4. Consumer beliefs given equilibrium prices are given by:
   a. In period 1, $E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)} & \text{if } p_1 > \overline{p_1^L} \\ q^L & \text{if } p_1 \leq \overline{p_1^L} \end{cases}$
   b. In period 2, $E[q_2|R_1, p_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} & \text{if } p_2 > \overline{p_2^L} \text{ and } R_1 = 1 \\ q^L & \text{for any other } p_2,\ R_1 \text{ combination} \end{cases}$
   c. In period 3, $E[q_3|R_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)} & \text{if } R_2 = \{11\} \\ q^L & \text{for any other } R_2 \end{cases}$

Furthermore, we show that $\delta_1^* \in (0, 1)$ and $\delta_2^* \in (0, 1)$ is an equilibrium if $(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\delta_1^*)(1-\gamma)(q^H - q^L)}{\gamma + (1-\delta_1^*)(1-\gamma)}$, and $\delta_1^* \in (0, 1)$ and $\delta_2^* = 1$ is an equilibrium if $(1-\kappa)q^L + \frac{e}{\Delta} \geq$

$\frac{(1-\delta_1^*)(1-\gamma)(q^H-q^L)}{\gamma+(1-\delta_1^*)(1-\gamma)}$. These equilibria exist if $(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\gamma)(q^H-q^L)}{2-\delta_2^*}$, $\gamma > \frac{1}{3}$, and $\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H+\delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}, \frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}\}$.

From Lemmas 10 and 11, we have shown statement 4 and the low-quality firm's prices in statement 2. What remains, is to show statements 1 and 3, the mixed-strategy in statement 2, and to show the existence of the equilibrium.

We prove statement 3. In period 3, firms are no longer incentivized by future ratings. We have also shown in Lemma 10 that by Restriction 2, in the final period consumers' beliefs are only dependent on past ratings and therefore independent of the price they observe in period 3. Firms set prices in period 3 equal to the expected quality conditional on past ratings. This concludes the proof of statement 3.

Next, we prove statement 1. From Lemma 10, we have shown that $p_t^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\}$. To show that $min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\} = E[q_t|R_{t-1}, p_t^H]$, suppose towards a contradiction that $min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\} = \kappa q^H - \frac{e}{\Delta}$. We do this in two parts, for period 2 then period 1.

Note first that in period 2, the low-quality firm with no rating in period 1 always sets the price q^L in periods 2 and 3. Thus, we only consider histories after which the low-quality firm received a good rating in period 1. We know from Lemma 10 that $p_{2,10}^L = E[q_2|R_1, p_2^H]$. Thus, in the candidate equilibrium we have $p_2^H \neq p_{2,10}^L$ and only high-quality firms set $p_2^H = \kappa q^H - \frac{e}{\Delta}$, which is why consumers believe that $E[q_2|R_1 = 1, p_2^H] = q^H$. Given Restriction 2, for any price above $\kappa q^L - \frac{e}{\Delta}$ consumers beliefs are the same. Hence $p_{2,10}^L = E[q_2|R_1 = 1, p_2^H] = q^H$. These beliefs are only correct in equilibrium if $\delta_2^* = 0$. But we cannot have $\delta_2^* = 1$ in equilibrium. To see why, note that in period 2 in the candidate equilibrium the low-quality firms charges a low price and receives a good rating, earning in periods 2 and 3 together $\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$. Deviating by setting a high price q_H gives no rating, but earns in periods 2 and 3 the profit $q^H + q^L$. Since the firm earns the same profits in either case in period 1, the deviation is profitable if it increases profits in periods 2 and 3. Because $q^L > \kappa q^L - \frac{e}{\Delta}$ and $q^H > \frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$, this deviation is profitable, contradicting $\delta_2^* = 0$ and that $p_2^H = \kappa q^H - \frac{e}{\Delta}$.

We conclude that $p_2^H = E[q_2|R_1, p_2^H]$.

In period 1, we know from Lemma 10 that $p_{1,0}^L = E[q_1|R_0, p_1^H]$. In the candidate equilibrium, we have $p_1^H \neq p_{1,0}^L$ and only high-quality firms charge $p_1^H = \kappa q^H - \frac{e}{\Delta}$, which is why consumers believe that $E[q_1|R_0, p_1^H] = q^H$. Given Restriction 2, for any price above $\kappa q^L - \frac{e}{\Delta}$ consumers beliefs are the same. Hence $p_{1,0}^L = E[q_1|R_0, p_1^H] = q^H$. These beliefs are only correct in equilibrium if $\delta_1^* = 0$. We now show that we cannot have $\delta_1^* = 0$ in equilibrium. To see why, note that in period 1 in the candidate equilibrium, the low-quality firm sets a low price, receives a good rating and earns total expected profits $\kappa q^L - \frac{e}{\Delta} + (1-\delta_2^*)(\kappa q^L - \frac{e}{\Delta} + E[q_3|R_2 = 11]) + \delta_2^*(E[q_2|R_1 = 1, p_2^H] + q^L)$. If the firm deviates in period 1 to price q_H and charges q_L in subsequent periods, it will earn $q^H + q^L + q^L$. Because $q^L > \kappa q^L - \frac{e}{\Delta}$, $q^L > (1-\delta_2^*)(\kappa q^L - \frac{e}{\Delta}) + \delta_2^*(q^L)$, and $q^H > (1-\delta_2^*)(E[q_3|R_2 = 11]) + \delta_2^*(E[q_2|R_1 = 1, p_2^H])$, this deviation is profitable, contradicting $\delta_1^* = 0$ and that $p_1^H = \kappa q^H - \frac{e}{\Delta}$.

We conclude that $p_1^H = E[q_1|R_0, p_1^H]$.

Overall, this prove statement 1, i.e. $p_t^H = min\{\kappa q^H - \frac{e}{\Delta}, E[q_t|R_{t-1}, p_t^H]\} = E[q_t|R_{t-1}, p_t^H]$.

We have now shown that statements 1,3 and 4 hold in equilibria that satisfy our restrictions. We continue to show statement 2 by characterizing the (mixed) strategies of low-quality firms.

We now characterize the low-quality firms' mixed-strategy in equilibrium in period 2, i.e. δ_2^*. Recall that δ_2^* is the probability of setting a high price of p_2^H that leads to a no rating. Reversely, $1-\delta_2^*$ is the probability of setting a low price of $\kappa q^L - \frac{e}{\Delta}$ that leads to a good rating. Note that low-quality firms only mix prices after histories where they received a good rating in period 1. This is because consumers' beliefs are such that on observing a period of no rating, they believe the firm to be a low-quality firm.

After a history of a good rating in period 1, when the low-quality firm sets a price of $\kappa q^L - \frac{e}{\Delta}$ in period 2, it obtains a good rating and earns in periods 2 and 3

$$\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}. \tag{A.13}$$

This is strictly increasing in δ_2^*.

Alternatively, if the low-quality firm sets the high price p_2^H it obtains no rating in period 2 and earns in periods 2 and 3

$$\frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} + q^L. \tag{A.14}$$

This is strictly decreasing in δ_2^*.

We now show that $\delta_2^* = 1$ is only possible if no mixed-strategy exists in period 2. First, consider that $\delta_2^* = 1$. Then (A.13) and (A.14) become $\kappa q^L - \frac{e}{\Delta} + q^H$ and $\frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)} + q^L$ respectively. In order for $\delta_2^* = 1$ to be optimal, we must have that at $\delta_2^* = 1$, (A.14) is weakly larger than (A.13), i.e.

$$\begin{aligned}\frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)} + q^L \geq \kappa q^L - \frac{e}{\Delta} + q^H \Leftrightarrow \\ (1-\kappa)q^L + \frac{e}{\Delta} \geq \frac{(1-\delta_1^*)(1-\gamma)(q^H - q^L)}{\gamma + (1-\delta_1^*)(1-\gamma)}\end{aligned} \tag{A.15}$$

Observe that (A.13) is strictly increasing in δ_2^* and (A.14) is strictly decreasing in δ_2^*, which is why (A.15) implies that (A.14) is larger than (A.13) for all $\delta_2^* \leq 1$. Therefore, if (A.15) is met, $\delta_2^* = 1$ is an equilibrium and no mixed-strategy equilibrium exists in period 2.

We now show that $\delta_2^* = 0$ is not a possible equilibrium in period 2. Suppose towards a contradiction that $\delta_2^* = 0$. Then (A.13) and (A.14) become $\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)}$ and $q^H + q^L$ respectively. Since $q^H > \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)}$ and $q^L > \kappa q^L - \frac{e}{\Delta}$, the firm is strictly better of by deviating, setting a high price and receiving no rating, which contradicts $\delta_2^* = 0$. We conclude that $\delta_2^* = 0$ cannot be an equilibrium.

We now show that for

$$(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\delta_1^*)(1-\gamma)(q^H - q^L)}{\gamma + (1-\delta_1^*)(1-\gamma)}, \tag{A.16}$$

there is a unique $\delta_2^* \in (0, 1)$ that characterizes the low-quality firms mixed-strategy in period 2.

Recall that a mixed-strategy equilibrium only exists when (A.13) and (A.14) are equal. Observe that for $\delta_2^* = 1$, (A.14) is strictly below (A.13) if and only if (A.16) holds. Next, observe that for $\delta_2^* = 0$, (A.14) is strictly above (A.13) because $q^L + q^H > \kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)}$.

Since (A.13) is strictly increasing and (A.14) is strictly decreasing in δ_2^*, there is a unique $\delta_2^* \in (0, 1)$ such that (A.13) equals (A.14) if and only if (A.16). Otherwise, i.e. if and only if (A.15), we have $\delta_2^* = 1$.

We now pin down δ_2^*. Note that if $\delta_2^* < 1$, it is determined by (A.13) equal (A.14), i.e.

$$\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)} = \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} + q^L$$

We have two candidates that solve this equation:

$$\delta_2^* = \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1-\delta_1^*)(1-\gamma)((1-\kappa)q^L + \frac{e}{\Delta})} \pm \frac{[(2\gamma(q^H - q^L))^2 + ((1-\delta_1^*)(1-\gamma) + 2\gamma)^2((1-\kappa)q^L + \frac{e}{\Delta})^2]^{\frac{1}{2}}}{2(1-\delta_1^*)(1-\gamma)((1-\kappa)q^L + \frac{e}{\Delta})}$$

Recall that $\delta_2^* \in (0, 1)$. In the subtraction case, $1 - \frac{\gamma(q^H - q^L)}{(1-\delta_1^*)(1-\gamma)((1-\kappa q^L) + \frac{e}{\Delta})} < 0$. Therefore, we conclude that

$$\delta_2^* = \frac{1}{2} - \frac{\gamma(q^H - q^L)}{(1-\delta_1^*)(1-\gamma)((1-\kappa)q^L + \frac{e}{\Delta})} + \frac{[(2\gamma(q^H - q^L))^2 + ((1-\delta_1^*)(1-\gamma) + 2\gamma)^2((1-\kappa)q^L + \frac{e}{\Delta})^2]^{\frac{1}{2}}}{2(1-\delta_1^*)(1-\gamma)((1-\kappa)q^L + \frac{e}{\Delta})}$$

Further, $\delta_2^* > \frac{1}{2} + \frac{\gamma(q^H-q^L)}{(1-\delta_1^*)(1-\gamma)((1-\kappa)q^L+\frac{e}{\Delta})} - \frac{2\gamma(q^H-q^L)}{2(1-\delta_1^*)(1-\gamma)((1-\kappa)q^L+\frac{e}{\Delta})} = \frac{1}{2}$.

Therefore, $\delta_2^* \in (\frac{1}{2}, 1)$ if and only if $(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\delta_1^*)(1-\gamma)(q^H-q^L)}{\gamma+(1-\delta_1^*)(1-\gamma)}$, and $\delta_2^* = 1$ otherwise.

We now turn our attention to period 1 and characterize δ_1^*. We first show that δ_1^* is unique. Recall that in period 1, the low-quality firm charges a low price $\kappa q^L - \frac{e}{\Delta}$ with probability $1-\delta_1^*$ and obtains a rating, and it charges a high price $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$ with probability δ_1^* and obtains no rating in period 1.

When the low-quality firm charges $\kappa q^L - \frac{e}{\Delta}$, the total continuation profit of the low-quality firm is

$$\pi_1^L(R_1 = 1) = \kappa q^L - \frac{e}{\Delta} + (1-\delta_2^*)[\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}] + \delta_2^*[\frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} + q^L], \tag{A.17}$$

where the first term are the profits in period 1, the second term are the continuation profits if the firm charges a low price and obtains a good rating in period 2, and the third term are the continuation profits if the firm charges a high price and obtains no rating in period 2.

We show that (A.17) is strictly increasing in δ_1^* if $\gamma > \frac{1}{3}$.

To show that (A.17) is strictly increasing in δ_1^*, we first show that $\frac{\partial \delta_2^*}{\partial \delta_1^*} > 0$.

Taking the derivative of (A.13) and (A.14) and solving for $\frac{\partial \delta_2^*}{\partial \delta_1^*}$. We find that

$$\frac{\partial \delta_2^*}{\partial \delta_1^*} = \frac{\delta_2^*(\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2 - (1-\delta_2^*)(\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma))^2}{(1-\delta_1^*)((\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2 + (\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma))^2)}. \tag{A.18}$$

Notice that the denominator is positive. The numerator simplifies to $(2\delta_2^* - 1)(\gamma^2 - (1-\delta_1)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2)$. We show that if $\gamma > \frac{1}{3}$, we have $(2\delta_2^* - 1)(\gamma^2 - (1-\delta_1)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2) > 0$. Since $\delta_2^* \in (\frac{1}{2}, 1]$, $2\delta_2^* - 1 > 0$. We know that $\delta_1^* \in (0, 1]$ and that $\delta_2^*(1-\delta_2^*)$ is maximum at $\delta_2^* = 0.5$. Therefore, $\gamma^2 - (1-\delta_1)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2 > \gamma^2 - 0.25(1-\gamma)^2$ and $\gamma^2 - 0.25(1-\gamma)^2 > 0$ whenever $\gamma > \frac{1}{3}$. Thus, if $\gamma > \frac{1}{3}$, we have $(2\delta_2^* - 1)(\gamma^2 - (1-\delta_1)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2) > 0$ and (A.18) > 0. $\gamma > \frac{1}{3}$ is a sufficient condition.

Next, we show that total derivative of (A.17) with respect to δ_1^* is greater than 0.

$$\begin{aligned}\frac{\partial \pi_1^L(R_1=1)}{\partial \delta_1^*} =& \frac{\partial \delta_2^*}{\partial \delta_1^*}[q^L(1-\kappa)+\frac{e}{\Delta}-\frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}+\\ &\frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}]+\frac{\gamma(1-\gamma)\delta_2^*(q^H-q^L)(\delta_2^*-\frac{\partial\delta_2^*}{\partial\delta_1^*}(1-\delta_1^*))}{(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2}+\\ &\frac{\gamma(1-\gamma)(1-\delta_2^*)(q^H-q^L)(\frac{\partial\delta_2^*}{\partial\delta_1^*}(1-\delta_1^*)+(1-\delta_2^*))}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2}\end{aligned}$$

We show that $\frac{\partial \pi_2^L}{\partial \delta_1^*}>0$ in three parts.

Consider first that $\frac{\partial \delta_2^*}{\partial \delta_1^*}[q^L(1-\kappa)+\frac{e}{\Delta}]$ and show that this is greater than 0. Since $\frac{\partial \delta_2^*}{\partial \delta_1^*}>0$ when $\gamma>\frac{1}{3}$ and $q^L>\kappa q^L-\frac{e}{\Delta}$, we can conclude that $\frac{\partial \delta_2^*}{\partial \delta_1^*}[\kappa q^L-\frac{e}{\Delta}-q^L]>0$.

Second, we reformulate $\frac{\partial \pi_2^L}{\partial \delta_1^*}$.

To do so, consider $\frac{\gamma(1-\gamma)(1-\delta_2^*)(q^H-q^L)(\frac{\partial\delta_2^*}{\partial\delta_1^*}(1-\delta_1^*)+(1-\delta_2^*))}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2}+\frac{\gamma(1-\gamma)\delta_2^*(q^H-q^L)(\delta_2^*-\frac{\partial\delta_2^*}{\partial\delta_1^*}(1-\delta_1^*))}{(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2}$. Substituting (A.18), simplifies to $\frac{\gamma(1-\gamma)(q^H-q^L)}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2+(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2}$. This leaves us with the simplified equation of

$$\begin{aligned}\frac{\partial \pi_1^L(R_1=1)}{\partial \delta_1^*} =& \frac{\partial \delta_2^*}{\partial \delta_1^*}[q^L(1-\kappa)+\frac{e}{\Delta}-\frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}+\\ &\frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}]+\\ &\frac{\gamma(1-\gamma)(q^H-q^L)}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2+(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2}.\end{aligned}$$

Finally, since $\frac{\partial \delta_2^*}{\partial \delta_1^*}[q^L(1-\kappa)+\frac{e}{\Delta}]>0$, we show that $\frac{\partial \delta_2^*}{\partial \delta_1^*}[\frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}-\frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}]+\frac{\gamma(1-\gamma)(q^H-q^L)}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2+(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2}>0$. Substituting (A.18), we get

$$\begin{aligned}&\frac{\gamma(1-\gamma)(q^H-q^L)}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2+(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2}[1+\\ &\frac{(1-\delta_1^*)(1-2\delta_2^*)(\delta_2^*(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2-(1-\delta_2^*)(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2)}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))}]\end{aligned}$$

We check that this is positive. Notice that $\frac{\gamma(1-\gamma)(q^H-q^L)}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2+(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2}>0$. Therefore, what remains is to show $1+\frac{(1-\delta_1^*)(1-2\delta_2^*)(\delta_2^*(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))^2-(1-\delta_2^*)(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))^2)}{(\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma))(\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma))}>0$, and this is true if and only if

$$\begin{aligned}&(1-\delta_1^*)(1-2\delta_2^*)^2[\gamma^2-(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2]\\ <&\gamma^2+\gamma(1-\gamma)(1-\delta_1^*)+(1-\gamma)^2(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*).\end{aligned}$$

We can verify that this is always true. Since $\delta_1^* \in [0,1]$ and $\delta_2^* \in [\frac{1}{2},1]$, $\delta_1^*(1-2\delta_2^*)^2 \in [0,1]$. Therefore, $(1-\delta_1^*)(1-2\delta_2^*)^2[\gamma^2-(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2] < \gamma^2-(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2$. It is easy to see that $\gamma^2-(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2 < \gamma^2+\gamma(1-\gamma)(1-\delta_1^*)+(1-\gamma)^2(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*)$. For $-(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*)(1-\gamma)^2 < \gamma(1-\gamma)(1-\delta_1^*)+(1-\gamma)^2(1-\delta_1^*)^2\delta_2^*(1-\delta_2^*)$, the left hand side is negative and the right hand side is positive.

We conclude that $\frac{\partial \pi_1^L(R_1=1)}{\partial \delta_1^*} > 0$ when $\gamma > \frac{1}{3}$. Note that $\gamma > \frac{1}{3}$ is a sufficient but not necessary condition.

Next consider the situation when the low-quality firm charges a high price $\frac{\gamma q^H+\delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$ with probability δ_1^* and obtains no rating in period 1. The firm then charges q^L in subsequent periods. Thus, the total continuation profits are given by

$$\pi_1^L(R_1=0) = \frac{\gamma q^H+\delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)} + q^L + q^L \tag{A.19}$$

This is strictly decreasing in δ_1^*.

We now show that $\delta_1^* = 1$ is only possible if no mixed-strategy equilibrium in period 1 exists. Suppose that $\delta_1^* = 1$. Then (A.17) and (A.19) become $\kappa q^L - \frac{e}{\Delta} + (1-\delta_2^*)[\kappa q^L - \frac{e}{\Delta} + q^H] + \delta_2^*[q^H+q^L]$ and $\gamma q^H + (1-\gamma)q^L + q^L + q^L$ respectively. For $\delta_1^* = 1$ to be optimal, we must have (A.17) lower that (A.19) for $\delta_1^* = 1$,

$$\begin{aligned}
\kappa q^L - \frac{e}{\Delta} + (1-\delta_2^*)[\kappa q^L - \frac{e}{\Delta}] + \delta_2^*[q^L] + q^H &\leq \gamma q^H + (1-\gamma)q^L + q^L + q^L \\
\iff (2-\delta_2^*)((1-\kappa)q^L + \frac{e}{\Delta}) &\geq (1-\gamma)(q^H - q^L).
\end{aligned}$$

Further, when $\gamma > \frac{1}{3}$, (A.17) increases in δ_1^* and (A.19) decreases in δ_1^*. Hence, when $(2-\delta_2^*)((1-\kappa)q^L + \frac{e}{\Delta}) > (1-\gamma)(q^H-q^L)$, (A.19) is larger than (A.17) for all $\delta_1^* \leq 1$. This indicates that $\delta_1^* = 1$ is an equilibrium and no mixed-strategy equilibrium exists in period 1. We conclude that $\delta_1^* = 1$ can only be an equilibrium if no mixed-strategy equilibrium exists.

We now show that $\delta_1^* = 0$ is not an equilibrium. Suppose towards a contradiction that $\delta_1^* = 0$. Then (A.17) and (A.19) become $\kappa q^L - \frac{e}{\Delta} + (1-\delta_2^*)[\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H+(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_2^*)(1-\gamma)}] + \delta_2^*[\frac{\gamma q^H+\delta_2^*(1-\gamma)q^L}{\gamma+\delta_2^*(1-\gamma)} + q^L]$ and $q^H+q^L+q^L$ respectively. Since $\kappa q^L - \frac{e}{\Delta} < q^L$, $\frac{\gamma q^H+(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_2^*)(1-\gamma)} < q^H$, and $\frac{\gamma q^H+\delta_2^*(1-\gamma)q^L}{\gamma+\delta_2^*(1-\gamma)} < q^H$ then (A.17) < (A.19). Therefore, if $\delta_1^* = 0$ and the low-quality firm participates in ratings harvesting with probability 1, deviating to the high price and obtaining no ratings is profitable. This contradicts $\delta_1^* = 0$. We conclude that $\delta_1^* = 0$ cannot be an equilibrium.

We now show that for

$$(2-\delta_2^*)((1-\kappa)q^L + \frac{e}{\Delta}) < (1-\gamma)(q^H - q^L), \tag{A.20}$$

there is a unique $\delta_1^* \in (0,1)$ that characterises the low-quality firms mixed-strategy in period 1. Recall that a mixed-strategy equilibrium only exists when (A.17) and (A.19) are equal. First, observe that for $\delta_1^* = 1$, (A.19) is strictly below (A.17). Next, observe that $\delta_1^* = 0$, (A.19) is strictly above (A.17). Since (A.17) is strictly increasing in δ_1^* when $\gamma > \frac{1}{3}$ and (A.19) is strictly decreasing in δ_1^*, there is a unique $\delta_1^* \in (0,1)$ such that (A.17) and (A.19) are equal if (A.20) and $\gamma > \frac{1}{3}$ hold.

We have shown that we cannot have an equilibrium where $\delta_1^* = 0$ or $\delta_2^* = 0$, and that $\delta_2^* = 1$ can be an equilibrium if no mixed-strategy equilibrium exists. Additionally, if $\gamma > \frac{1}{3}$, then $\delta_1^* = 1$ can be an equilibrium if no mixed-strategy equilibrium exists and if a mixed-strategy exists, $\delta_1^* \in (0,1)$ exists such that (A.17) and (A.19) are equal, it is unique if $\gamma > \frac{1}{3}$. And, if $\delta_2^* \in (0,1)$ exists such that (A.13) and (A.14) are equal, it must be unique. Therefore, we either have a unique pure-strategy in equilibrium or a unique mixed-strategy in equilibrium in each period when $\gamma > \frac{1}{3}$. We have also characterised the mixed-strategy equilibrium for period 2 and shown that $\delta_2^* \in (\frac{1}{2}, 1]$. We conclude that statement 2 holds, and therefore conclude that statements 1-4 hold in equilibrium.

We now show that equilibria satisfying statements 1-4 indeed exist.

We begin by considering the scenario where (A.16), (A.20), $\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}, \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}\}$ and $\gamma > \frac{1}{3}$ hold together.

Consumers' beliefs are as follows. In period 1,

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)} & \text{if } p_1 > \kappa q^L - \frac{e}{\Delta} \\ q^L & \text{if } p_1 \leq \kappa q^L - \frac{e}{\Delta}, \end{cases}$$

and in period 2,

$$E[q_2|R_1, p_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} & \text{if } p_2 > \kappa q^L - \frac{e}{\Delta} \text{ and } R_1 = 1 \\ q^L & \text{for any other } p_2,\ R_1 \text{ combination,} \end{cases}$$

and in period 3,

$$E[q_3|R_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)} & \text{if } R_2 = \{11\} \\ q^L & \text{for any other } R_2. \end{cases}$$

These beliefs follow Bayes rule on the candidate equilibrium path of play. The candidate equilibrium is consistent with our restrictions. Because high-quality firms obtain a good rating with probability 1, the candidate equilibrium is consistent with Restriction 1. Further, whenever low-quality firms obtain no rating, consumers' beliefs are independent of prices (in both periods 1 and 2), and in period 3, beliefs are the same for all period 3 prices and they only depend on the history of ratings. Therefore, the candidate equilibrium is consistent with Restriction 2.

The candidate equilibrium is such that high-quality firms always play $p_t^H = E[q_t|R_{t-1}, p_t^H]$ for all t and get a good rating. Low-quality firms mix between playing a low price $\kappa q_L - \frac{e}{\Delta}$ and getting a good rating with probability δ_t^*, and setting a high price p_t^H and getting no rating with probability $(1-\delta_t^*)$ in each period $t \in \{1, 2\}$.

More precisely, in the candidate equilibrium in period 1, the high-quality firm sets a price $p_{1,1}^H = \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}$ and obtains a good rating with probability 1. In period 2, conditional on receiving a good rating in period 1, the high-quality firm sets a price $p_{2,11}^H = \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}$ and obtains a good rating with probability 1. If the high-quality firm did not receive a good rating in period 1, consumers believe the firm is of a low quality and thus the maximum price that the high-quality firm can set is q^L and receives a good rating. To see that the firm receives a good rating, note that by assumption $\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}, \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}\}$ and both of $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}$ and $\frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}$ are larger than q^L. Therefore, when consumers pay q^L for a high quality product, they receive a sufficient amount of excess surplus and leave a good rating. In the third period of the candidate equilibrium, having received a continuation of good ratings, i.e. a rating history of $R_2 = 11$, the high-quality firm sets a price of $\frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$. If the firm did not receive a continuation of good ratings, i.e. a rating history of $R_2 \in \{00, 01, 10\}$, then it sets a maximum price of q^L.

In period 1 of the candidate equilibrium, the low-quality firm sets $p_{1,0}^L = \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}$ and obtains no rating with some probability δ_1^*, and $p_{1,1}^L = \kappa q^L - \frac{e}{\Delta}$ and obtains a good rating with some probability $1 - \delta_1^*$. In period 2, conditional on having obtained a good rating in period 1, the low-quality firm sets a price $p_{2,11}^L = \kappa q^L - \frac{e}{\Delta}$ with some probability $1 - \delta_2^*$ and obtains a good rating, and a price $p_{2,10}^L = \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)}$ with some probability δ_2^* and obtains no rating. If the low-quality firm obtained no rating in period 1, then it sets a price q^L in period 2 and obtains no rating. In period 3, if the low-quality firm received a continuation of good ratings, i.e. a rating history of $R_2 = 11$, it sets a price of $\frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$, and when it receives a rating history of $R_2 = \{00, 10, 01\}$, it sets a price of q^L. This fully characterizes the firms' prices and consumers' beliefs.

We now show that the firms have no profitable deviations.

In the candidate equilibrium, the high-quality firm earns a total profit of $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)} + \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma + (1-\delta_1^*)\delta_2^*(1-\gamma)} + \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$.

In period 3, deviations to a higher price would reduce demand to zero and earn a total maximum profit of $\frac{\gamma q^H+\delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)} + \frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)} + 0$, and deviations to a lower price would reduce profit margins in the third period without increasing demand. Neither deviation increases profits.

In period 2, deviations to a higher price would reduce demand to zero and result in no rating, this leads to a total maximum profit of $\frac{\gamma q^H+\delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)} + 0 + q^L$, and deviations to a lower price would reduce profit margins in the second period without increasing demand. Neither deviation increases profits.

In period 1, deviations to a higher price would reduce demand to zero and result in no rating, this leads to a total maximum profit of $0 + q^L + q^L$, and deviations to a lower price would reduce profit margins in the first period without increasing demand. Neither deviation increases profits.

Therefore, there are no profitable deviations for the high-quality firm in any period.

Next, we show that low-quality firms have no profitable deviation.

In the candidate equilibrium, the low-quality firm earns a total profit of $\kappa q^L - \frac{e}{\Delta} + \kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$. For the equilibrium $\delta_1^* \in (0,1)$ and $\delta_2^* \in (0,1)$, the low-quality firm is indifferent between setting the lower price that obtains good rating and setting a higher price that obtains no rating in all periods. Hence, the firm is indifferent between the total profits of $\kappa q^L - \frac{e}{\Delta} + \kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$, $\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)} + q^L$ and $\frac{\gamma q^H+\delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)} + q^L + q^L$. Recall that (A.16) implies $\delta_2^* \in (0,1)$ and (A.20) implies $\delta_1^* \in (0,1)$, and $\gamma > \frac{1}{3}$ implies $\delta_1^* \in (0,1)$ is unique. Showing that there exists no profitable deviation from any of these cases shows that there is no profitable deviation for the low-quality firm.

In period 3, deviations towards a price above $\frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$ would result in zero demand and a maximum total profit of $\kappa q^L - \frac{e}{\Delta} + \kappa q^L - \frac{e}{\Delta} + 0$, and a deviation towards a lower price results in smaller profit margins without increasing demand. Hence, there is no profitable deviation in period 3.

In period 2, the low-quality firm is indifferent between setting the price $\kappa q^L - \frac{e}{\Delta}$ and $\frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}$, where $\frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)} > \kappa q^L - \frac{e}{\Delta}$ and her total profits are $\kappa q^L - \frac{e}{\Delta} + \kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$ and $\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)} + q^L$ respectively. Deviations towards a price above $\frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}$ leads to zero demand in period 2 and the maximal profits of $\kappa q^L - \frac{e}{\Delta} + 0 + q^L$, which is not a profitable deviation. Deviations towards a price $p_2 \in (\kappa q^L - \frac{e}{\Delta}, \frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)})$, leads to the same rating as when the firm sets $\frac{\gamma q^H+(1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}$; hence, any deviation to p_2 lowers margin without improving demand or future profit, and are therefore not profitable. Deviations towards a price below $\kappa q^L - \frac{e}{\Delta}$ decreases profits in period 2 but does not increase demand or change the rating the firm receives when setting a price of $\kappa q^L - \frac{e}{\Delta}$, which is why this is also not a profitable deviation. Therefore, in period 2, there is no profitable deviation for the firm.

In period 1, the low-quality firm is indifferent between setting the price $\kappa q^L - \frac{e}{\Delta}$ and $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$, where its total profits are $\kappa q^L - \frac{e}{\Delta} + \kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\delta_2^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\delta_2^*)(1-\gamma)}$ and $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)} + q^L + q^L$ respectively. If it deviates to a price above $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$, demand falls to zero and it makes a total profit of $0 + q^L + q^L$, which is not a profitable deviation. If it deviates to a price $p_1 \in (\kappa q^L - \frac{e}{\Delta}, \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)})$, then it receives the same rating as when it sets the price of $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$; hence such a deviation reduces margins in the first period without improving demand or future profit and is not profitable. When deviating to a price below $\kappa q^L - \frac{e}{\Delta}$, the firm receives a good rating; however, the deviation does not increase demand and her margins are lower than when setting the price of $\kappa q^L - \frac{e}{\Delta}$, which is why the deviation is not profitable.

We conclude that there are no profitable deviations for either high- or low-quality firms from the candidate equilibrium.

We conclude that if $(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\delta_1^*)(1-\gamma)(q^H-q^L)}{\gamma+(1-\delta_1^*)(1-\gamma)}$, $(1-\kappa)q^L + \frac{e}{\Delta} < \frac{(1-\gamma)(q^H-q^L)}{2-\delta_2^*}$, $\gamma > \frac{1}{3}$ and $\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}, \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}\}$, the candidate equilibrium exists.

Consider next the case where (A.15), (A.20), $\gamma > \frac{1}{3}$ and $\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}, \frac{\gamma q^H + (1-\delta_1^*)\delta_2^*(1-\gamma)q^L}{\gamma+(1-\delta_1^*)\delta_2^*(1-\gamma)}\}$. In this scenario, $\delta_2^* = 1$. The last three inequalities are identical to the previous case and play the same role in this case.

In the candidate equilibrium, consumers' beliefs are as follows. In period 1,

$$E[q_1|p_1] = \begin{cases} \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)} & \text{if } p_1 > \kappa q^L - \frac{e}{\Delta} \\ q^L & \text{if } p_1 \leq \kappa q^L - \frac{e}{\Delta}, \end{cases}$$

and in period 2,

$$E[q_2|R_1, p_2] = \begin{cases} \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)} & \text{if } p_2 > \kappa q^L - \frac{e}{\Delta} \text{ and } R_1 = 1 \\ q^L & \text{for any other } p_2,\ R_1 \text{ combination,} \end{cases}$$

and in period 3,

$$E[q_3|R_2] = \begin{cases} q^H & \text{if } R_2 = \{11\} \\ q^L & \text{for any other } R_2. \end{cases}$$

These beliefs follow Bayes rule on the candidate equilibrium path of play. The candidate equilibrium is consistent with our restrictions. Because high-quality firms obtain a good rating with probability 1, the candidate equilibrium is consistent with Restriction 1. Further, whenever low-quality firms obtain no rating, consumers' beliefs are independent of prices (in both periods 1 and 2), and in period 3, beliefs are the same for all prices. Therefore, the candidate equilibrium is consistent with Restriction 2.

In the candidate equilibrium in period 1, the high-quality firm sets a price $p^H_{1,1} = \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$ and obtains a good rating with probability 1. In period 2, conditional on receiving a good rating in period 1, the high-quality firm sets a price $p^H_{2,11} = \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)}$ and obtains a good rating with probability 1. If the high-quality firm did not receive a rating in period 1, consumers believe the firm is of a low quality and thus the maximum price that the high-quality firm can set is q^L and the firm receives a good rating. To see that the firm receives a good rating, note that by assumption $\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}, \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)}\}$ and both of $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$ and $\frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)}$ are larger than q^L. In the third period of the candidate equilibrium, having received a continuation of good ratings, i.e. a rating history of $R_2 = 11$, the high-quality firm sets a price of q^H. If the firm did not receive a continuation of good ratings, i.e. a rating history of $R_2 \in \{00, 01, 10\}$, then it sets a maximum price of q^L.

In period 1 of the candidate equilibrium, the low-quality firm sets a price $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}$ and obtains no rating with some probability δ_1^* and $\kappa q^L - \frac{e}{\Delta}$ and obtains a good rating with some probability $1-\delta_1^*$. In period 2, the low-quality firm sets a price of $\frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)}$ with probability 1 and receives no rating. In period 3, if the low-quality firm received a continuation of good ratings, i.e. a rating history of $R_2 = 11$, it sets a price of q^H, and when it receives a rating history of $R_2 = \{00, 10, 01\}$, it sets a price of q^L.

We now show that the firms have no profitable deviations.

In the candidate equilibrium, the high-quality firm earns a total profit of $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)} + \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)} + q^H$. In period 3, deviations to a higher price reduces demand to zero, and deviations to a lower price reduces profit margins without increasing demand, thus there is no profitable deviation in period 3. In period 2, deviations to a higher price reduces demand to zero, which induces no rating and also reduces continuation profits; deviations to a lower price reduces profit margins without increasing demand. Thus, neither deviation is profitable in period 2. In period 1, deviations to a higher price reduces demand to zero, induce no rating and therefore reduces continuation profits in all future periods to q_L, and deviations to a lower price reduces profit margins without increasing demand. Neither deviation is profitable in period 1. Therefore, there are no profitable deviations for the high-quality firm.

In the candidate equilibrium, the low-quality firm earns a total profit of $\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)} + q^L$. For the equilibrium $\delta_1^* \in (0,1)$, the low-quality firm is indifferent between setting a lower price that obtains a good rating and setting a higher price that obtains no rating in period 1, and for $\delta_2^* = 1$, the low-quality firm always sets a high price and obtains no rating in period 2. The low-quality firm is therefore indifferent between the total profits of $\kappa q^L - \frac{e}{\Delta} + \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)} + q^L$ and $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)} + q^L + q^L$. Recall that (A.15) implies $\delta_2^* = 1$, (A.20) implies $\delta_1^* \in (0,1)$, and $\gamma > \frac{1}{3}$ implies $\delta_1^* \in (0,1)$ is unique. In period 3, deviations towards a higher price, above q^L, reduces demand to zero, and deviations towards a lower price results in a lower profit margin without improving demand, which is why there is no profitable deviation in period 3.

In period 2, conditional on receiving a good rating in period 1, deviations towards a price above $\frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)}$ leads to zero demand and results in the same rating (and hence future profit) as setting a price of $\frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)}$, which is why it is not a profitable deviations. Further, deviations towards a price of $p_2 \in (\kappa q^L - \frac{e}{\Delta}, \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)})$ results in the same rating (and hence future profit) as setting a price of $\frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma + (1-\delta_1^*)(1-\gamma)}$, therefore this reduces margins without providing any additional continuation profit, and is not a profitable deviation. Deviations towards a price $p_2 \leq \kappa q^L - \frac{e}{\Delta}$ leads to a maximal total profit of $\kappa q^L - \frac{e}{\Delta} + \kappa q^L - \frac{e}{\Delta} + q^H$, but since (A.15) holds, this is not a profitable deviation. Therefore, there are no profitable deviations in period 2 conditional on receiving a good rating in period 1. In period 2, conditional on having received no rating in period 1, deviations towards a price above q^L leads to zero demand and does not change rating and continuation profits, which is why this is not a profitable deviation. Further, since consumers beliefs on observing a single period of no rating is that the firm is of a low-quality, setting a price lower than q^L only reduce the margins in period 2 without increasing demand or future profits, thus this is not a profitable deviation. Therefore, there are no profitable deviations in period 2 conditional on not having received a rating in period 1.

In period 1, the low-quality firm is indifferent between setting the price $\kappa q^L - \frac{e}{\Delta}$ and $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}$. If it deviates to a price above $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}$, demand falls to zero, it gets no rating and it makes a total profit of $0 + q^L + q^L$, which is not a profitable deviation. If it deviates to a price $p_1 \in (\kappa q^L - \frac{e}{\Delta}, \frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)})$, then it receives the same rating as when it sets the price of $\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma + \delta_1^*(1-\gamma)}$; hence such a deviation reduces margins in the first period without improving demand or future profit and is not a profitable deviation. When deviating to a price below $\kappa q^L - \frac{e}{\Delta}$, the firm receives a good rating; however, the deviation does not increase demand and her margins are lower than when setting the price of $\kappa q^L - \frac{e}{\Delta}$. Therefore, we conclude that there is no profitable deviation in period 1.

We conclude that there are no profitable deviations for either the high- or low- quality firm from the candidate equilibrium.

We conclude that if $(1-\kappa)q^L + \frac{e}{\Delta} \geq \frac{(1-\delta_1^*)(1-\gamma)(q^H - q^L)}{\gamma + (1-\delta_1^*)(1-\gamma)}$, $(1-\kappa)q^L + \frac{e}{\Delta} < (1-\gamma)(q^H - q^L)$,

$\gamma > \frac{1}{3}$ and $\kappa q^H - \frac{e}{\Delta} \geq max\{\frac{\gamma q^H + \delta_1^*(1-\gamma)q^L}{\gamma+\delta_1^*(1-\gamma)}, \frac{\gamma q^H + (1-\delta_1^*)(1-\gamma)q^L}{\gamma+(1-\delta_1^*)(1-\gamma)}\}$, the candidate equilibrium exists.

This concludes the proof. □

A.2.3 Deriving Rating Utility

We derive the rating utility function from the utility function proposed in Rabin (1993). In his paper, Rabin proposes a utility function which incorporates a reciprocity term in addition to a consumption utility. This reciprocity term depends on the additional surplus that some player i is allowed to obtain given the actions of another player j relative to a predefined equity point.

$$U_i = \pi_i + [\frac{\pi_i - \pi_i^e}{\pi_i^H - \pi_i^{min}}][1 + \frac{\pi_j - \pi_j^e}{\pi_j^H - \pi_j^{min}}]$$

- For all $h \in \{i, j\}$.
- π_h is the utility.
- π_h^{min} is the lowest possible payoff to player h.
- π_h^H is the highest possible pareto efficient payoff to player h.
- $\pi_h^e = \frac{\pi_h^H + \pi_h^L}{2}$, where π_h^L denotes the lowest possible pareto efficient payoff to player h. π_h^e is the equitable reference point.

At this junction, allow us to provide some intuition. Suppose that player i is the consumer. Then this function takes into account the consumption utility and some additional reciprocity term. The additional term is what we consider the rating utility. Suppose a firm sets a low price, such that $\pi_i - \pi_i^e > 0$, the consumer would believe that the firm is treating him kindly. In response, consumers will receive a higher overall utility if he is kind to the firm, $\pi_j - \pi_j^e > 0$. In our context, a good rating will result in π_j being higher and therefore by leaving a good rating, consumers would be being kind to the firms.

On the contrary, a firm charging a high price such that pay off for the consumer is below the equitable point would result in consumers punishing the firm by lowering their profits in the future periods - perhaps through a negative rating. For simplicity, we remove the ability of consumers to punish a firm and assume that consumers are only able to provide good or no ratings. Specifically, in our model, we show that a good rating results in a better future pay off and hence satisfies this feature.

In what follows, I show how we adapt this framework to better fit our context. Firstly, we have assumed that providing a rating is costly for consumers. This seems to be an intuitive feature of our model and follows from the literature of costly provision of reviews and ratings (Avery et al., 1999; Miller et al., 2005). Secondly, we make some simplifications that make the framework more tractable in our setup. We remove the normalization terms in the denominator and remove the "1".[6] Thirdly, we consider that the rating component of the utility only comes into effect when a good rating is provided. This leaves us with the following function:

$$U_i = \pi_i + \mathbb{1}_{\{R_t=1\}}\{[\pi_i - \pi_i^e][\pi_j - \pi_j^e] - e\}$$

Since $\pi_i = q^j - p_1^j$, $\pi_i^e = \frac{(q^j - q^j) + (q^j - 0)}{2}$. The highest possible pareto payoff to consumers being $q^j - 0$ where sellers set a price of 0 and the lowest possible being 0, where sellers set a price of q^j.

Moreover, $\pi_j = p_1 + p_2$, profits of the firm being the sum of profits in two periods, given 0 marginal cost, profits is the sum of prices in both periods. And $\pi_j^e = \frac{(p_1 + q^H) + (p_1 + q^L)}{2}$. The seller, setting some price p_1 in period 1, is able to get a maximum benefit of q^H and a minimum benefit of q^L in period 2.

This leaves us with:

$$U_i = q^j - p_1 + \mathbb{1}_{\{R_t=1\}}\{[\frac{q^j}{2} - p_1][p_2 - \frac{q^H + q^L}{2}] - e\}$$

Next, we replace $\frac{q^j}{2}$ with κq^j, where $\kappa \in [0, 1]$ and $[p_2 - \frac{q^H + q^L}{2}]$ with Δ. This reflects the notion that an equitable payoff may not be one of equal split, allowing us to generalise the equitable point. Hence, when κ is sufficiently high, firms are able to charge some price slightly below quality and still receive a positive rating if e is sufficiently small. We do not make any assumptions over Δ, except that it is positive. This allows us to capture that consumers may not fully understand how firms benefit from ratings, only that a good rating is beneficial for a firm, and a bad rating can harm the firm. Thus, capturing kindness from consumers which enables firms to gain some benefits in subsequent periods.

Finally, we split the consumption utility and the rating utility. This allows for more compatible purchase decision across periods as consumer's purchase decision does not depend on whether they anticipate giving a good rating.

[6] Rabin notes that doing so does not affect the behavior of the utility function.

Appendix B
Ratings with Heterogeneous Preferences

B.1 Experiment Instructions

B.1.1 Consent

Consent Form

Carleton College, Informed Consent Form

Overview and Procedure

Thank you for considering participation in this study. Today's session should take approximately 15 minutes to complete. The purpose of the study is to collect information on how people make decisions and share information in economic environments. If you agree to participate you will be asked to read instructions and then complete a series of simple economic decisions involving sharing information and choosing small amounts of money. You will be paid a fixed fee of $3.00 for participating, plus additional potential earnings based on your choices. Average earnings are expected to be approximately $6.00 per participant.

Risks and Benefits

There are no direct risks to you as a participant in this study. Other than monetary compensation there are no direct benefits to you from the study. The data collected from this study may help to further general understanding of economic decision making.

Confidentiality

Your privacy will be protected and at no time will your name or any other identifying information be attached to your responses. Any information obtained from you during this experiment will remain confidential and will be used solely for research purposes.

Your Rights

Your participation in this study is voluntary. You may withdraw from the study at any time without penalty.

If you have questions or concerns about your rights in this study, please contact the principle investigator, Jonathan Lafky at jlafky@carleton.edu or Carleton College, One North College Street, Northfield MN, 55057. If you would like to contact someone other than the researcher, please contact the Institutional Review Board for Research with Human Subjects at Carleton College, c/o Office of the Associate Provost, Carleton College, One North college Street, Northfield MN, 55057; telephone (507) 222-4301.

I certify that I have read all of the information in this consent form. I agree to participate in this study of decision making, and I understand that I can withdraw at any time without penalty.

Do you consent to participating in this experiment?

- ◯ I do not consent to participate in this study.
- ◯ I consent to participate in this study.

To ensure your payment, please enter your Prolific ID

Please double check that your ID is correct, otherwise you may not receive payment.

Begin

Figure B.1: Consent Form

B.1.2 Raters

Instructions

Instructions

Welcome and thank you for participating in this experiment on decision making. Please read these instructions carefully as they explain how you will earn money.

In today's experiment, your identity and decisions will be kept anonymous. This means that you will not know the identity or decisions of any other participant, nor will any other participant know your identity or decisions at any time during or after the experiment.

The experiment will be broken into 20 identical rounds of decision making.

Next

Instructions

Today, you will receive a fixed fee of $3.00, plus potential additional earnings based on the decisions you make. The experiment will be broken into 20 identical rounds of decision making. During each round you can potentially earn tokens, with each token worth $0.50.

At the end of the experiment, one of the 20 rounds will be randomly selected for payment. The tokens you earned in that round, and only that round, plus your initial $3.00 fee, determine your final payment.

Previous Next

Instructions

Before the first round, you will be randomly assigned a type, either type 1 or type 2. You may be assigned either type with equal chance. In other words, there is a 50% chance you will be type 1, and a 50% chance you will be type 2. Your type will be made known to you.

Your type will determine two "weights" called w_1 and w_2, that can influence your earnings. If you are type 1, you will receive the weights $w_1 = 1$ and $w_2 = 0.1$. If, instead, you were type 2, you will receive the weights $w_1 = 0.1$ and $w_2 = 1$.

The effects these weights have are described on the next page. Your type will not change between rounds.

Previous Next

Figure B.2: Rater's Instructions (i)

Instructions

In each round, you will be assigned a new randomly generated "prize". A prize is made up of two numbers, X_1 and X_2. Each of X_1 and X_2 are randomly chosen whole numbers between 1 and 10 inclusive. Each number is equally likely to be drawn. Together, X_1, X_2 and the weights w_1, w_2 determine the value of the prize.

The value of the prize (in tokens) is: $w_1 \times X_1 + w_2 \times X_2$.

In other words, the larger w_1 is, the more X_1 influences the value of the prize, and the larger w_2 is, the more X_2 influences the value of the prize.

An example of how to compute the value of a prize is provided on the next page.

Previous | Next

Instructions

Below is a brief example of how your type determines the value of a prize.

For example, suppose that you are assigned to be type 1.

If the prize takes the following numbers:

X_1	6
X_2	2

Then the value of the prize to you is $1 \times 6 + 0.1 \times 2 = 6.2$.

If, instead, you were type 2, the value of the same prize would be $0.1 \times 6 + 1 \times 2 = 2.6$.

Previous | Next

Instructions

After you receive a random prize and learn the exact numbers that X_1 and X_2 take, you will then be asked to rate the prize. A rating takes place on a scale from 1 to 5, where 1 is the worst rating and 5 is the best.

After you have decided your rating for the prize, you will be given the choice to pay 0.1 tokens to send the rating and your type to a future group of participants. If you choose not to send your rating, you pay no costs, and no one will see your rating and type.

After the experiment is completed, only one of the 20 rounds will be randomly chosen for payment. Your payment will depend on your earnings from that round. As a reminder, your earnings within a round is the value of the prize after subtracting any cost incurred from sending a rating in that round.

Previous | Next

Figure B.3: Rater's Instructions (ii)

Instructions

One or more participants in a future session will be offered the same prize as you. Each of the future participants has a 50% chance to be type 1 and a 50% chance to be type 2. Their type will be determined independently, and may be different from your own. These participants know the range of possible values that X_1 and X_2 can take.

They will also be told the actual number drawn for X_1 but not be told the actual number drawn for X_2. If you choose to send a rating, these participants will learn your rating as well as your type.

The future participant will be given a series of choices between obtaining either the prize or increasingly large amounts of money. One of these choices will be chosen at random to determine their earnings for the round. This means that the more money they are willing to forgo, the more likely they will obtain the prize. If they obtain the prize, their payoffs are calculated using the same weights as before. That is, the value of the prize to them is $w_1 \times X_1 + w_2 \times X_2$.

Previous | Next

Instructions

As a reminder, one and only one of the 20 rounds will be randomly chosen for payment after all rounds are completed. Your total payment for today's study will be equal to your fixed payment of $3.00, plus your earnings from the one round randomly selected for payment. Your earnings from that round will be converted into dollars at a rate of one token = $0.50.

Previous | Next

Instructions

To verify your understanding of the instructions, please answer the following questions:

If you were assigned to be type 2, what is the value of a prize of $X_1 = 5$ and $X_2 = 3$?

If you were assigned to be type 2, what is the probability that a future participant observing the same prize is type 1? Please enter a percentage between 0 and 100.

How many tokens does it cost to send a rating?

(Enter only the numerical value of your answers.)

Previous | Submit

Figure B.4: Rater's Instructions (iii)

Actions

Round 1

You are **type 1**. Your weights are:

w_1	1
w_2	0.1

In this round, the prize is made up of:

X_1	5
X_2	3

The value of this prize for a **type 1** participant is 5.3 tokens.
The value of this prize for a **type 2** participant is 3.5 tokens.

Next

How do you rate this prize?

◯ 1 ◯ 2 ◯ 3 ◯ 4 ◯ 5

Previous Next

Do you wish to pay 0.10 tokens to share your **rating and type** with future participants?

◯ Yes
◯ No

Previous Submit

Figure B.5: Rater's Actions

B.1.3 Consumers

Instructions

Instructions

Welcome and thank you for participating in this experiment on decision making. Please read these instructions carefully as they explain how you will earn money.

In today's experiment, your identity and decisions will be kept anonymous. This means that you will not know the identity or decisions of any other participant, nor will any other participant know your identity or decisions at any time during or after the experiment.

The experiment will be broken into 20 identical rounds of decision making.

Next

Instructions

Today, you will receive a fixed fee of $3.00, plus potential additional earnings based on the decisions you make. The experiment will be broken into 20 identical rounds of decision making. During each round you can potentially earn tokens, with each token worth $0.50.

At the end of the experiment, one of the 20 rounds will be randomly selected for payment. The tokens you earned in that round, and only that round, plus your initial $3.00 fee, determine your final payment.

Previous Next

Instructions

Before the first round, you will be randomly assigned a type, either type 1 or type 2. You may be assigned either type with equal chance. In other words, there is a 50% chance you will be type 1, and a 50% chance you will be type 2. Your type will be made known to you.

Your type will determine two "weights" called w_1 and w_2, that can influence your earnings. If you are type 1, you will receive the weights $w_1 = 1$ and $w_2 = 0.1$. If, instead, you are type 2, you will receive the weights $w_1 = 0.1$ and $w_2 = 1$.

The effects these weights have are described on the next page. Your type will not change between rounds.

Previous Next

Figure B.6: Consumer's Instructions (i)

Instructions

In each round, you will be shown a rating for a "prize". Each prize is randomly selected from a pool of prizes that have been rated by previous participants. The prizes were randomly generated before being shown to the previous participants.

A prize is made up of two numbers, X_1 and X_2. Each of X_1 and X_2 are randomly chosen whole numbers between 1 and 10 inclusive. Each number was equally likely to be drawn. Together, X_1, X_2 and the weights w_1, w_2 determine the value of the prize.

The value of the prize (in tokens) is: $w_1 \times X_1 + w_2 \times X_2$.

In other words, the larger w_1 is, the more X_1 influences the value of the prize, and the larger w_2 is, the more X_2 influences the value of the prize.

Each of the previous participants had a 50% chance of being type 1 and a 50% chance of being type 2. They were asked to rate a series of prizes using a scale of 1 to 5, where 1 is the worst rating and 5 is the best. The participants were then given the opportunity to share their rating with you at a cost to themselves. If they chose to share their rating, they paid a cost of 0.1 tokens. You will also learn of their type. You will only see prizes for which ratings were sent.

In each round, you will be shown one of the ratings that were sent. You will be asked how much you would be willing to pay for the prize that was rated. You will know the number that X_1 takes at the start of each round, but you will not know the number that X_2 takes until after you have made your decision for that round.

An example of how to compute the value of a prize is provided on the next page.

Previous Next

Instructions

Below is a brief example of how your type determines the value of a prize.

For example, suppose that you are assigned to be type 1.

If the prize takes the following numbers:

X_1	6
X_2	2

Then the value of the prize to you is $1 \times 6 + 0.1 \times 2 = 6.2$.

If, instead, you were type 2, the value of the same prize would be $0.1 \times 6 + 1 \times 2 = 2.6$.

Previous Next

Figure B.7: Consumer's Instructions (ii)

Instructions

To understand how much you would be willing to pay for the prize, you will be asked the following series of questions:

Question #		Option A		Option B
Question 1:	Would you rather have	the prize	or	0.1 tokens
Question 2:	Would you rather have	the prize	or	0.2 tokens
Question 3:	Would you rather have	the prize	or	0.3 tokens
⋮	⋮	⋮	⋮	⋮
Question 108:	Would you rather have	the prize	or	10.8 tokens
Question 109:	Would you rather have	the prize	or	10.9 tokens
Question 110:	Would you rather have	the prize	or	11.0 tokens

For each question, you will be asked if you prefer the prize (Option A) or a fixed amount of tokens (Option B). At the end of the experiment, a round will be randomly chosen. From this round, one question will be chosen at random and you will be paid the option that you chose on that question. Each round and each question are equally likely to be chosen for payment.

You have no incentive to lie on any question, because if that question gets chosen for payment, then you would end up with the option you like less. Note that the larger the question number, the larger the amount of tokens offered by Option B.

You will probably prefer the prize (Option A) in at least the first few questions, but at some point switch to prefering the fixed amount of tokens (Option B). So, to save time you will simply state the question at which you would first switch to prefering the fixed amount of tokens (Option B). The answers to the rest of the questions will be automatically "filled out" based on this switch point. This means you will be choosing the prize (Option A) for all questions before the switch point, and the fixed amount of tokens (Option B) for all questions at or after.

Previous Next

Figure B.8: Consumer's Instructions (iii)

Instructions

For example, suppose that in the round chosen for payment, the value of the prize is 3.0 tokens, and you chose to switch to Option B at Question 55. Question 46 is randomly chosen for payment. Because 46 is less than your switch point of 55, your payment will be based on the value of the prize, which is 3.0 tokens.

Suppose instead that Question 71 is randomly chosen for payment. Because 71 is greater than your switch point of 55, your payment will be based on Option B from Question 71, or 7.1 tokens.

Question #		**Option A**		**Option B**
Question 1:	Would you rather have	the prize	or	0.1 tokens
Question 2:	Would you rather have	the prize	or	0.2 tokens
Question 3:	Would you rather have	the prize	or	0.3 tokens
⋮	⋮	⋮	⋮	⋮
Question 108:	Would you rather have	the prize	or	10.8 tokens
Question 109:	Would you rather have	the prize	or	10.9 tokens
Question 110:	Would you rather have	the prize	or	11.0 tokens

Previous Next

Instructions

As a reminder, the answer to one randomly chosen question from one and only one of the 20 rounds will be chosen for payment after all rounds are completed. Your total payment for today's study will be equal to your fixed payment of $3.00, plus your earnings from the one round randomly selected for payment. Your earnings from that round will be converted into dollars at a rate of one token = $0.50.

Previous Next

Figure B.9: Consumer's Instructions (iv)

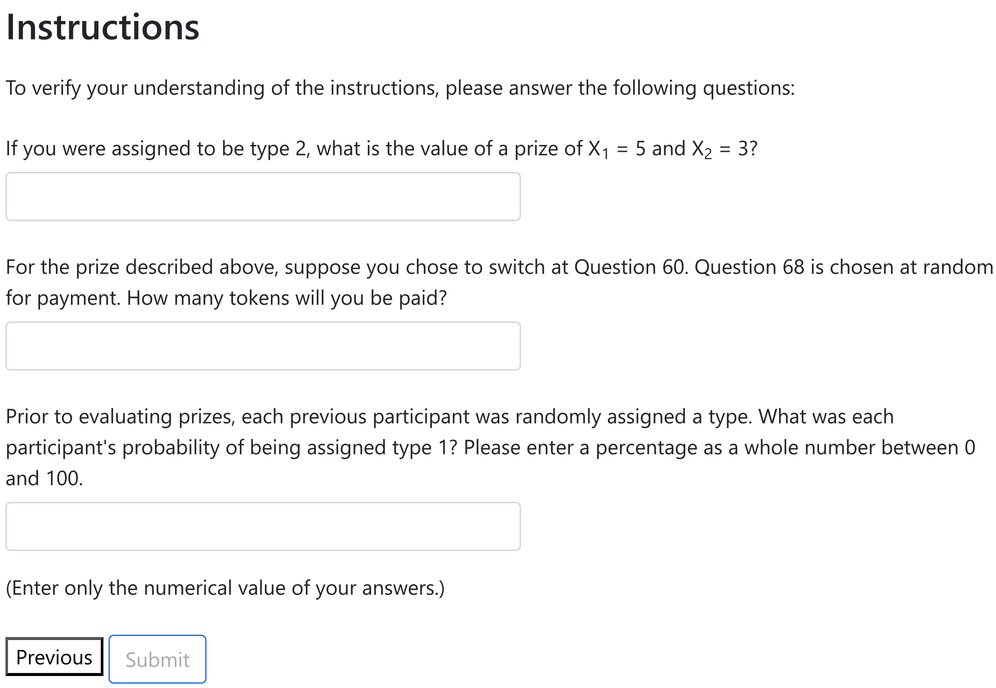

Instructions

To verify your understanding of the instructions, please answer the following questions:

If you were assigned to be type 2, what is the value of a prize of X_1 = 5 and X_2 = 3?

For the prize described above, suppose you chose to switch at Question 60. Question 68 is chosen at random for payment. How many tokens will you be paid?

Prior to evaluating prizes, each previous participant was randomly assigned a type. What was each participant's probability of being assigned type 1? Please enter a percentage as a whole number between 0 and 100.

(Enter only the numerical value of your answers.)

Previous Submit

Figure B.10: Consumer's Instructions (v)

Actions

Round 1

You are **type 1**. Your weights are:

w_1	1
w_2	0.1

A previous participant evaluated the prize for this round and sent a rating of **5**.

The previous participant was **type 2**.

For this prize, X_1 took the number **5**. You will learn the number X_2 took at the end of the round.

As a reminder, ratings are on a scale of 1 to 5, where 1 is the worst rating and 5 is the best.
The list of questions is included below for your reference.

At which question will you switch to Option B?

Be sure to enter a number between 1 and 110 inclusive.

Submit

Question #		Option A		Option B
Question 1:	Would you rather have	the prize	or	0.1 tokens
Question 2:	Would you rather have	the prize	or	0.2 tokens
Question 3:	Would you rather have	the prize	or	0.3 tokens
⋮	⋮	⋮	⋮	⋮
Question 108:	Would you rather have	the prize	or	10.8 tokens
Question 109:	Would you rather have	the prize	or	10.9 tokens
Question 110:	Would you rather have	the prize	or	11.0 tokens

Figure B.11: Consumer's Actions (i)

Round 1 Outcome

You are **type 1**. Your weights are:

w_1	1
w_2	0.1

In this round, the prize was made up of:

x_1	5
x_2	7

The value of this prize to you was 5.7 tokens.
If the prize is chosen for payment at the end of the experiment, you will earn 5.7 tokens.

Next Round

Figure B.12: Consumer's Actions (ii)

B.1.4 Earnings

Thank you

This concludes the experiment. Thank you for participating.

Round 5 and Question 88 have been chosen for your payment.
You chose to switch at Question 1.
You will receive an additional payment of 8.80 tokens to be converted to dollars at the rate described in the instructions.

Your total payment is $7.40.
This comprises the fixed payment of $3.00 and a bonus payment of $4.40.

Please enter the following code into Prolific 1307BDF6.

Figure B.13: Earnings Screen

Appendix C

Free and Open-Source Software: Coordination and Competition

C.1 Proofs

C.1.1 Proof of Lemma 1

To show Lemma 1, I first show the following Lemma:

Lemma 2. *The indifferent user is given by* $\bar{x} = \frac{v_p - p_p + p_o - \gamma + t(1-l)}{t(2-l)-\gamma}$*.* $\frac{\partial \bar{x}}{\partial \gamma} < 0$ *and* $\frac{\partial \bar{x}}{\partial l} < 0$*.*

This is because of the backward induction nature of the game, where the proprietary firm sets prices before users decide which software to use.

Proof of Lemma 2.
To find the indifferent user, note that the indifferent user derives the same utility from the consumption of the proprietary software and the FOSS. Therefore,

$$\begin{aligned} v_p - p_p - t\bar{x} &= \gamma(1-\bar{x}) - t|1 - l(1-\bar{x}) - \bar{x}| - p_o \\ &= (1-\bar{x})(\gamma - t(1-l)) - p_o \\ \bar{x} &= \frac{v_p - p_p + p_o - \gamma + t(1-l)}{t(2-l) - \gamma} \end{aligned}$$

and this is represented by (3.1).

By Restriction 2, $\bar{x} \in (0,1)$ and an interior profit-maximising solution implies

$$t(2-l) - \gamma > 0 \tag{c.1}$$

$$v_P - p_p + p_o - \gamma + t(1-l) > 0 \tag{c.2}$$

$$p_p + t - v_p - p_o > 0 \tag{c.3}$$

Therefore,

$$\begin{aligned} \frac{\partial \bar{x}}{\partial \gamma} &= \frac{v_p - t + p_o - p_p}{(t(2-l)-\gamma)^2} \\ \frac{\partial \bar{x}}{\partial l} &= \frac{t(v_p - t + p_o - p_p)}{(t(2-l)-\gamma)^2} \end{aligned}$$

applying (c.3), both $\frac{\partial \bar{x}}{\partial \gamma}$ and $\frac{\partial \bar{x}}{\partial l}$ are < 0.

This concludes the proof. $\square$

I now show Lemma 1.

Proof of Lemma 1.
To find the equilibrium price of the proprietary software, consider the profit function of the firm.

$$\begin{aligned}\pi_p &= p_p \cdot \bar{x} \\ &= p_p \frac{v_p - p_p + p_o - \gamma + t(1-l)}{t(2-l) - \gamma}\end{aligned}$$

Notice that for an interior profit maximising solution, $t(2-l) - \gamma > 0$ is required. This ensures that $\frac{\partial^2 \pi_p}{\partial p_p^2} < 0$, and implies that $v_p - p_p + p_o - \gamma + t(1-l) > 0$.

Therefore, $p_p^* = \frac{v_p + p_o - \gamma + t(1-l)}{2}$ when Restriction 2 holds. Restriction 2 implies that

$$t(2-l) - \gamma > 0$$

$$v_p + p_o + t(1-l) - \gamma > 0 \tag{c.4}$$

$$t(3-l) - p_o - v_p - \gamma > 0 \tag{c.5}$$

I now show that p_p^* increases in v_p and p_o, and decreases in γ and l.

$$\begin{aligned}\frac{\partial p_p^*}{\partial v_p} &= \frac{1}{2} > 0 \\ \frac{\partial p_p^*}{\partial p_o} &= \frac{1}{2} > 0 \\ \frac{\partial p_p^*}{\partial \gamma} &= -\frac{1}{2} < 0 \\ \frac{\partial p_p^*}{\partial t} &= -\frac{t}{2} < 0\end{aligned}$$

This concludes the proof. $\square$

C.1.2 Proof of Corollary 1

Proof of Corollary 1.
To show Corollary 1, recall that $\pi_p = p_p \cdot \bar{x}$. Considering (3.1) and p_p^*,

$$
\begin{aligned}
\pi_p &= \frac{[v_p + p_o - \gamma + t(1-l)]^2}{4(t(2-l)-\gamma)} \\
\frac{\partial \pi_p}{\partial \gamma} &= \frac{v_p + p_o + t(1-l) - \gamma}{4(t(2-l)-\gamma)^2}(v_p + p_o + \gamma - t(3-l)) \\
\frac{\partial \pi_p}{\partial l} &= t\frac{v_p + p_o + t(1-l) - \gamma}{4(t(2-l)-\gamma)^2}(v_p + p_o + \gamma - t(3-l)) = t\frac{\partial \pi_p}{\partial \gamma}
\end{aligned}
$$

First notice that $t(2-l) - \gamma > 0$ (c.1) and $v_p + p_o + t(1-l) - \gamma > 0$ (c.2). This means that the sign of $\frac{\partial \pi_p}{\partial \gamma}$ depends on the sign of $v_p + p_o + \gamma - t(3-l)$. From (c.5), $v_p + p_o + \gamma - t(3-l) < 0$. Therefore $\frac{\partial \pi_p}{\partial \gamma} < 0$, and by extension $\frac{\partial \pi_p}{\partial l} < 0$.

This concludes the proof. □

C.1.3 Proof of Proposition 1

I solve the game by backward induction. Notice that in Section C.1.1, I have solved the last two stages of the game. And this is common to all coordinators. Here, I provide the proof for what happens in the first stage of the game when the coordinator is a Founder.

Proof of Proposition 1.
To proof Proposition 1, I first propose a candidate equilibrium, ruling out all other possible equilibrium. I then show that such an equilibrium satisfies the Restrictions, and maximises the Founder's utility.

Recall that $\pi_o = (\gamma - tl)(1 - \bar{x})$. Considering (3.1), and remembering that $p_o = 0$ when the coordinator is a Founder,

$$\pi_o = \frac{(\gamma - tl)(t(3-l) - v_p - \gamma)}{2(t(2-l)-\gamma)}$$

Taking the first and second order conditions,

$$
\begin{aligned}
\frac{\partial \pi_o}{\partial \gamma} &= \frac{t(t(6-6l+l^2) - 2v_p(1-l)) + \gamma^2 - 2\gamma t(2-l)}{2(t(2-l)-\gamma)^2} \\
&\quad \text{evaluated at } \gamma = 0, \\
&= \frac{t(6-6l+l^2) - 2v_p(1-l)}{2t(2-l)^2}, > 0 \text{ at } t(3 + \frac{l^2}{2(1-l)}) > v_p \\
\frac{\partial^2 \pi_o}{\partial \gamma^2} &= -\frac{2t(1-l)(v_p - t)}{(t(2-l)-\gamma)^3}, < 0 \text{ at } t(2-l) - \gamma > 0.
\end{aligned}
$$

Notice that $\gamma \in [0,1]$ implies the FOSS is only active if $3t + \frac{tl^2}{2(1-l)} > v_p$.

Solving for γ^*

$$\gamma^* = t(2-l) \pm \sqrt{2t(v_p - t)(1-l)}$$

Notice that $\sqrt{2t(v_p - t)(1-l)}$ is real only if $v_p \geq t$, because $l \in (0,1)$. And when $\sqrt{2t(v_p - t)(1-l)}$ is real, it is ≥ 0. Because $t(2-l) > \gamma$ (c.1), I reject $t(2-l) + \sqrt{2t(v_p - t)(1-l)}$ as a candidate for γ^*. Moreover, because $\gamma \in [0,1]$, $\gamma^* = 1$ when $t > \frac{1}{2-l}$ and $v_p \leq \frac{1-2t(2-l)+t^2(6-6l+l^2)}{2t(1-l)} = 3t + \frac{tl^2}{2(1-l)} + \frac{1}{2t(1-l)} - \frac{2-l}{1-l}$. In all other cases, $\gamma^* = t(2-l) - \sqrt{2t(v_p - t)(1-l)}$. Therefore, I propose $\gamma^* = min\{t(2-l) - \sqrt{2t(v_p - t)(1-l)}, 1\}$ as a candidate solution.

I show when this solution satisfies Restriction 1, Consider that the Founder and the indifferent user must have weakly positive utility. The Founder's utility is given by π_o, and this is weakly positive whenever $\gamma \geq tl$. To see that $\gamma^* \geq tl$,

$$\begin{aligned} t(2-l) - \sqrt{2t(v_p - t)(1-l)} &\geq tl \\ 2t(1-l) &\geq \sqrt{2t(v_p - t)(1-l)} \\ 2t(1-l) &\geq (v_p - t) \\ t(3-2l) &\geq v_p \\ \frac{3t - v_p}{2t} &\geq l \end{aligned}$$

and $1 \geq tl \iff l \leq \frac{1}{t}$ when $\gamma^* = 1$.

The indifferent user's utility is $\frac{(\gamma - t(1-l))(t(3-l) - v_p - \gamma)}{2(t(2-l)-\gamma)}$. Because $t(2-l) - \gamma > 0$ (c.1) and $t(3-l) - v_p - \gamma > 0$ (c.5), it remains that the indifferent user's utility is weakly positive if $\gamma \geq t(1-l)$. To see that $\gamma^* \geq t(1-l)$,

$$\begin{aligned} t(2-l) - \sqrt{2t(v_p - t)(1-l)} &\geq t(1-l) \\ t &\geq \sqrt{2t(v_p - t)(1-l)} \\ t &\geq 2(v_p - t)(1-l) \\ l &\geq \frac{2v_p - 3t}{2(v_p - t)} \end{aligned}$$

and $1 \geq t(1-l) \iff l \geq 1 - \frac{1}{t}$ when $\gamma^* = 1$.

Combining $\frac{3t - v_p}{2t} \geq l$ with $l \geq \frac{2v_p - 3t}{2(v_p - t)}$,

$$\frac{3t - v_p}{2t} \geq \frac{2v_p - 3t}{2(v_p - t)}$$
$$\frac{3t - v_p}{t} \geq \frac{2v_p - 3t}{v_p - t}$$
$$(3t - v_p)(v_p - t) \geq t(2v_p - 3t)$$
$$v_p(2t - v_p) \geq 0$$

implies that $2t \geq v_p$.

Together, this means that when $\gamma^* = t(2-l) - \sqrt{2t(v_p - t)(1-l)}$, $l \in [\frac{2v_p-3t}{2(v_p-t)}, \frac{3t-v_p}{2t}]$ such that $2t \geq v_p$. And $1 - \frac{1}{t} \leq l \leq \frac{1}{t}$, which is true when $t[1, 2]$.

Next, I show when this solution satisfies Restriction 2. It is immediate that the interior candidate solution always satisfies (c.1). To see when it satisfies (c.4),

$$v_p + t(1-l) - t(2-l) + \sqrt{2t(v_p - t)(1-l)} > 0$$
$$v_p - t + \sqrt{2t(v_p - t)(1-l)} > 0$$

This is always true.

To see when it satisfies (c.5),

$$t(3-l) - v_p - t(2-l) + \sqrt{2t(v_p - t)(1-l)} > 0$$
$$t - v_p + \sqrt{2t(v_p - t)(1-l)} > 0$$
$$2t(1-l) > v_p - t$$
$$t(3 - 2l) > v_p$$

And the corner solution of $\gamma^* = 1$ satisfies these equations when $t(2-l) > 1$, $v_p + t(1-l) > 1$, and $t(3-l) - v_p > 1$.

Given these Restrictions, I now show that the candidate solution is utility maxinimising. For the candidate solution to be utility maximising, it must be that $\frac{\partial \pi_o}{\partial \gamma} > 0$ at $\gamma = 0$ and $\frac{\partial^2 \pi_o}{\partial \gamma^2} < 0$.

At $\gamma = 0$, $\frac{\partial \pi_o}{\partial \gamma}$ is $\frac{t(6-6l+l^2)-2v_p(1-l)}{2t(2-l)^2}$. This is positive whenever $t(3 + \frac{l^2}{2(1-l)}) > v_p$. Recall from (c.5) that $t + \sqrt{2t(v_p - t)(1-l)} > v_p$. I now show directly that $t(3 + \frac{l^2}{2(1-l)}) \geq t + \sqrt{2t(v_p - t)(1-l)}$:

$$t(3 + \frac{l^2}{2(1-l)}) \geq t + \sqrt{2t(v_p - t)(1-l)}$$
$$t(2 + \frac{l^2}{2(1-l)}) \geq \sqrt{2t(v_p - t)(1-l)}$$

Recall that $t(3-2l) > v_p \iff 2t(1-l) > v_p - t$. This means that $2t(1-l) > \sqrt{2t(v_p-t)(1-l)}$ and showing $t(2+\frac{l^2}{2(1-l)}) \geq 2t(1-l)$ implies that $t(2+\frac{l^2}{2(1-l)}) > \sqrt{2t(v_p-t)(1-l)}$.

$$\begin{aligned} t(2+\frac{l^2}{2(1-l)}) &\geq 2t(1-l) \\ \frac{4(1-l)+l^2}{2(1-l)} &\geq 2(1-l) \\ 4 &\geq 3l \end{aligned}$$

Since $l \in (0,1)$, this is always true.

It is immediate that $\frac{\partial^2 \pi_o}{\partial \gamma^2} < 0$ for the range of possible solutions. This is because it is negative whenever $v_p > t$ is required for a real solution, and (c.1) must be satisfied, $l \in (0,1)$, and $t > 0$.

Therefore, whenever Restriction 2 is satisfied, the candidate solution maximises the utility of the Founder.

I now show that the Founder does not always play $\gamma = 1$ when it is active. For the Founder to always play $\gamma^* = 1$, it must be that $3t + \frac{tl^2}{2(1-l)} + \frac{1}{2t(1-l)} - \frac{2-l}{1-l} \geq 3t + \frac{tl^2}{2(1-l)}$, and $t > \frac{1}{2-l}$. Suppose to a contradiction that the Founder always plays $\gamma^* = 1$.

The first criteria reduces to

$$\begin{aligned} 3t + \frac{tl^2}{2(1-l)} + \frac{1}{2t(1-l)} - \frac{2-l}{1-l} eq 3t + \frac{tl^2}{2(1-l)} \frac{1}{2t(1-l)} - \frac{2-l}{1-l} &\geq 0 \\ \frac{1-2t(2-l)}{t(1-l)} &\geq 0 \\ \frac{1}{2(2-l)} &\geq t \\ \frac{1}{2(2-l)} &> \frac{1}{2-l} \\ \frac{1}{2} &> 1 \end{aligned}$$

A contradiction. This means that there exists a range of v_p and t for which the Founder will play the $t(2-l) - \sqrt{2t(v_p-t)(1-l)}$.

Therefore, I conclude that there exists a unique equilibrium where $\gamma^* = min\{t(2-l) - \sqrt{2t(v_p-t)(1-l)}, 1\}$ maximises the Founder's utility subject to Restriction 1 and 2, which require $l \in [\frac{2v_p-3t}{2(v_p-t)}, \frac{3t-v_p}{2t}]$. In particular, $\gamma^* = 1$ if $v_p \leq 3t + \frac{tl^2}{2(1-l)} + \frac{1}{2t(1-l)} - \frac{2-l}{1-l}$ and $t > \frac{1}{2-l}$, and $l \in [1-\frac{1}{t}, \frac{1}{t}]$.

To see that $\frac{\partial \gamma^*}{\partial l} \leq 0$ and $\frac{\partial \gamma^*}{\partial v_p} \leq 0$, at $\gamma^* = t(2-l) - \sqrt{2t(v_p-t)(1-l)}$,

$$\frac{\partial \gamma^*}{\partial l} = \frac{t(v_p - t - \sqrt{2t(v_p - t)(1-l)})}{\sqrt{2t(v_p-t)(1-l)}}$$
$$\frac{\partial \gamma^*}{\partial v_p} = -\frac{t(1-l)}{\sqrt{2t(v_p-t)(1-l)}}$$

From (c.5), $t - v_p + \sqrt{2t(v_p - t)(1-l)} > 0$, hence $\frac{\partial \gamma^*}{\partial l} < 0$. And $l \in (0,1)$ implies $\frac{\partial \gamma^*}{\partial v_p} < 0$.

When $\gamma^* = 1$, changes to l or v_p does not affect γ^*. Therefore $\frac{\partial \gamma^*}{\partial l} = 0$ and $\frac{\partial \gamma^*}{\partial v_p} = 0$.

This concludes the proof. □

C.1.4 Proof of Corollary 2

Proof of Corollary 2.
To show Corollary 2, recall that

$$\pi_o = \frac{(\gamma - tl)(t(3-l) - v_p - \gamma)}{2(t(2-l) - \gamma)}$$

Notice that the Founder only benefits if $\gamma > tl$ and $t(3-l) - v_p - \gamma > 0$. This means that the Founder only manages the FOSS if $v_p < t(3-2l)$, and the most permissive license that an active FOSS can use is $l < \frac{3t - v_p}{2t} = \bar{l}$. If a license is too permissive, proprietary firms for a monopolist in the market.

This concludes the proof. □

C.1.5 Proof of Proposition 2

Recall that in Section C.1.1 that I have solved the last two stages of the game, hence I rely on these results to proof the first stage of the game.

Proof of Proposition 2.
To proof Proposition 2, I first propose a candidate equilibrium, ruling out all other possible equlibria. I then show that the candidate equilibrium satisfies the Restrictions, and maximises the Altruist's utility.

Recall that $\pi_o^A = \int_{\bar{x}}^1 \gamma(1-\bar{x}) - t|1 - l(1-\bar{x}) - x|dx$. Accounting for (3.1), $\pi_o^A = \frac{(2\gamma - t(1-2l+2l^2))(v_p + \gamma - t(3-l))^2}{8(t(2-l)-\gamma)^2}$.

Taking the first order condition,

$$\frac{\partial \pi_o^A}{\partial \gamma} = \frac{(v_p + \gamma - t(3-l))(\gamma(v_p + t(3-2l) - \gamma) + t(v_p(1+l-2l^2) - t(5-3l-l^2)))}{4(t(2-l)-\gamma)^3}$$
$$\frac{\partial^2 \pi_o^A}{\partial \gamma^2} = \frac{(v_p - t)(4t\gamma(1-l^2) - t^2(17 - 4l - 14l^2 + 4l^3) + 2\gamma v_p + tv_p(5 + 2l - 6l^2))}{4(t(2-l)-\gamma)^4}$$

First, I find the candidate γ^A. Solving $\frac{\partial \pi_o^A}{\partial \gamma} = 0$,

$$\gamma^A = \begin{cases} t(3-l) - v_p \\ \frac{t(3-2l)+v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2} \\ \frac{t(3-2l)+v_p+\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2} \end{cases}$$

I eliminate the first candidate because it violates $t(3-l) - v_p - \gamma > 0$ (c.5). I eliminate the third candidate because $v_p + t(1-l) - \gamma > 0$,

$$\begin{aligned} v_p + t(1-l) - \frac{t(3-2l)+v_p+\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2} &> 0 \\ \frac{v_p - t - \sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2} &> 0 \\ (v_p-t)^2 > (v_p-t)(11t-8tl^2+v_p) \\ 8tl^2 &> 12t \\ 2l^2 &> 3 \end{aligned}$$

This is not possible because $l \in (0,1)$. This leaves $\frac{t(3-2l)+v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2}$ as the only candidate. Since $(11t - 8tl^2 + v_p) > 0$, a real solution only exists if $v_p > t$.

I now show that at the candidate $\gamma^A = \frac{t(3-2l)+v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2}$ that $\frac{\partial \pi_o^A}{\partial \gamma} < 0$. Notice that $\frac{\partial \pi_o^A}{\partial \gamma} > 0$ only if $v_p < t\frac{5-2l^2}{2}$ and $\gamma > \frac{t(t(17-4l-14l^2+4l^3)-v_p(5+2l-6l^2))}{2(v_p+2(1-l^2)t)}$, and $\frac{\partial \pi_o^A}{\partial \gamma} < 0$ whenever $v_p > t\frac{5-2l^2}{2}$ or $\gamma < \frac{t(t(17-4l-14l^2+4l^3)-v_p(5+2l-6l^2))}{2(v_p+2(1-l^2)t)}$. I show that the candidate $\gamma^A < \frac{t(t(17-4l-14l^2+4l^3)-v_p(5+2l-6l^2))}{2(v_p+2(1-l^2)t)}$ when $v_p < t\frac{5-2l^2}{2}$. Therefore there is no scenario where $\frac{\partial^2 \pi_o^A}{\partial \gamma^2} > 0$.

$$\begin{aligned} \frac{t(3-2l)+v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2} &< \frac{t(t(17-4l-14l^2+4l^3)-v_p(5+2l-6l^2))}{2(v_p+2(1-l^2)t)} \\ \frac{(v_p-t)(11t-8tl^2+v_p)}{v_p+2t(1-l^2)} &< \sqrt{(v_p-t)(11t-8tl^2+v_p)} \\ \frac{(v_p-t)(11t-8tl^2+v_p)}{(v_p+2t(1-l^2))^2} &< 1 \\ (v_p-t)(11t-8tl^2+v_p) &< (v_p+2t(1-l^2))^2 \\ t(3-2l^2)(2v_p+t(-5+2l^2)) &< 0 \text{ notice that } 3-2l^2>0 \\ 2v_p - t(5-2l^2) &< 0 \end{aligned}$$

Therefore, whenever $v_p < t\frac{5-2l^2}{2}$, $\gamma^A < \frac{t(t(17-4l-14l^2+4l^3)-v_p(5+2l-6l^2))}{2(v_p+2(1-l^2)t)}$. And I conclude that there is no scenario where $\frac{\partial^2 \pi_o^A}{\partial \gamma^2} > 0$ subject to my Restrictions. And the candidate equilibrium maximises the Altruist's utility.

However, notice that there exists some range where $\gamma^A = 1$. This upper bound is reached when $\frac{t(3-2l)+v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2} \geq 1 \iff v_p \leq \frac{1+t^2(5-3l-l^2)-t(3-2l)}{1+t(1+l-2l^2)}$ and $t > \frac{1}{2-l}$.

I conclude that the Altruists is always selects a positive level of coordination, $\gamma^A > 0$. And $\gamma^A = min\{\frac{t(3-2l)+v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2}, 1\}$.

To see that $\frac{\partial\gamma^A}{\partial v_p} \leq 0$, at $\frac{t(3-2l)+v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}}{2}$,

$$\frac{\partial\gamma^A}{\partial v_p} = \frac{1}{2} - \frac{v_p + t(5-4l^2)}{2\sqrt{(v_p-t)(11t-8tl^2+v_p)}}$$

To see that

$$\begin{aligned}
\frac{1}{2} - \frac{v_p + t(5-4l^2)}{2\sqrt{(v_p-t)(11t-8tl^2+v_p)}} &< 0 \\
\sqrt{(v_p-t)(11t-8tl^2+v_p)} &< v_p + t(5-4l^2) \\
(v_p-t)(11t-8tl^2+v_p) &< (v_p + t(5-4l^2))^2 \\
0 &< t^2(3-2l^2)^2
\end{aligned}$$

this is always true because $t > 0$ and $l \in (0,1)$.

And when $\gamma^A = 1$, changes to v_p does not affect γ^A. Therefore, $\frac{\partial\gamma^A}{\partial v_p} \leq 0$.

This concludes the proof. □

C.1.6 Proof of Corollary 4

Proof of Corollary 4.
To see this, I compare γ^* with γ^A whenever $\gamma^* < 1$ and $\gamma^A < 1$.

$$\begin{aligned}
\gamma^* &> \gamma^A \\
t(2-l) - \sqrt{2t(v_p-t)(1-l)} &> \frac{t(3-2l) + v_p - \sqrt{(v_p-t)(v_p+11t-8tl^2)}}{2} \\
t - 2\sqrt{2t(v_p-t)(1-l)} &> v_p - \sqrt{(v_p-t)(v_p+11t-8tl^2)} \\
t - v_p + \sqrt{(v_p-t)(v_p+11t-8tl^2)} &> 2\sqrt{2t(v_p-t)(1-l)}
\end{aligned}$$

Notice that $\sqrt{(v_p-t)(v_p+11t-8tl^2)} > \sqrt{2t(v_p-t)(1-l)}$ and $t - v_p + \sqrt{2t(v_p-t)(1-l)} > 0$. Thus, both the left and right side of the equation are positive. Therefore,

$$
\begin{aligned}
(t - v_p + \sqrt{(v_p - t)(v_p + 11t - 8tl^2)})^2 &> 8t(v_p - t)(1 - l) \\
(v_p - t)(t + v_p - \sqrt{(v_p - t)(v_p + 11t - 8tl^2)} + 4tl(1 - l)) &> 0 \\
t + v_p + 4tl(1 - l) &> \sqrt{(v_p - t)(v_p + 11t - 8tl^2)} \\
4t((3 + 2l - 8l^3 + 4l^4)t - 2v_p(1 - l)) &> 0 \\
(3 + 2l - 8l^3 + 4l^4)t &> 2v_p(1 - l)
\end{aligned}
$$

Notice that $l \geq \frac{2v_p - 3t}{2(v_p - t)}$ is necessary for the interior solution for the self-interested Founder. Hence, $2v_p(1 - l) \leq t(3 - 2l)$, and

$$
\begin{aligned}
(3 + 2l - 8l^3 + 4l^4)t &> t(3 - 2l) \\
4l - 8l^3 + 4l^4 &> 0 \\
1 - l^2 + l^3 &> 0
\end{aligned}
$$

And this is true for all $l \in [0.1]$. Therefore,

$$
(3 + 2l - 8l^3 + 4l^4)t > t(3 - 2l) \geq 2v_p(1 - l)
$$

And indeed $t(2 - l) - \sqrt{2t(v_p - t)(1 - l)} > \frac{t(3-2l)+v_p-\sqrt{(v_p-t)(v_p+11t-8tl^2)}}{2}$ is always true as long as an interior solution is played by both the self-interested Founder and the Altruist.

This concludes the proof. □

C.1.7 Proof of Proposition 3

I solve the game by backward induction. Notice that in Section C.1.1, I have solved the indifferent user, and the pricing strategy of the proprietary firm. Hence, to show Proposition 3 I first show Lemma 3.

Lemma 3. *The pricing strategy of the profit-driven Manager is* $p_o^M = \frac{t(3-l)-\gamma-v_p}{3}$ *and the proprietary firm is* $p_p^M = \frac{v_p-2\gamma+t(3-2l)}{3}$.

Proof of Lemma 3.
To show Lemma 3, recall that both the Manager and the proprietary firm set prices simultaneously, and the firm's pricing strategy is given in Lemma 1. I solve for the Manager's pricing strategy, then show the equilibrium pricing strategy.

Recall that the Manager's profit function is given by $\pi_o^M = p_o \cdot (1 - \bar{x})$, where $\bar{x}$ is given in (3.1).

$$\pi_o^M = p_o \frac{t + p_p - p_o - v_p}{t(2-l) - \gamma}$$

and profits are concave in price. Therefore,

$$p_o^M = \frac{p_p + t - v_p}{2}$$

accounting for $p_p = \frac{v_p + p_o - \gamma + t(1-l)}{2}$,

$$p_o^M = \frac{t(3-l) - \gamma - v_p}{3}$$
$$p_p^M = \frac{v_p - 2\gamma + t(3-2l)}{3}$$

This concludes the proof. □

Further, accounting for the pricing strategies, Restriction 1 implies that

$$t(2-l) - \gamma > 0 \tag{c.6}$$
$$v_p + t(3-2l) - 2\gamma > 0 \tag{c.7}$$
$$t(3-l) - v_p - \gamma > 0 \tag{c.8}$$

Using these pricing strategies, I now show Proposition 3.

Proof of Proposition 3.
I first propose a candidate equilibrium which maximises the Manager's profits. I then show how such an equilibrium satisfies the Restrictions of the model.

Accounting for the pricing strategies, the indifferent user is $\bar{x} = \frac{v_p + t(3-2l) - 2\gamma}{3(t(2-l)-\gamma)}$ and the Managers profit is now

$$\pi_o^M = \frac{(t(3-l) - v_p - \gamma)^2}{9(t(2-l) - \gamma}$$
$$\frac{\partial \pi_o^M}{\partial \gamma} = \frac{(t(3-l - v_p - \gamma)(\gamma - v_p - t(1-l))}{9(t(2-l) - \gamma)^2}$$
$$\frac{\partial^2 \pi_o^M}{\partial \gamma^2} = \frac{2(v_p - t)^2}{9(t(2-l) - \gamma)^3}, > 0$$

First, observe that within the restrictions of the model, (c.6) and (c.8) imply that it is always profitable for the Manager to coordinate the FOSS.

Second, observe that the second order condition implies that the FOSS profit is not concave in the level of coordination. Therefore, whenever $\frac{\partial \pi_o^M}{\partial \gamma} > 0$, the Manager prefers the highest level of coordination. Since $t(3-l) - \gamma - v_p > 0$ and $t(2-l) - \gamma > 0$, $\frac{\partial \pi_o^M}{\partial \gamma} > 0 \iff \gamma - v_p - t(1-l) > 0$ and 0 with equality.

Suppose first that $v_p < t$. When this is true, I show that $\frac{\partial \pi_o^M}{\partial \gamma} > 0$. From $\gamma - v_p - t(1-l) > 0$,

$$\begin{aligned} \gamma - v_p - t(1-l) > 0 \text{ adding } t(2-l) - \gamma > 0, \\ t(2-l) - \gamma + \gamma - v_p - t(1-l) > 0 \\ t > v_p \end{aligned}$$

Therefore, whenever $v_p < t$, $\frac{\partial \pi_o^M}{\partial \gamma} > 0$ and the Manager prefers the highest level of coordination, $\gamma^M = 1$.

Suppose next that $v_p \geq t$. When this is true, I show that $\frac{\partial \pi_o^M}{\partial M} < 0$. Suppose to a contradiction that $\gamma - v_p - t(1-l) > 0$.

$$\begin{aligned} \gamma - v_p - t(1-l) > 0 \\ t - v_p > t(2-l) - \gamma \end{aligned}$$

the left hand side is weakly negative, and the right hand side is positive. Hence, a contradiction. This implies that $\frac{\partial \pi_o^M}{\partial \gamma} < 0$, and the Manager is still active, but prefers the lowest level of coordination.

I now consider the Restrictions of the model. Recall that for Restriction 1 the indifferent user and the extreme user located at 1 must have weakly positive utility, and for Restriction 2 the equations (c.6), (c.7) and (c.8) hold.

The indifferent user's utility is given by

$$u_o(\bar{x}) = (1-\bar{x})(\gamma - t(1-l)) - p_o = \frac{(t(3-l) - v_p - \gamma)(2\gamma - t(3-2l))}{3(t(2-l) - \gamma)}$$

, and the for the extreme user $u_o(1) = \frac{2(t(3-l)-v_p-\gamma)(\gamma-t)}{3(t(2-l)-\gamma))}$. For the indifferent user to receive weakly positive utility, $\gamma \geq \frac{t(3-2l)}{2}$, and for the extreme user, $\gamma \geq t$.

Because $\gamma \in [0,1]$, Restriction 1 implies that $1 \geq \frac{t(3-2l)}{2}$ or $l \geq \frac{3t-2}{2t}$ is a sufficient condition for the Manager to be active and cover the market.

I now show when the Restrictions are satisfied by the candidate equilibrium.

Suppose first that $v_p < t$, and the candidate equilibrium is $\gamma^M = 1$. The following equations need to be satisfied.

$$
\begin{aligned}
1 &\geq t \\
1 &\geq \frac{t(3-2l)}{2} \\
t(2-l) &> 1 \\
v_p + t(3-2l) &> 2 \\
t(3-l) - v_p &> 1
\end{aligned}
$$

Observe that when $v_p < t$ then whenever $t(2-l) > 1$ is satisfied, both of $v_p + t(3-2l) > 2$ and $t(3-l) - v_p > 1$ are satisfied. Notice that $1 \geq t$ is consistent whenever

$$
\begin{aligned}
&t(2-l) > 1 \geq t \\
&t(2-l) > t \\
&t(1-l) > 0
\end{aligned}
$$

which is always true because $t > 0$ and $l \in (0,1)$.

Finally, $1 \geq \frac{t(3-2l)}{2} \iff 2 \geq t(3-2l)$. And this is consistent whenever

$$
\begin{aligned}
v_p + t(3-2l) > 2 \geq t(3-2l) \\
v_p + t(3-2l) > t(3-2l) \\
v_p > 0
\end{aligned}
$$

And this is implicitly true, or the proprietary firm is inactive. See Etzion and Pang (2014) for details.

Therefore, whenever Restriction 1 and 2 are satisfied and $v_p < t$, the Manager chooses $\gamma^M = 1$.

Suppose next that $v_p \geq t$, and the candidate equilibrium is the smallest level of γ that satisfies the following equations.

$$
\begin{aligned}
\gamma &\geq t \\
\gamma &\geq \frac{t(3-2l)}{2} \\
t(2-l) &> \gamma \\
v_p + t(3-2l) &> 2\gamma \\
t(3-l) - v_p &> \gamma
\end{aligned}
$$

Notice first that $\gamma \geq t$ and $\gamma \geq \frac{t(3-2l)}{2}$ imply that for the candidate equilibrium to satisfy Restriction 1, $\gamma^M = max\{t, t(1.5-l)\}$. Observe that $\gamma^M = t$ if $l \geq 0.5$ and $\gamma^M = t(1.5-l)$ if $l < 0.5$.

Using this, I show that the candidate equilibrium satisfies Restriction 2.

Suppose first that $l \geq 0.5$ then (c.6), (c.7) and (c.8) become

$$\begin{aligned} t(2-l) &> t \\ v_p + t(3-2l) &> 2t \\ t(3-l) - v_p &> t \end{aligned}$$

which are all satisfied because

$$\begin{aligned} t(1-l) &> 0 \\ v_p + t(1-l) &> 0 \\ t(2-l) - v_p &> 0 \end{aligned}$$

are always true because $t > 0$ and $l \in [0.5, 1)$.

Suppose next that $l < 0.5$ then (c.6), (c.7) and (c.8) become

$$\begin{aligned} t(2-l) &> t(1.5-l) \\ v_p + t(3-2l) &> t(3-2l) \\ t(3-l) - v_p &> t(1.5-l) \end{aligned}$$

which are all satisfied because

$$\begin{aligned} 0.5t &> 0 \\ v_p &> 0 \\ 1.5t - v_p &> 0 \end{aligned}$$

which are always true because $t > 0$.

Therefore, following Restrictions 1 and 2, $l \geq \frac{3t-2}{2t}$ is a sufficient condition for the Manager to be active. And when $v_p \geq t$, $\gamma^M = \max\{t, t(1.5-l)\}$, when $v_p < t$, $\gamma^M = 1$.

This concludes the proof. $\square$

C.1.8 Proof of Corollary 5

Before embarking on the surplus analysis, I first provide the expressions for consumer surplus, CS_i, and mean consumer surplus cs_i where $i \in \{o, p, b\}$ represent the FOSS, the proprietary firm, and both respectively.

$$\begin{aligned}
CS_o &= \int_{\bar{x}}^{1} \gamma(1-\bar{x}) - t|(1-l(1-\bar{x})-x| - p_o dx \\
&= \frac{(1-\bar{x})^2(2\gamma + 2tl(1-l) - t)}{2} - p_o(1-\bar{x}) \\
cs_o &= \frac{(1-\bar{x})(2\gamma + 2tl(1-l) - t)}{2} - p_o \\
CS_p &= \frac{\bar{x}(2v_p - t\bar{x} - 2p_p)}{2} \\
cs_p &= \frac{2v_p - t\bar{x} - 2p_p}{2} \\
CS_b &= CS_o + CS_p
\end{aligned}$$

Proof of Corollary 5.
I evaluate CS_o and cs_o at the Founder equilibrium, where

$$\begin{aligned}
p_o &= 0 \\
\bar{x} &= \frac{v_p - \gamma^* + t(1-l)}{2(t(2-l)-\gamma^*)} \\
\gamma^* &= t(2-l) - \sqrt{2t(v_p - t)(1-l)}
\end{aligned}$$

Therefore, at $\gamma^* = t(2-l) - \sqrt{2t(v_p-t)(1-l)}$

$$\begin{aligned}
CS_o &= \frac{(t(3-2l^2) - 2\sqrt{2t(v_p-t)(1-l)})(v_p + t(1-2l) - 2\sqrt{2t(v_p-t)(1-l)})}{16t(1-l)} \\
\frac{\partial CS_o}{\partial l} &< 0
\end{aligned}$$

and

$$\begin{aligned}
cs_o &= \frac{\sqrt{2}(t(3-2l^2) - 2\sqrt{2t(v_p-t)(1-l)})(v_p - t - \sqrt{2t(v_p-t)(1-l)})}{8\sqrt{t(v_p-t)(1-l)}} \\
\frac{\partial cs_o}{\partial l} &= t\frac{\sqrt{2}(v_p-t)(1+4l-6l^2) - 16(1-l)l\sqrt{t(v_p-t)(1-l)}}{16(1-l)\sqrt{t(v_p-t)(1-l)}}
\end{aligned}$$

$\frac{\partial cs_o}{\partial l} > 0 \iff l < \frac{2+\sqrt{10}}{6}$ and $v_p > t(1 + \frac{2^8 l^2 (1-l)^3}{(1+4l-6l^2)})$.

This concludes the proof. □

C.1.9 Proof of Corollary 6

Proof of Corollary 6.
I evaluate CS_p and cs_p at the Founder equilibrium, where

$$
\begin{aligned}
p_o &= 0 \\
\bar{x} &= \frac{v_p - \gamma^* + t(1-l)}{2(t(2-l)-\gamma^*)} \\
\gamma^* &= t(2-l) - \sqrt{2t(v_p - t)(1-l)} \\
p_p^* &= \frac{v_p + p_o - \gamma + t(1-l)}{2}
\end{aligned}
$$

Then,

$$
\frac{\partial CS_p}{\partial l} = \frac{\sqrt{t(v_p-t)(1-l)}(\sqrt{2}v_p(1-l) + 2\sqrt{2}t(1-l)^2 - \sqrt{t(v_p-t)(1-l)})}{16t(1-l)^3} > 0
$$

Evaluating cs_p in equilibrium when there is a self-interested Founder,

$$
\begin{aligned}
cs_p &= \frac{1}{8}(4v_p + 2t - \frac{\sqrt{2}(5-4l)t}{\sqrt{t(v_p-t)(1-l)}}(v_p - t)) \\
\frac{\partial cs_p}{\partial l} &= \frac{(3-4l)\sqrt{t(v_p-t)(1-l)}}{8\sqrt{2}(1-l)^2}
\end{aligned}
$$

$\frac{\partial cs_p}{\partial l} > 0 \iff l < \frac{3}{4}$ and $\frac{\partial cs_p}{\partial l} = 0$ when $l = \frac{3}{4}$.

This concludes the proof.

□

C.1.10 Proof of Corollary 7

Proof of Corollary 7.

I evaluate CS_o, cs_o, CS_p, cs_p at the Altruists equilibrium, where

$$
\begin{aligned}
p_o &= 0 \\
\bar{x} &= \frac{v_p - \gamma^* + t(1-l)}{2(t(2-l)-\gamma^*)} \\
\gamma^* &= \frac{t(3-2l) + v_p - \sqrt{(v_p - t)(11t - 8tl^2 + v_p)}}{2} \\
p_p^* &= \frac{v_p + p_o - \gamma + t(1-l)}{2}
\end{aligned}
$$

Then,

$$
\begin{aligned}
CS_o = (v_p + 2t(1-l^2) - \sqrt{(v_p-t)(v_p+11t-8tl^2)})\times \\
\frac{(3(v_p-t) - \sqrt{(v_p-t)(v_p+11t-8tl^2)})^2}{8(v_p - t - \sqrt{(v_p-t)(v_p+11t-8tl^2)})^2} \\
\frac{\partial CS_o}{\partial l} \leq 0
\end{aligned}
$$

and

$$cs_o = (v_p + 2t(1-l^2) - \sqrt{(v_p-t)(v_p+11t-8tl^2)})\times$$
$$\frac{(3(v_p-t) - \sqrt{(v_p-t)(v_p+11t-8tl^2)})}{8(v_p-t-\sqrt{(v_p-t)(v_p+11t-8tl^2)})}$$
$$\frac{\partial cs_o}{\partial l} \leq 0$$
$$CS_p = \frac{(v_p-t+\sqrt{(v_p-t)(v_p+11t-8tl^2)})}{8(t+\sqrt{(v_p-t)(v_p+11t-8tl^2)}-v_p)^2}\times$$
$$[(t^2(13-8l^2)+4v_p(\sqrt{(v_p-t)(v_p+11t-8tl^2)})-$$
$$t(v_p(9-8l^2)+\sqrt{(v_p-t)(v_p+11t-8tl^2)})))]$$
$$\frac{\partial CS_p}{\partial l} > 0$$

and

$$cs_p = \frac{1}{4(t+\sqrt{(v_p-t)(v_p+11t-8tl^2)}-v_p)}\times$$
$$[t^2(13-8l^2)+4v_p(\sqrt{(v_p-t)(v_p+11t-8tl^2)})- \quad \frac{\partial cs_p}{\partial l} \quad > 0$$
$$t(v_p(9-8l^2)+\sqrt{(v_p-t)(v_p+11t-8tl^2)}))]$$

This concludes the proof.

□

C.1.11 Proof of Corollary 8

Proof of Corollary 8.

To show Corollary 8, I first evaluate total consumer surplus, CS_b at $\bar{x}$, (3.1), and p_p^*, (3.2), setting $p_o = 0$. I then take the first derivative of the resulting consumer surplus and total surplus with respect to γ. I evaluate this at γ^A, and show that this is positive whenever $\gamma^A < 1$.

$$\frac{\partial CS_b}{\partial \gamma} = \frac{2(v_p-t)}{(t-v_p+\sqrt{(v_p-t)(11t-8tl^2-v_p)})^3}\times$$
$$[tl^2(3v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}-3t)$$
$$-t(4v_p-\sqrt{(v_p-t)(11t-8tl^2+v_p)}-4t)$$
$$+v_p(\sqrt{(v_p-t)(11t-8tl^2+v_p)}-v_p)+t^2]$$

Recall from (c.5) that $t(3-l) > v_p+\gamma$, so $0 > 3(v_p-t) - \sqrt{(v_p-t)(v_p+11t-8tl^2)}$, and from (c.1), $t(2-l)-\gamma > 0$ implies $t+\sqrt{(v_p-t)(v_p+11t-8tl^2)}-v_p > 0$, and the denominator is positive. Therefore I only need to show

$$t(3v_p - 3t - \sqrt{(v_p - t)(11t - 8tl^2 + v_p)})(l^2 - 1) + 2t^2 +$$
$$v_p(\sqrt{(v_p - t)(11t - 8tl^2 + v_p)} - v_p - t) > 0$$

And $t(3v_p - 3t - \sqrt{(v_p - t)(11t - 8tl^2 + v_p)})(l^2 - 1) > 0$ because $l^2 < 1$ and $2t^2 > 0$ because $t > 0$. Finally, to see

$$\begin{aligned}
v_p(\sqrt{(v_p - t)(11t - 8tl^2 + v_p)} - v_p - t) &> 0 \; v_p > t > 0 \\
\sqrt{(v_p - t)(11t - 8tl^2 + v_p)} - v_p - t &> 0 \\
11t - 8tl^2 + v_p &> v_p - t \\
12t &> 8tl^2
\end{aligned}$$

this is always true because $l \in (0, 1)$.

Therefore, $\frac{\partial CS_b}{\partial \gamma} > 0$ in equilibrium.

The total surplus in the market is given as

$$\begin{aligned}
TS &= CS_b + p_p \cdot \bar{x} \\
\frac{\partial TS}{\partial \gamma} &= \frac{(v_p - t)(t^2 + (2v_p - t)(\sqrt{(v_p - t)(11t - 8tl^2 + v_p)} - v_p))}{(t + \sqrt{(v_p - t)(11t - 8tl^2 + v_p)} - v_p)^3} \\
&> 0 \text{ when evaluated at } \gamma^A.
\end{aligned}$$

This concludes the proof. □

C.1.12 Proof of Corollary 9

Proof of Corollary 9.

I show that profits of both the profit-driven Manager and the proprietary firm are decreasing in the FOSS license permissiveness. To see this, I evaluate the profits individually in equilibrium, taking the first derivative with respect to l.

Recall that

$$p_o^M = \frac{t(3-l)-\gamma-v_p}{3}$$
$$p_p^M = \frac{v_p-2\gamma+t(3-2l)}{3}$$
$$\bar{x} = \frac{v_p+t(3-2l)-2\gamma}{3(t(2-l)-\gamma)}$$
$$\gamma^* = \begin{cases} 1 & \text{if } v_p < t \\ max\{t, t(1.5-l)\} & \text{if } v_p \geq t \end{cases}$$
$$\pi_o^M = \frac{(t(3-l)-v_p-\gamma)^2}{9(t(2-l)-\gamma)}$$
$$\frac{\partial \pi_o^M}{\partial l} = \begin{cases} \frac{(t(3-l)-v_p-1)(1-v_p-t(1-l))}{9(t(2-l)-1)^2} > 0 & \text{if } \gamma^M = 1 \\ \frac{(tl-v_p)(t(2-l)-v_p)}{9t(1-l)^2} < 0 & \text{if } \gamma^M = t \\ 0 & \text{if } \gamma^M = t(1.5-l) \end{cases}$$
$$\pi_p^M = \frac{(v_p+t(3-2l)-2\gamma)^2}{9(t(2-l)-\gamma)}$$
$$\frac{\partial \pi_p^M}{\partial l} = \begin{cases} \frac{t(t(3-2l)-v_p-2)(2+v_p-t(5-2l)))}{9(t(2-l)-1)^2} & \text{if } \gamma^M = 1 \\ -\frac{(t(1-2l)+v_p)(t(3-2l)-v_p)}{9t(1-l)^2} < 0 & \text{if } \gamma^M = t \\ 0 & \text{if } \gamma^M = t(1.5-l) \end{cases}$$

This concludes the proof.

□

C.1.13 Proof of Corollary 10

Proof of Corollary 10.

To show that the total consumer surplus decreases in the permissiveness of the FOSS license, first recall that

$$p_o^M = \frac{t(3-l)-\gamma-v_p}{3}$$
$$p_p^M = \frac{v_p-2\gamma+t(3-2l)}{3}$$
$$\bar{x} = \frac{v_p+t(3-2l)-2\gamma}{3(t(2-l)-\gamma)}$$
$$\gamma^* = \begin{cases} 1 & \text{if } v_p < t \\ max\{t, t(1.5-l)\} & \text{if } v_p \geq t \end{cases}$$

I show that the total consumer surplus in unambiguously decreasing in the permissiveness of the FOSS license, when Restriction 1 and 2 hold.

$$CS_b = CS_o + CS_p$$

$$\frac{\partial CS_b}{\partial l} \begin{cases} < 0 & \text{when } \gamma^M = 1 \\ = \frac{2(t^2(-5+9l-6l^2+l^3)+tv_p(3-2l+l^2)-v_p^2)}{9t(1-l)^2} < 0 & \text{when } \gamma^M = t \text{ and } l \in [0.5, 1) \\ = -\frac{2l(3t-2v_p)^2}{9t} < 0 & \text{when } \gamma^M = t(1.5-l) \end{cases}$$

This concludes the proof.

□

C.1.14 Proof of Corollary 11

Proof of Corollary 11.

The total surplus in the market is given as

$$TS = CS_b + p_p \cdot \bar{x} + p_o \cdot (1 - \bar{x})$$

$$\frac{\partial TS}{\partial \gamma} = \begin{cases} \frac{t^2(7-l^2)+2(-5-2l+l^2)v_p t+2v_p^2(2+l)}{9t^2(1-l)^2} & > 0 \text{ when evaluated at } \gamma^M = t \\ \frac{t^2(33-24l^2)+8(-7+5l^2)v_p t+8v_p^2(3-2l^2)}{9t^2} & > 0 \text{ when evaluated at } \gamma^M = t(1.5-l) \end{cases}$$

This concludes the proof. □

C.1.15 Proof of Proposition 4

Proof of Proposition 4.

To show this, I solve the game by backward induction. Since the last two stages of the game is identical to the main model, the proof for the indifferent user and the proprietary firm's action can be found in Lemma 2 and 1 respectively. Recall that the objective function of the Founder is

$$\begin{aligned} \pi_o &= \gamma(1-\bar{x}) - t(l(1-\bar{x})) \\ &= (\gamma - tl)(\frac{t(3-l)-\gamma-v}{2(t(2-l)-\gamma)}) \\ \frac{\partial \pi_o}{\partial l} &= \frac{-t(\gamma^2 + t((6-4l+l^2)t - 2v) + 2\gamma((-3+l)t+v))}{2(t(2-l)-\gamma)^2} \\ \frac{\partial^2 \pi_o}{\partial l^2} &= \frac{2t^2(v_p - t)(t-\gamma)}{(t(2-l)-\gamma)^3} \end{aligned}$$

The Founder has an incentive to select a somewhat permissive license whenever $\frac{\partial \pi_o}{\partial l} > 0$ at $l = 0$. Evaluating $\frac{\partial \pi_o}{\partial l}$ at $l = 0$ gives

$$\frac{\partial \pi_o}{\partial l} = \frac{-t(\gamma^2 + 2(t-\gamma)(3t - v_p)}{2(t(2-l)-\gamma)^2}$$

if $v_p > \frac{t^2(6-4l+l^2)+\gamma^2-2t\gamma(3-l)}{2(t-\gamma)}$, and this is always larger than 0 if $(t-\gamma)(t(2-l)-\gamma)^2 > 0$.

Otherwise, if $(t-\gamma)(t(2-l)-\gamma)^2 > 0 \iff t > \gamma$, then $\frac{\partial \pi_o}{\partial l} > 0$ whenever $v_p > 0$.

If $v_p < \frac{t^2(6-4l+l^2)+\gamma^2-2t\gamma(3-l)}{2(t-\gamma)}$ and $\gamma < t$, then $\frac{\partial \pi_o}{\partial l} < 0$ and $\frac{\partial \pi_o^2}{\partial l^2} < 0$, and $l = 0$.

This concludes the proof. □

C.1.16 Proof of Proposition 5

Proof of Proposition 5.

This game is identical to the main model, with the following change: users have heterogeneous skill level, and those located at 1 have a higher skill level than those located at 0. I solve this game by backward induction.

First, I look at the decision made by the marginal user.

$$v_p - p_p - t\bar{x} = \frac{(1-\bar{x})(\alpha(1+\bar{x})(\gamma+tl)-2t)}{2}$$

From this, we can find partial derivatives for $\bar{x}$ with respect to p_p and γ.

$$\frac{\partial \bar{x}}{\partial p_p} = \frac{1}{\bar{x}\alpha(\gamma+tl)-2t}$$
$$\frac{\partial \bar{x}}{\partial \gamma} = \frac{\alpha(1-\bar{x}^2)}{2(\bar{x}\alpha(\gamma+tl)-2t)}$$

Solving explicitly for $\bar{x}$,

$$\bar{x} = \frac{2t \pm \sqrt{4t^2+\alpha^2(\gamma+tl)^2-2\alpha(\gamma+tl)(v+t-p)}}{\alpha(\gamma+tl)}$$

Moving to the next stage, we look at the decision made by the firm. Notice that for any real interior solution, this means that $\frac{\partial \bar{x}}{\partial p} < 0$. This implies that $2t > \bar{x}\alpha(\gamma+tl)$. And additionally that we reject the positive of $\bar{x}$, and $\bar{x} = \frac{2t-\sqrt{4t^2+\alpha^2(\gamma+tl)^2-2\alpha(\gamma+tl)(v+t-p)}}{\alpha(\gamma+tl)}$. This means for positive demand, $p < v + t - \frac{\alpha(\gamma+tl)}{2}$ and for a real solution to $\bar{x}$, $p \geq v + t - \frac{4t^2+\alpha^2(\gamma+tl)^2}{4\alpha(\gamma+tl)}$. Hence, for a duopoly it must be that $2t > \alpha(\gamma+tl)$.

Solving for the optimal price,

$$\pi_p = p\bar{x}$$
$$\frac{\partial \pi_p}{\partial p} = \bar{x} + \frac{\partial \bar{x}}{\partial p}p$$

For a solution, $\frac{\partial \pi_p}{\partial p} = 0$ and this implies that $p_p = -\frac{\bar{x}}{\frac{\partial \bar{x}}{\partial p_p}} = \bar{x}(2t - \bar{x}\alpha(\gamma+tl))$.

$$\frac{\partial^2 \pi_p}{\partial p_p^2} = 2\frac{\partial \bar{x}}{\partial p} + p\frac{\partial^2 \bar{x}}{\partial p^2}$$
$$= \frac{3\bar{x}\alpha(\gamma + tl) - 4t}{(2t - \bar{x}\alpha(\gamma + tl))^2} < 0$$

For an interior solution, $3\bar{x}\alpha(\gamma + tl) - 4t < 0$ or $\bar{x} < \frac{4t}{3\alpha(\gamma+tl)}$ or $\gamma < t(\frac{4}{3\bar{x}\alpha} - l)$.

These two conditions together imply that $2t > \alpha(\gamma + tl)$ or that $\alpha < \frac{2t}{\gamma+tl}$.

Together, the three conditions $\alpha < \frac{4t}{3\bar{x}(\gamma+tl)}$, $\alpha < \frac{2t}{\bar{x}(\gamma+tl)}$ and $\alpha < \frac{2t}{\gamma+tl}$ suggest that an interior solution is only possible if α is sufficiently small, bound above by $\frac{2t}{\gamma+tl}$ if $\bar{x} \leq \frac{2}{3}$ and $\frac{4t}{3\bar{x}(\gamma+tl)}$ if $\bar{x} > \frac{2}{3}$. In all other cases, there will be zero demand for the proprietary software. That is, if users are exceptionally skilled, they simply create a unique product on their own, driving out the proprietary firm. Given this, I now turn to the Founder's choice. Observe that $\alpha < \frac{2t}{\bar{x}(\gamma+tl)}$ is the least restrictive condition. Also notice that rewriting $\bar{x} < \frac{4t}{3\alpha(\gamma+tl)}$ and $\bar{x} < \frac{2t}{\alpha(\gamma+tl)}$ suggests that $\bar{x}$ decreases in α.

Moving to the first stage, the decision by the Founder is

$$\pi_o = \frac{\alpha(\gamma - tl)(1 - \bar{x}^2)}{2}$$

And subject to $\gamma \leq 1$, letting λ be the shadow price,

$$\frac{\partial \pi_o}{\partial \gamma} = \frac{\alpha(1 - \bar{x}^2 - 2(\gamma - tl)\bar{x}\frac{\partial \bar{x}}{\partial \gamma})}{2} - \lambda$$
$$= \frac{\alpha t(1 - \bar{x}\alpha l)(1 - \bar{x}^2)}{2t - \bar{x}\alpha(\gamma + tl)} - \lambda$$

$\lambda \geq 0$, $\lambda(1 - \gamma) = 0$, $1 - \gamma \geq 0$.

Notice that for $\frac{\alpha t(1-\bar{x}\alpha l)(1-\bar{x}^2)}{2t-\bar{x}\alpha(\gamma+tl)} - \lambda = 0$, either $\bar{x} = \frac{1}{\alpha l}$ and $\lambda = 0$ or $\lambda = \frac{\alpha t(1-\bar{x}\alpha l)(1-\bar{x}^2)}{2t-\bar{x}\alpha(\gamma+tl)}$. In the first case, $\lambda = 0$ implies that this solution can be true for any γ. In the second case, $\lambda > 0$ implies that $\gamma = 1$.

Observe that the coordinator is only active if $\gamma > tl$. Suppose that $\gamma \neq 1$. Then $\frac{\partial \pi_o}{\partial \gamma} \geq 0 \iff (1 - \bar{x}\alpha l) \geq 0$, and 0 only with equality This means should there be an interior solution for the FOSS, that $\bar{x} = \frac{1}{\alpha l}$.

However, I show that $\bar{x} = \frac{1}{\alpha l}$ leads to a contradiction between an active FOSS $\gamma > tl$ and an active firm $2t - \bar{x}\alpha(\gamma + tl) = 2t - \frac{\alpha(\gamma+tl)}{\alpha l} > 0 \iff tl > \gamma$. Hence, a duopoly cannot exist if $\bar{x} = \frac{1}{\alpha l}$.

Therefore, the optimal solution for the Founder is $\gamma = 1$.

I conclude that a skilled self-interested Founder chooses to develop a FOSS if $1 > tl$, and they prefer that contributions of each developer is unique, $\gamma = 1$.

This concludes the proof. □

C.1.17 Proof of Proposition 6

Proof of Proposition 6.

I use backward induction to solve for the equilibrium. This is done in the following steps: (1) determine the marginal user indifferent between consumption of the proprietary software and the FOSS; (2) determine the marginal developer indifferent between other leisure and the development of the FOSS; (3) solve for the price strategy of the firm; (4) solve for the optimal level of coordination for developers of the FOSS (γ).

In deciding which software to use, users compare between the following utilities, $u_p = v_p - p_p - tx$ and $u_o = \gamma(1-\bar{o}) - t|(1 - l(1-\bar{o}) - x)|$. The demand here arises from the uniform distribution. The indifferent user being $\bar{x}$, and $\bar{o}$ the indifferent developer.

$$\begin{aligned} v_p - p_p - t\bar{x} &= (1-\bar{o})(\gamma + tl) - t(1-\bar{x}) \\ \bar{x} &= \frac{v_p - p_p - (1-\bar{o})(\gamma + tl) + t}{2t} \end{aligned}$$

From Restriction 1, $\bar{x}$ gives the demand for the proprietary software, and $(1-\bar{x})$ the demand of the FOSS. I turn now to the decision of the developers.

Developers decide between contributing to the FOSS and receiving a utility of $w_o = s_o - k|L_o - o| = \beta(1-\bar{x}) - k|(1 - l(1-\bar{o}) - o)|$. Thus the indifferent developer is

$$\begin{aligned} \beta(1-\bar{x}) - k(1 - l(1-\bar{o}) - \bar{o}) = 0 \\ \bar{o} &= \frac{\beta(\bar{x}-1) + k(1-l)}{k(1-l)} \\ &= 1 - \frac{\beta(1-\bar{x})}{k(1-l)} \end{aligned}$$

This means that

$$\begin{aligned} \bar{x} &= \frac{k(1-l)(v+t-p) - \beta(\gamma + tl)}{2tk(1-l) - \beta(\gamma + tl)} \\ \bar{o} &= \frac{2tk(1-l) + \beta(v - p - t(1-l) - \gamma)}{2tk(1-l) - \beta(\gamma + tl)} \end{aligned}$$

We will need $k > \frac{\beta(\gamma+tl)}{2t(1-l)}$ for concavity of the profit function of the firm.

Now we turn to the decision of the firm.

$$\pi_p = p_p \bar{x}$$
$$p_p = \frac{1}{2}(v_p + t - \frac{\beta(\gamma + tl)}{k(1-l)})$$
$$\bar{x} = \frac{k(1-l)(v+t) - \beta(\gamma + tl)}{2(2tk(1-l) - \beta(\gamma + tl))}$$
$$\bar{o} = \frac{4tk^2(1-l)^2 + \beta^2(\gamma + tl) + k\beta(1-l)(v_p + 2tl - 3t - 2\gamma)}{2k(1-l)(2tk(1-l) - \beta(\gamma + tl))}$$
$$\pi_p = \frac{((v+t)k(1-l) - \beta(\gamma + tl))^2}{4k(1-l)(2tk(1-l) - \beta(\gamma + tl))}$$

Finally, we turn our attention to the coordinator of the FOSS. Recall that the Founder is selfish, and is only motivated by maximising its individual utility.

$$\pi_o = \beta(1 - \bar{x}) - kl(1 - \bar{o})$$
$$= \frac{\beta(\beta(2l-1)(\gamma + tl) + k(1-l)(t(3 - 6l + 4l^2) + v_p(2l-1)))}{2(1-l)(2tk(1-l) - \beta(\gamma + lt))}$$
$$\frac{\partial \pi_o}{\partial \gamma} = \frac{\beta^2 k(4tl^2 + (2l-1)(v_p - t))}{2(2tk(1-l) - \beta(\gamma + tl))^2)}$$
$$\frac{\partial^2 \pi_o}{\partial \gamma^2} = \frac{\beta^3 k(4tl^2 + (2l-1)(v_p - t))}{(2tk(1-l) - \beta(\gamma + tl))^3}$$

Observe that both $\frac{\partial \pi_o}{\partial \gamma}$ and $\frac{\partial^2 \pi_o}{\partial \gamma^2} > 0 \iff$ both $2l - 1 > 0$ and $v_p - t > 0$ simultaneously. Hence, the FOSs is only active if there is a sufficiently permissive license when the proprietary software has a high quality, and a sufficiently restrictive license when the proprietary software has a low quality. In either case, the Founder selects $\gamma = 1$ preferring that contributions by each developer is unique.

Notice that β does not influence the outcome of this decision.

This concludes the proof.

$\square$

Bibliography

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

Adem Ait Fonollà. (2022). Survival rate of github projects – an empirical study. Retrieved July 2023, from https://livablesoftware.com/survival-rate-github-projects-empirical/

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

Almeida, D. A., Murphy, G. C., Wilson, G., & Hoye, M. (2017). Do software developers understand open source licenses? *2017 IEEE/ACM 25th International Conference on Program Comprehension (ICPC)*, 1–11.

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

Andersen-Gott, M., Ghinea, G., & Bygstad, B. (2012). Why do commercial companies contribute to open source software? *International journal of information management*, *32*(2), 106–117.

Anderson, S. P., & Renault, R. (2006). Advertising content. *The American Economic Review*, *96*(1), 93–113.

André Staltz. (2019). Software below the poverty line. Retrieved May 2023, from https://staltz.com/software-below-the-poverty-line.html

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.

Athey, S., & Ellison, G. (2014). Dynamics of open source movements. *Journal of Economics & Management Strategy*, *23*(2), 294–316.

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., & Hortaçsu, 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

Balter, B. (2015). Open source license usage on github.com. Retrieved March 2023, from https://github.blog/2015-03-09-open-source-license-usage-on-github-com/

Bar-Isaac, H., & Tadelis, S. (2008). Seller Reputation. *Foundations and Trends® in Microeconomics*, *4*(4), 273–351. https://doi.org/10.1561/0700000027

Bar-Isaac, H., Tadelis, S., et al. (2008). Seller reputation. *Foundations and Trends® in Microeconomics*, *4*(4), 273–351.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral science*, *9*(3), 226–232.

Belleflamme, P., & Peitz, M. (2015). *Industrial organization: Markets and strategies*. Cambridge University Press.

Benjamin Cedric Larsen. (2022). The geopolitics of ai and the rise of digital sovereignty. Retrieved March 2023, from https://www.brookings.edu/articles/the-geopolitics-of-ai-and-the-rise-of-digital-sovereignty/

Besley, T. (1995). Property rights and investment incentives: Theory and evidence from ghana. *Journal of political Economy*, *103*(5), 903–937.

Blind, K., & Schubert, T. (2023). Estimating the gdp effect of open source software and its complementarities with r&d and patents: Evidence and policy implications. *The Journal of Technology Transfer*, 1–26.

Bloxam, G. (1957). Letters patent for inventions: Their use and misuse. *The Journal of Industrial Economics*, *5*(3), 157–179.

Bolton, G., Greiner, B., & Ockenfels, A. (2013a). Engineering Trust: Reciprocity in the Production of Reputation Information. *Management Science*, *59*(2), 265–285. https://doi.org/10.1287/mnsc.1120.1609

Bolton, G., Greiner, B., & Ockenfels, A. (2013b). Engineering trust: Reciprocity in the production of reputation information. *Management science*, *59*(2), 265–285.

Bolton, G. E., Katok, E., & Ockenfels, A. (2004a). 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., Katok, E., & Ockenfels, A. (2004b). How effective are electronic reputation mechanisms? an experimental investigation. *Management science*, *50*(11), 1587–1602.

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

Boone, J. (2001). Intensity of competition and the incentive to innovate. *International Journal of Industrial Organization*, *19*(5), 705–726.

Boston Consulting Group. (2021). Why you need an open source software strategy. Retrieved May 2023, from https://www.bcg.com/publications/2021/open-source-software-strategy-benefits

Bowles, S., & Polania-Reyes, S. (2012). Economic incentives and social preferences: Substitutes or complements? *Journal of Economic Literature*, *50*(2), 368–425.

Breton, A., & Breton, R. (1969). An economic theory of social movements. *The American Economic Review*, *59*(2), 198–205.

Brewster, R. (2009). A new license for paint.net v3.5. Retrieved March 2023, from https://blog.getpaint.net/2009/11/06/a-new-license-for-paintnet-v35/

Brian Proffitt. (2016). Hierarchy in open source. Retrieved March 2023, from https://www.redhat.com/en/blog/hierarchy-open-source

Cabral, L. (2000a). Stretching Firm and Brand Reputation. *The RAND Journal of Economics*, *31*(4), 658–673. https://doi.org/10.2307/2696353

Cabral, L., & Hortacsu, A. (2010). The dynamics of seller reputation: Evidence from ebay. *The Journal of Industrial Economics*, *58*(1), 54–78.

Cabral, L., & Li, L. (2015a). A dollar for your thoughts: Feedback-conditional rebates on ebay. *Management Science*, *61*(9), 2052–2063.

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. (2015b). A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay. *Management Science*, *61*(9), 2052–2063. https://doi.org/10.1287/mnsc.2014.2074

Cabral, L. M. (2000b). Stretching firm and brand reputation. *RAND Journal of Economics*, 658–673.

Cai, H., Jin, G. Z., Liu, C., & Zhou, L.-a. (2014a). 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

Cai, H., Jin, G. Z., Liu, C., & Zhou, L.-a. (2014b). Seller reputation: From word-of-mouth to centralized feedback. *International Journal of Industrial Organization*, *34*, 51–65.

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*.

Carnehl, C., Stenzel, A., & Schmidt, P. (2023). Pricing for the stars: Dynamic pricing in the presence of rating systems. *Management Science*.

Casadesus-Masanell, R., & Ghemawat, P. (2006). Dynamic mixed duopoly: A model motivated by linux vs. windows. *Management Science*, *52*(7), 1072–1084.

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*.

Chakraborty, I., Kim, M., & Sudhir, K. (2022). Attribute sentiment scoring with online text reviews: Accounting for language structure and missing attributes. *Journal of Marketing Research*, *59*(3), 600–622.

Chen, D. L., Schonger, M., & Wickens, C. (2016). Otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97.

Chen, Y., Harper, F. M., Konstan, J., & Li, S. X. (2010a). 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

Chen, Y., Harper, F. M., Konstan, J., & Li, S. X. (2010b). Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review*, *100*(4), 1358–98.

Chen, Y., Kartik, N., & Sobel, J. (2008). Selecting cheap-talk equilibria. *Econometrica*, *76*(1), 117–136.

Chen, Y., & Puttitanun, T. (2005). Intellectual property rights and innovation in developing countries. *Journal of development economics*, *78*(2), 474–493.

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.

Chiles, B. (2021). Shrouded prices and firm reputation: Evidence from the us hotel industry. *Management Science*, *67*(2), 964–983.

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

Chris Stokel-Walker. (2014). The internet is being protected by two guys named steve. Retrieved May 2023, from https://www.buzzfeed.com/chrisstokelwalker/the-internet-is-being-protected-by-two-guys-named-st

Chris Williams. (2016). How one developer just broke node, babel and thousands of projects in 11 lines of javascript. Retrieved May 2023, from https://www.theregister.com/2016/03/23/npm_left_pad_chaos/

Competition and Markets Authority (UK). (2017). *Digital comparison tools market study*.

Connors, L., Mudambi, S. M., & Schuff, D. (2011). Is it the review or the reviewer? a multi-method approach to determine the antecedents of online review helpfulness. *2011 44th Hawaii International Conference on System Sciences*, 1–10.

Cooper, R., DeJong, D. V., Forsythe, R., & Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, *107*(2), 739–771.

Crawford, G. S., Crémer, J., Dinielli, D., Fletcher, A., Heidhues, P., Luca, M., Salz, T., Schnitzer, M., Scott Morton, F. M., Seim, K., et al. (2021). Consumer protection for online markets and large digital platforms. *Policy Discussion Paper*, (1).

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, *50*(6), 1431–1451. https://doi.org/10.2307/1913390

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/

d'Aspremont, C., Gabszewicz, J. J., & Thisse, J.-F. (1979). On hotelling's "stability in competition". *Econometrica*, *47*(5), 1145–1150.

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

Dellarocas, C., Zhang, X., & Awad, N. F. (2007). Exploring the value of online product reviews in forecasting sales: The case of motion pictures. *Journal of Interactive marketing*, *21*(4), 23–45.

Denis Pushkarev. (2023). 2023-02-14-so-whats-next.md. Retrieved May 2023, from https://github.com/zloirock/core-js/blob/master/docs/2023-02-14-so-whats-next.md

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/

Diekmann, A., Jann, B., Przepiorka, W., & Wehrli, S. (2014a). Reputation Formation and the Evolution of Cooperation in Anonymous Online Markets. *American Sociological Review*, *79*(1), 65–85. https://doi.org/10.1177/0003122413512316

Diekmann, A., Jann, B., Przepiorka, W., & Wehrli, S. (2014b). Reputation formation and the evolution of cooperation in anonymous online markets. *American sociological review*, *79*(1), 65–85.

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

Emma Roth. (2022). Open source developer corrupts widely-used libraries, affecting tons of projects. *The Verge*. Retrieved October 2022, from https://www.theverge.com/2022/1/9/22874949/developer-corrupts-open-source-libraries-projects-affected

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

Etzion, H., & Pang, M.-S. (2014). Complementary online services in competitive markets: Maintaining profitability in the presence of network effects. *Mis Quarterly*, *38*(1), 231–248.

Fainmesser, I. P., & Galeotti, A. (2020). Pricing network effects: Competition. *American Economic Journal: Microeconomics*, *12*(3), 1–32.

Fielding, R. T. (1999). Shared leadership in the apache project. *Communications of the ACM*, *42*(4), 42–43.

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., Grewal, E., & Holtz, D. (2021a). Reciprocity and unveiling in two-sided reputation systems: Evidence from an experiment on airbnb. *Marketing Science*, *40*(6), 1013–1029. https://doi.org/10.1287/mksc.2021.1311

Fradkin, A., Grewal, E., & Holtz, D. (2021b). Reciprocity and unveiling in two-sided reputation systems: Evidence from an experiment on airbnb. *Marketing Science*, *40*(6), 1013–1029.

Fradkin, A., & Holtz, D. (2022). Do incentives to review help the market? evidence from a field experiment on airbnb. *American Economic Review: Insights*.

Franke, N., & Von Hippel, E. (2003). Satisfying heterogeneous user needs via innovation toolkits: The case of apache security software. *Research policy*, *32*(7), 1199–1215.

Gamblin, T. (2021). *Picking an open source license at llnl: Guidance and recommendations from the computing directorate* (tech. rep.). Lawrence Livermore National Lab.(LLNL). Livermore, CA (United States).

Gao, B., Li, X., Liu, S., & Fang, D. (2018a). How power distance affects online hotel ratings: The positive moderating roles of hotel chain and reviewers' travel experience. *Tourism Management*, *65*, 176–186. https://doi.org/10.1016/j.tourman.2017.10.007

Gao, B., Li, X., Liu, S., & Fang, D. (2018b). How power distance affects online hotel ratings: The positive moderating roles of hotel chain and reviewers' travel experience. *Tourism management*, *65*, 176–186.

Gaudeul, A. (2005). Public provision of a private good: What is the point of the bsd license? *Available at SSRN 933631*.

Gaudeul, A. (2008). Open source licensing in mixed markets, or why open source software does not succeed. *CCP Working Paper*.

Geanakoplos, J., Pearce, D., & Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and economic Behavior*, *1*(1), 60–79.

Gesche, T. (2022). Reference-price shifts and customer antagonism: Evidence from reviews for online auctions. *Journal of Economics & Management Strategy*, *31*(3), 558–578.

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. (2019a). 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

Halliday, S. D., & Lafky, J. (2019b). Reciprocity through ratings: An experimental study of bias in evaluations. *Journal of Behavioral and Experimental Economics*, *83*, 101480.

Hars, A., & Ou, S. (2002). Working for free? motivations for participating in open-source projects. *International journal of electronic commerce*, *6*(3), 25–39.

He, S., Hollenbeck, B., & Proserpio, D. (2022). The market for fake reviews. *Marketing Science*. https://doi.org/10.1287/mksc.2022.1353

Healy, P. J. (2018). *Explaining the bdm—or any random binary choice elicitation mechanism—to subjects* (tech. rep.). mimeo.

Hefti, A., Liu, S., & Schmutzler, A. (2022). Preferences, confusion and competition. *The Economic Journal*, *132*(645), 1852–1881.

Hotelling, H. (1929). Stability in competition. *The Economic Journal*, *39*(153), 41–57.

Hoyer, B., & van Straaten, D. (2022). Anonymity and self-expression in online rating systems—an experimental analysis. *Journal of Behavioral and Experimental Economics*, *98*, 101869.

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

Hui, X., Saeedi, M., Spagnolo, G., & Tadelis, S. (2021). Raising the bar: Certification thresholds and market outcomes. *American Economic Journal: Microeconomics*.

Hui, X., Saeedi, M., Spagnolo, G., & Tadelis, S. (2023). Raising the bar: Certification thresholds and market outcomes. *American Economic Journal: Microeconomics*, *15*(2), 599–626.

Introducing steam reviews. (2021). Retrieved June 2021, from https://store.steampowered.com/reviews/

Janis Lesinski. (2020). The mit licence is community hostile. Retrieved May 2023, from https://www.lesinskis.com/MIT-licence-community-hostile.html

Jin, G. Z., & Kato, A. (2006a). 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

Jin, G. Z., & Kato, A. (2006b). Price, quality, and reputation: Evidence from an online field experiment. *The RAND Journal of Economics*, *37*(4), 983–1005.

John Biggs. (2014). Heartbleed, the first security bug with a cool logo. Retrieved May 2023, from https://techcrunch.com/2014/04/09/heartbleed-the-first-consumer-grade-exploit/

Johnen, J., & Leung, B. T. K. (2022). Distracted from comparison: Product design and advertisement with limited attention.

Johnen, J., & Ng, R. (2023a). Ratings and reciprocity. *CEPR Discussion Paper No. 18021*.

Johnen, J., & Ng, R. (2023b). Ratings and Reciprocity Web Appendix. https://robin-ng.com/download/ratings_and_reciprocity_appendix_B.pdf

Johnen, J., & Somogyi, R. (2021). Deceptive features on platforms. https://ssrn.com/abstract=3886615

Johnson, J. P. (2006). Collaboration, peer review and open source software. *Information Economics and Policy*, *18*(4), 477–497.

Johnson, J. P. (2002). Open source software: Private provision of a public good. *Journal of Economics & Management Strategy*, *11*(4), 637–662.

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. (2014a). New, Like New, or Very Good? Reputation and Credibility. *The Review of Economic Studies*, *81*(4), 1543–1574. https://doi.org/10.1093/restud/rdu012

Jullien, B., & Park, I.-U. (2014b). New, like new, or very good? reputation and credibility. *The Review of Economic Studies*, *81*(4), 1543–1574.

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

Klug, D., Bogart, C., & Herbsleb, J. D. (2021). "they can only ever guide" how an open source software community uses roadmaps to coordinate effort. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW1), 1–28.

Kovbasyuk, S., & Spagnolo, G. (2021). Memory and markets. *Available at SSRN 2756540*.

Kreps, D. M., & Sobel, J. (1994). Signalling. *Handbook of game theory with economic applications*, *2*, 849–867. https://doi.org/10.1016/S1574-0005(05)80057-8

Lafky, J. (2014a). Why do people rate? Theory and evidence on online ratings. *Games and Economic Behavior*, *87*, 554–570. https://doi.org/10.1016/j.geb.2014.02.008

Lafky, J. (2014b). Why do people rate? theory and evidence on online ratings. *Games and Economic Behavior*, *87*, 554–570.

Lagunoff, R., & Matsui, A. (2004). Organizations and overlapping generations games: Memory, communication, and altruism. *Review of Economic Design*, *8*, 383–411.

Lakhani, K. R., & Hippel, E. v. (2004). How open source software works: "free" user-to-user assistance. In *Produktentwicklung mit virtuellen communities* (pp. 303–339). Springer.

Lerner, J., & Tirole, J. (2002). Some simple economics of open source. *The journal of industrial economics*, *50*(2), 197–234.

Lerner, J., & Tirole, J. (2005a). The economics of technology sharing: Open source and beyond. *Journal of Economic Perspectives*, *19*(2), 99–120.

Lerner, J., & Tirole, J. (2005b). The scope of open source licensing. *Journal of Law, Economics, and Organization*, *21*(1), 20–56.

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., Tadelis, S., & Zhou, X. (2020a). Buying reputation as a signal of quality: Evidence from an online marketplace. *The RAND Journal of Economics*, *51*(4), 965–988.

Li, L., & Xiao, E. (2014a). Money talks: Rebate mechanisms in reputation system design. *Management Science*, *60*(8), 2054–2072.

Li, L. I., Tadelis, S., & Zhou, X. (2020b). 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, L. I., & Xiao, E. (2014b). Money Talks: Rebate Mechanisms in Reputation System Design. *Management Science*, *60*(8), 2054–2072. https://doi.org/10.1287/mnsc.2013.1848

Li, X., & Hitt, L. M. (2010a). Price Effects in Online Product Reviews: An Analytical Model and Empirical Analysis. *MIS Quarterly*, *34*(4), 809–831. https://doi.org/10.2307/25750706

Li, X., & Hitt, L. M. (2010b). Price effects in online product reviews: An analytical model and empirical analysis. *MIS quarterly*, 809–831.

Li, Y., Tan, C.-H., & Teo, H.-H. (2012). Leadership characteristics and developers' motivation in open source software development. *Information & Management*, *49*(5), 257–267.

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. (2021a). The effect of price on firm reputation. *Management Science*, *67*(7), 4408–4419. https://doi.org/10.1287/mnsc.2021.4049

Luca, M., & Reshef, O. (2021b). The effect of price on firm reputation. *Management Science*, *67*(7), 4408–4419.

Luke Little. (2021). Google chrome: Its history and rise to market domination. Retrieved November 2022, from https://www.androidauthority.com/google-chrome-history-1025602/

Madiega, T. A. (2020). Digital sovereignty for europe. *EPRS: European Parliamentary Research Service*.

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

Masterov, D. V., Mayer, U. F., & Tadelis, S. (2015a). Canary in the e-Commerce Coal Mine: Detecting and Predicting Poor Experiences Using Buyer-to-Seller Messages. *Proceedings of the Sixteenth ACM Conference on Economics and Computation - EC '15*, 81–93. https://doi.org/10.1145/2764468.2764499

Masterov, D. V., Mayer, U. F., & Tadelis, S. (2015b). Canary in the e-commerce coal mine: Detecting and predicting poor experiences using buyer-to-seller messages. *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 81–93.

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/

Mayzlin, D., Dover, Y., & Chevalier, J. (2014). Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review*, *104*(8), 2421–55.

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

Michlmayr, M., Hunt, F., & Probert, D. (2007). Release management in free software projects: Practices and problems. *Open Source Development, Adoption and Innovation: IFIP Working Group 2.13 on Open Source Software, June 11–14, 2007, Limerick, Ireland 3*, 295–300.

Microsoft. (2023a). Microsoft 365 roadmap. Retrieved May 2023, from https://www.microsoft.com/en-gb/microsoft-365/roadmap?filters=

Microsoft. (2023b). Microsoft release cycle. Retrieved May 2023, from https://learn.microsoft.com/en-us/windows/deployment/update/release-cycle

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

Moser, P. (2005). How do patent laws influence innovation? evidence from nineteenth-century world's fairs. *American economic review*, *95*(4), 1214–1236.

Mustonen, M. (2003). Copyleft—the economics of linux and other open source software. *Information Economics and Policy*, *15*(1), 99–121.

Nelson, P. (1974). Advertising as information. *Journal of political economy*, *82*(4), 729–754. https://doi.org/10.1086/260231

Neumann, J., & Gutt, D. (2017). 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.*

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.

Nicolas Suzor. (2013). What motivates free software developers to choose between copyleft and permissive licences? Retrieved May 2023, from https://opensource.com/law/13/8/motivation-free-software-licensing

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

Open Source Initiative. (2022a). About. Retrieved November 2022, from https://opensource.org/about

Open Source Initiative. (2022b). Licenses & standards. Retrieved November 2022, from https://opensource.org/licenses

Palfrey, T. R., & Rosenthal, H. (1991). Testing for effects of cheap talk in a public goods game with private information. *Games and economic behavior*, *3*(2), 183–220.

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. (2018a). 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

Proserpio, D., Xu, W., & Zervas, G. (2018b). You get what you give: Theory and evidence of reciprocity in the sharing economy. *Quantitative Marketing and Economics*, *16*(4), 371–407.

Qian, Y. (2007). Do national patent laws stimulate domestic innovation in a global patenting environment? a cross-country analysis of pharmaceutical patent protection, 1978–2002. *The Review of Economics and Statistics*, *89*(3), 436–453.

Qiao, D., Lee, S.-Y., Whinston, A. B., & Wei, Q. (2020). Financial incentives dampen altruism in online prosocial contributions: A study of online reviews. *Information Systems Research*, *31*(4), 1361–1375.

Qiu, L., Pang, J., & Lim, K. H. (2012). Effects of conflicting aggregated rating on ewom review credibility and diagnosticity: The moderating role of review valence. *Decision Support Systems*, *54*(1), 631–643. https://doi.org/10.1016/j.dss.2012.08.020

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.

Raymond, E. (1999). The cathedral and the bazaar. *Knowledge, Technology & Policy*, *12*(3), 23–49.

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

Robles, G., Steinmacher, I., Adams, P., & Treude, C. (2019). Twenty years of open source software: From skepticism to mainstream. *IEEE Software*, *36*(6), 12–15.

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. (2021a). 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

Schneider, C., Weinmann, M., Mohr, P. N., & vom Brocke, J. (2021b). When the stars shine too bright: The influence of multidimensional ratings on online consumer ratings. *Management Science*, *67*(6), 3871–3898.

Sean Fleming. (2021). What is digital sovereignty and why is europe so interested in it? Retrieved March 2023, from https://www.weforum.org/agenda/2021/03/europe-digital-sovereignty/

Shah, S. K. (2006). Motivation, governance, and the viability of hybrid forms in open source software development. *Management science*, *52*(7), 1000–1014.

Shy, O., & Thisse, J.-F. (1999). A strategic approach to software protection. *Journal of Economics & Management Strategy*, *8*(2), 163–190.

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

Stackexchange. (2015). How can a project be relicensed? Retrieved May 2023, from https://opensource.stackexchange.com/questions/33/how-can-a-project-be-relicensed

Statcounter. (2022). Desktop browser market share worldwide. Retrieved November 2022, from https://gs.statcounter.com/browser-market-share/desktop/worldwide%5C#monthly-200901-202210

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

Subramaniam, C., Sen, R., & Nelson, M. L. (2009). Determinants of open source software project success: A longitudinal study. *Decision Support Systems*, *46*(2), 576–585.

Tadelis, S. (1999a). 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

Tadelis, S. (1999b). What's in a name? reputation as a tradeable asset. *American Economic Review*, *89*(3), 548–563.

TaobaoTranslate. (2021). *Taobao item review guide*. Retrieved June 2021, from https://www.taobaotranslate.com/taobao-item-review-rating/taobao-item-rating.html

The Apache Software Foundation. (2022). Apache in 2021 – by the digits. Retrieved March 2023, from https://news.apache.org/foundation/entry/apache-in-2021-by-the

The Apache Software Foundation. (2023a). Apache corporate governance. Retrieved March 2023, from https://www.apache.org/foundation/governance/

The Apache Software Foundation. (2023b). Foundation reports and statements. Retrieved May 2023, from https://www.apache.org/foundation/reports.html

The Document Foundation. (2023). The document foundation. Retrieved March 2023, from https://www.documentfoundation.org/governance/bodies/

The Linux Foundation. (2023). Linux foundation leadership teams. Retrieved March 2023, from https://www.linuxfoundation.org/about/leadership

Thomas Greene. (2001). Ballmer: 'linux is a cancer'. Retrieved March 2023, from https://www.theregister.com/2001/06/02/ballmer_linux_is_a_cancer/

Tom Warren. (2018). Google's chrome browser is now 10 years old. Retrieved November 2022, from https://www.theverge.com/2018/9/2/17811844/google-chrome-browser-10-years-history

Tom Warren. (2020). Microsoft: We were wrong about open source. Retrieved March 2023, from https://www.theverge.com/2020/5/18/21262103/microsoft-open-source-linux-history-wrong-statement

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

Vendome, C., Linares-Vásquez, M., Bavota, G., Di Penta, M., German, D., & Poshyvanyk, D. (2015a). License usage and changes: A large-scale study of java projects on github. *2015 IEEE 23rd International Conference on Program Comprehension*, 218–228.

Vendome, C., Linares-Vásquez, M., Bavota, G., Di Penta, M., German, D. M., & Poshyvanyk, D. (2015b). When and why developers adopt and change software licenses. *2015 IEEE international conference on software maintenance and evolution (ICSME)*, 31–40.

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

Woodside, A. G., & Davenport Jr, J. W. (1974). The effect of salesman similarity and expertise on consumer purchasing behavior. *Journal of Marketing Research*, *11*(2), 198–202.

Yelp. (2022). *Yelp elite squad*. Retrieved July 2022, from https://www.yelp.com/elite

Zervas, G., Proserpio, D., & Byers, J. W. (2021a). 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

Zervas, G., Proserpio, D., & Byers, J. W. (2021b). A first look at online reputation on airbnb, where every stay is above average. *Marketing Letters*, *32*(1), 1–16.

Zhang, X., Luo, J., & Li, Q. (2012a). Do different reputation systems provide consistent signals of seller quality: A canonical correlation investigation of Chinese C2C marketplaces. *Electronic Markets*, *22*(3), 155–168. https://doi.org/10.1007/s12525-012-0092-4

Zhang, X., Luo, J., & Li, Q. (2012b). Do different reputation systems provide consistent signals of seller quality: A canonical correlation investigation of chinese c2c marketplaces. *Electronic Markets*, *22*, 155–168.

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/