Strategic Review Presentation*

Leif Brandes[†]
University of Lucerne

David Godes[‡]
Johns Hopkins University

Dina Mayzlin[§]
University of Southern California

December 16, 2024

Abstract

Online reviews represent a key information source for consumers to distinguish between products of different quality levels. An implicit assumption in much of the extant literature is that marketers benefit from more precise, high quality review information (see, for example, Tadelis, 2015). By now, review platforms have many opportunities to facilitate online review systems: AI-based topic extraction, enabling the upload of additional content (e.g., videos, images), personalizing the set of displayed reviews, etc. However, we did not find a single platform that employs all these opportunities. Here we demonstrate that review platforms may benefit from the use of online review systems that obfuscate quality differences across sellers. Our research makes four substantial contributions. First, we provide the novel insight that marketers may not always want to facilitate review evaluation for customers. This result stands in stark contrast to the prevailing views in the literature on online feedback systems and online reviews. Second, we demonstrate an inherent tension between price competition and review precision: under perfectly precise reviews, no pooling equilibrium exists. Third, our analytical model integrates a platform's review funnel design with a signaling game between two sellers. This is a key distinction from most existing models in marketing. Finally, we add a new mechanism to the literature on search obfuscation. While past work has focused on the role of price and others, we are the first to show that the design of online review systems has a direct effect on price competition intensity on a website.

Keywords: microeconomics, electronic commerce, online reviews, online word of mouth

*The authors are thankful to the participants at the Northeastern marketing seminar and the participants at the 2024 *Marketing Science* Conference in Sydney, Australia.

[†]leif.brandes@unilu.ch

[‡]dgodes@jhu.edu

[§]mayzlin@marshall.usc.edu

1 Introduction

Online reviews represent a key information source for consumers to distinguish between products of different quality levels. Studies have investigated ways for marketers to facilitate the provision, exposure and evaluation of consumers to online reviews (see Babic Rosario et al. 2020 for an overview). An implicit assumption in much of the extant literature is thus that marketers benefit from this facilitation. In this paper we ask the extent to which this assumption always holds.

Consider the review evaluation stage, the focus of the present article. By now, review platforms have many opportunities to facilitate review evaluation: AI-based topic extraction, enabling the upload of additional content (e.g., videos, images), personalizing the set of displayed reviews, etc. Previous work on online marketplaces emphasizes the value of better reputation systems for platforms (e.g., Tadelis, 2015). However, we did not find a single platform that employs all these opportunities. We suggest that there be a strategic reason for this behavior.

In this article, we demonstrate that review platforms may benefit from the use of online review systems that obfuscate quality differences across sellers. Our research makes four substantial contributions. First, we provide the novel insight that marketers may not always want to facilitate review evaluation for customers. This result stands in stark contrast to the prevailing views in the literature on online feedback systems and online reviews. Second, we demonstrate an inherent tension between price competition and review precision: under perfectly precise reviews, no pooling equilibrium exists. Third, our analytical model integrates a platform's review funnel design with a signaling game between two sellers. This is a key distinction from most existing models in marketing. Finally, we add a new mechanism to the literature on search obfuscation. While past work has focused on the role of price and others, we are the first to show that the design of online review systems has a direct effect on price competition intensity on a website.

2 Preliminary Analysis

2.1 Monopoly E-Commerce Platform (e.g., Amazon)

Consider two sellers, one of which offers high value V^H and the other low value V^L, and a retail platform that knows V and provides information (processed reviews) that can inform consumers about the sellers' types. Let $\theta \in \{H, L\}$ be the seller types. The stages of the game are as follows:

1. Nature picks $V \in \{V^H, V^L\}$. This is observable to the firms and to the platform but not to consumers
2. Platform announces review presentation policy
3. Firms choose prices P^H and P^L (marginal costs are zero)
4. Platform provides reviews consistent with announced policy
5. Commerce ensues with consumer buying product with highest expected value net of price
6. Platform earns $\gamma \cdot REV$, and successful seller earns $REV \cdot (1 - \gamma)$

The key aspect to notice is that the decision by the platform as to whether or not there is residual uncertainty over V following provision of reviews will determine whether the pricing game is one of complete or incomplete information. The two non-standard specifications here are that (i) the platform chooses the level of uncertainty and (ii) there is the potential for signaling (due to private information) in a game with two sellers (usually, there's one seller of unknown type). This will require us to think through how various equilibrium concepts/refinements carry over.

Let $\sigma \in \{H, L\}$ be consumer's inference on firm type, based on reviews provided according to the implemented system on the platform. I.e., after reading the reviews – which might be on any possible scale, might include ML processing, might allow for Q&A, etc. An easy way to think about it would be that σ is the top ranked product in

the category. The specific mechanism leading to the inference is not modeled (could be of course) but it takes the form of of a posterior belief such that:

$$\pi \equiv \Pr[\sigma = H | \theta = H] \tag{1}$$

where the platform's main choice variable (which is meant to capture all these possible review presentation strategies) is π.[1] There's probably a way to characterize this in horizontal differentiation (big vs small segments). Anyway, the focal question is what is the platform's choice of π? Does it always pass through reviews with perfect precision?

We could also allow for a platform-specific benefit of review precision such that the objective function includes an additive term $\rho \cdot \pi$. As well, we could reflect the fact that there may be a (fixed) cost associated with providing better information: $C(\pi)$ such that $C' > 0$. We'll abstract away from these at the outset.

Formally, an equilibrium is defined in terms of a) the platform's review presentation strategy, π, b) the pair of prices chosen by the sellers, $(\tilde{P}^H, \tilde{P}^L)$, and c) the consumer's belief after reading reviews that the firm choosing price P^i is high quality, $\mu(P^i, P^j, \pi)$. The fact that the beliefs are a function of two price observations is the key distinction between this set-up and what we typically see in the (marketing) literature. Following the exposition in Hertzendorf and Overgaard (2001), the equilibrium must satisfy:

1. $\pi \in \text{argmax } \Pi^P(P^H, \tilde{P}^L, \mu(P^H, \tilde{P}^L, \pi), V^H, V^L, \pi)$
2. $\tilde{P}^H \in \text{argmax } \Pi^H(P^H, \tilde{P}^L, \mu(P^H, \tilde{P}^L), V^H, V^L, \pi)$
3. $\tilde{P}^L \in \text{argmax } \Pi^L(\tilde{P}^H, P^L, \mu(\tilde{P}^H, P^L), V^H, V^L, \pi)$
4. If $\tilde{P}^H \neq \tilde{P}^L$, then $\mu(\tilde{P}^H, \tilde{P}^L | V^H, V^L, \pi) = 1$

[1] Suppose that an inspection of reviews on the platform produces a relative ranking of the two products. That is, suppose that the consumer searches for pressure cookers and examines the reviews presented by the platform. We model the result of the search as a message of the type, "Product 1 is H" or "Product 2 is H." Suppose that the prior belief that either product is high quality is $\frac{1}{2}$. The platform controls the precision of the message received by the consumer after she examines the reviews on the platform. For example, the platform can display the reviews in a way that makes it super clear which is the better product or it can make the information appear more ambiguous: P("Product 1 is H"|Product 1 is H) = P("Product 1 is L"|Product 1 is L) = π. Using Bayes' Rule, the consumer's inference is P(Product 1 is H|"Product 1 is H") = $\frac{\frac{1}{2}\pi}{\frac{1}{2}\pi + \frac{1}{2}(1-\pi)} = \pi$.

5. If $\tilde{P}^H = \tilde{P}^L$ and $\pi < 1$, then $\mu(\tilde{P}^H, \tilde{P}^L | V^H, V^L, \pi) = \pi$

6. If $\tilde{P}^H = \tilde{P}^L$ and $\pi = 1$, then $\mu(\tilde{P}^H, \tilde{P}^L | V^H, V^L, \pi) = 1$

7. $\mu(P^H, P^L | V^H, V^L, \pi) + \mu(P^L, P^H | V^H, V^L, \pi) = 1$ [2]

An important distinction that we will need to highlight to differentiate from Hertzendorf and Overgaard is that they use dissipative advertising (burning money) whereas our review platforms provide real information that impacts signaling but isn't necessarily a signal in itself (except in the FPR case).

Full Information: $\pi = 1$

The simplest (and perhaps expected) case would be where the platform eliminates all consumer uncertainty and sets $\pi = 1$.[3] In this case, the consumer chooses between $U^H \equiv V^H - \tilde{P}^H$ and $U^L \equiv V^L - \tilde{P}^L$. The first question is whether L has any incentive to mimic H and the answer is obviously no. What prices sustain such an equilibrium? Of course, in this case, it would lead to a Bertrand game between the firms in which $\tilde{P}^L = 0$ and $\tilde{P}^H = V^H - V^L$. In this case, the platform earns $\gamma \cdot (V^H - V^L)$, the high type earns $(1-\gamma) \cdot (V^H - V^L)$, the low type earns zero and consumers earn surplus of V^L.

The key point here is that there can be no pooling on price when $\pi = 1$. I think this is a potentially novel insight in that there's an inherent tension between price competition and review precision.

Incomplete Information: $\pi \in (0, 1)$

If there remains uncertainty after reading the reviews, then the pricing game is, again,

[2] This seems superfluous given 3 and 4 but H & O state it.

[3] We still need to consider the case where this isn't possible, but the more I think about it, the more this just eliminates the $\pi = 1$ case and the rest is just incomplete information.

one of incomplete information. Let's start by considering a **a separating equilibrium** where $\tilde{P}^H \neq \tilde{P}^L$. Of course, as above, the consumer buys from the firm offering the highest expected value. Thus, the only possible equilibrium here could be the Bertrand outcome with $\tilde{P}^L = 0$ and $\tilde{P}^H = V^H - V^L$. But, even this will not be sustainable because firm L has a clear profitable deviation to $P^{L'} = V^H - V^L$. Any $\mu(V^H - V^L, V^H - V^L)$ would yield positive profits (whenever $\sigma = L$) and thus kill the equilibrium. We don't need to consider off-path beliefs. Essentially, in this case, the imperfect signal from the reviews is overwhelmed by the information delivered by the separating equilibrium in prices. It is likely the case in general that many separating equilibria disappear in a multi-firm setting. This result is consistent with the intuition that price signaling in the monopoly setting requires certain elements such as differential costs, demand or profit across types.

What about a **pooling equilibrium** in prices? Assuming $\pi > .5$, a consumer's optimal strategy in this equilibrium would be to choose the seller indicated by σ, which would yield $E[U] = \pi V^H + (1-\pi)V^L - \tilde{P}$. Can we sustain a pooling equilibrium in which $\tilde{P}^H = \tilde{P}^L = \pi V^H + (1-\pi)V^L$?[4]

From the platform's perspective, would this be better than the firms separating in prices under full information? We'll ignore γ since the platform optimally maximizes price in this simple game (this will change when/if we introduce platform competition). The platform prefers the pricing game of incomplete information to the game of complete information if the following condition holds:

$$\pi V^H + (1-\pi)V^L > V^H - V^L \tag{2}$$

$$\Leftrightarrow \frac{V^L}{V^H - V^L} > 1 - \pi \tag{3}$$

$$\Leftrightarrow \frac{V^H}{V^L} < \frac{(2-\pi)}{(1-\pi)} \tag{4}$$

[4]I think that any price $V^L \leq \tilde{P} \leq \pi V^H + (1-\pi)V^L$ is sustainable in the incomplete information pooling equilibrium. Should we have a formal result on this? Are there refinements we can use to rule out any pooling prices? I think that we probably could in the case of monopoly

We can show that the RHS of Equation (4) is increasing in π. Hence, an increase in π implies that the platform prefers the pricing game of incomplete information. A lower difference between V^H and V^L also implies that the platform prefers the pricing game of incomplete information.

The sellers earn the following profits on the equilibrium path:

$$\tilde{\Pi}^{H*} = \pi(\pi V^H + (1-\pi)V^L)$$
$$\tilde{\Pi}^{L*} = (1-\pi)(\pi V^H + (1-\pi)V^L)$$

Note that as $\pi \to 1$, the low type earnings vanish.

The key (as always) is to establish the nature of "reasonable" off-path beliefs that support the equilibrium. As noted by Hertzendorf and Overgaard, a challenge is that the typical available refinements are designed for a monopoly setting. Since neither of their suggestions are appropriate here, we propose a criterion which is very much an application of the Intuitive Criterion. That is, the beliefs following an out of equilibrium action assign zero probability to types that could never benefit from the observed deviation. We continue to assume that an equilibrium is defined as the absence of a profitable unilateral deviation, i.e. we assume that only one firm deviates.

First consider a deviation to a higher price: $P' > \tilde{P}$. Since both firms would benefit from this, we can't rule either of them out (a la the IC), and thus it is uninformative and therefore unprofitable.

Now consider a deviation to a lower price: $P' < \tilde{P}$. Here, the first question is whether and when a given firm could ever benefit? I.e., is there some set of beliefs that could imply a gain from such a deviation? Let's start with H and assume that the inference is this is the high type with probability 1? The two questions here are (a) would this

increase sales and (b) if so, would it increase revenue? Consider a small cut and let's ask what's the best that could happen for the firm? Of course, the best would be that $\mu^H = 1$ and all customers would buy regardless of reviews. As long as the cut is small enough of course it would benefit both firms and thus the only reasonable such inference would thus be that the deviation is uninformative and thus unprofitable.

Imagine now a larger cut such that firm H would never benefit. This would imply that

$$\pi \cdot \tilde{P} \geq P'$$
$$\leftrightarrow \pi \cdot (\pi \cdot V^H + (1-\pi) \cdot V^L) \geq P'$$

In fact, deviating to any price such that

$$\pi \cdot (\pi \cdot V^H + (1-\pi) \cdot V^L) \geq P' > (1-\pi) \cdot (\pi \cdot V^H + (1-\pi) \cdot V^L)$$

would, in the best case, benefit L but not H. As a result, conditional on seeing such a deviation, the customer would be justified in assuming it's L with probability 1. This is if course a potential threat to the equilibrium. First, we check whether the consumer would actually purchase in this case (we had a assumed the best case, but this may not always be true). Recall that all is now revealed (since H would never implement such a deviation) and so the customer's decisions is between

$$V^H - (\pi \cdot V^H + (1-\pi) \cdot V^L)$$

and

$$V^L - P^{L'}$$

implying that under full information, consumers will only purchase from the deviating low firm iff

$$P^{L'} < (2-\pi) \cdot V^L - (1-\pi) \cdot V^H$$

Finally, we want to check whether this is makes L better off. We'll do so by checking the "best" such price deviation:

$$\begin{aligned} &(2-\pi) \cdot V^L - (1-\pi) \cdot V^H > (1-\pi) \cdot (\pi \cdot V^H + (1-\pi) \cdot V^L) \\ \leftrightarrow &\frac{V^H}{V^L} < \frac{1+\pi-\pi^2}{1-\pi^2} \end{aligned}$$

Note that the RHS of this inequality is strictly increasing in π, which implies that in order to prevent this deviation the platform needs to lower the precision π. This is somewhat intuitive as very high precision hurts L.

So, to collect constraints here, in order to sustain the pooling equilibrium, we need:

1. The platform prefers the pool over the Bertrand Competition associated with perfect information:

$$\frac{V^H}{V^L} < \frac{(2-\pi)}{(1-\pi)}$$

2. Any deviation to a higher price or to a slightly lower price would impact both firms equally and thus would be uninformative

3. The high type would never deviate to the observed price even under the best circumstances:

$$\pi \cdot (\pi \cdot V^H + (1-\pi) \cdot V^L) \geq P'$$

4. The low type is (weakly) worse off by deviating to a lower price that reveals its type (ie satisfying the Intuitive Criterion):

$$(1-\pi) \cdot (\pi \cdot V^H + (1-\pi) \cdot V^L) < P'$$

Focusing on the best possible deviation – $P' = (2-\pi) \cdot V^L - (1-\pi) \cdot V^H$ – allows us to rewrite (3) and (4) as, respectively:

$$\frac{V^H}{V^L} \geq \frac{2 - 2 \cdot \pi + \pi^2}{1 - \pi + \pi^2}$$
$$\frac{V^H}{V^L} \geq \frac{1 + \pi - \pi^2}{1 - \pi^2}$$

It is easy to show that (as we would expect, given the asymmetry) the latter dominates (has a strictly higher RHS for a given π than) the former. And so we can write the condition for the pooling equilibrium as follows:

$$\frac{1 + \pi - \pi^2}{1 - \pi^2} < \frac{V^H}{V^L} < \frac{2 - \pi}{1 - \pi} \tag{5}$$

It is easy to check that this is possible since $\frac{1+\pi-\pi^2}{1-\pi^2} < \frac{2-\pi}{1-\pi}$. Since $\frac{2-\pi}{1-\pi} \in (2, \infty)$, This right-hand inequality is not really binding. The left-hand inequality though does matter, This is the one that captures the low type's motivation to deviate. Intuitively, the higher π is, the more incentive the low type will have to deviate because it is less likely to be highly ranked and thus less likely to earn a sale in the pooling equilibrium. Since $\frac{1+\pi-\pi^2}{1-\pi^2} \in (1, \infty)$ is strictly increasing in π the platform needs to select π appropriately and in particular not too high given the relative values V. It should choose the highest possible π that enforces the pooling equilibrium.

So, the basic intuition here is that the platform can induce a game between sellers by adding some uncertainty via imperfect reviews in which the low type doesn't have to resort to Bertrand price competition. The benefit of this noise is that it allows low type sellers to make a sale every now and then, preventing them from starting the process of price unraveling.