Seller Curation in Platforms*

Preliminary Version

Ben Casner†

February 14, 2019

Abstract

This article explores why market platforms do not screen out low-quality sellers in situations where screening costs are minimal. Consumers in a platform's market must search for a seller whose product is a good match. The presence of low-quality sellers reduces search intensity which softens competition between sellers, increasing equilibrium price and the platform's revenue per sale. If the platform's market is sufficiently competitive, then it admits a positive proportion of low-quality sellers. Surprisingly, recommending a high-quality seller and search obfuscation are complementary strategies because the low-quality sellers enable the recommended seller to attract many consumers at a high price.

Keywords: Search Obfuscation, Two-sided Markets, Platforms, Screening, Recommendations

JEL Codes: D21, D83, L22, L15

*I wish to thank Jim Peck, Huanxing Yang, PJ Healy, Volcker Nocke, Teddy Kim, Laura Doval, Parker Wheatley, Yaron Azrieli, Adam Dearing, Kirby Nielsen, Jinsoo Bae, Soyoung Lee and many others for their helpful comments. All mistakes are my own.

†The Ohio State University. Mailing address: 410 Arps Hall, 1945 N. High St. Columbus OH, 43210. Email: casner.15@osu.edu.

1 Introduction

Platforms which connect buyers and sellers suffer a loss of reputation if they host a significant portion of low-quality sellers, yet many platforms seem unwilling to curate the quality of sellers on their platform[1]. While the cost of screening could potentially explain platforms' behavior, this reluctance appears even in cases when simple low cost measures could significantly increase the average quality of products on offer[2].

I present a model where a market platform hosts sellers in exchange for a percentage of their revenue. Sellers can either be low- or high-quality, and the platform decides the proportion of low-quality to high-quality sellers on its market. Consumers participating in the platform's market engage in a standard Wolinsky (1986) style search process before they purchase from a seller.

I show that the low-quality sellers can act as a search obfuscation mechanism similar to the one defined by Ellison and Ellison (2009). Consumers never knowingly purchase from a low-quality seller, so if a consumer searches and encounters a low-quality seller then they will pay the search cost to visit another seller, if that next seller is low-quality then the consumer pays the search cost yet again to visit yet another seller and so on until they encounter a seller they believe to be high-quality. Because of the probability of encountering a low-quality seller upon searching, the value of searching decreases as the proportion of low-quality sellers increases, which in turn leads to higher equilibrium prices[3]. By softening competition between sellers

[1] See for example https://www.npr.org/sections/money/2018/06/27/623990036/episode-850-the-fake-review-hunter for a description of low-quality sellers imitating high-quality sellers via fake reviews in the nutritional supplement industry. Amazon is disinclined to comment on the presence of these fake reviews despite credible evidence that they exist.

[2] The computer game sales platform Steam has notoriously low barriers to entry, with some games being sold without an executable file. It would be a relatively simple matter to ensure that a game must at least run before the platform allows it to be sold. See https://www.gamerevolution.com/news/373453-steam-poorly-moderated-selling-empty-game-folder for more detail.

[3] One could equivalently think of this as an increase in the effective search cost.

the platform can increase their profits, and consequently its own revenue, at the cost of driving some consumers away from the platform. I refer to this softening of competition and consequent increase in prices resulting from reduced consumer search intensity as the *obfuscation effect*.

The low-quality sellers in my model attempt to fool consumers into believing that they are high-quality sellers. Consumers know the probability that low-quality sellers successfully imitate high-quality sellers, but cannot distinguish between a high-quality seller and a low-quality seller who is successful in their deception. Thus, in addition to the search obfuscation effect, the presence of low-quality sellers also reduces consumers' confidence in the quality of any prospect they are considering because it is possible that the prospect is a low-quality seller in disguise. I define this loss of confidence as the *lemon effect*, where consumers have diminished willingness to pay for sellers' products and are less eager to participate in the platform's market because of the possibility that they may be scammed. The lemon effect is purely harmful to the platform, so if it is significantly stronger than the search obfuscation effect then the platform does not admit any low-quality sellers.

The common practice among such platforms of highlighting recommended sellers might initially seem to be contradictory to the idea that the platform is obfuscating search, but in fact the ability to recommend a seller (modeled here as the platform choosing a seller to make "prominent" in the sense of Armstrong, Vickers, and Zhou (2009)) and obfuscating search via low-quality sellers are *complementary strategies*. It is profit maximizing for the platform to recommend a high quality seller, so the recommended product is not subject to the lemon effect and consumers are willing to pay a higher price for the recommended seller's product than the product of a random seller in a market with low quality sellers. Both consumers and the platform are better off with a recommended seller for any fixed proportion of low-quality sellers. However, by mitigating the lemon effect the recommended seller reduces the cost to the platform of admitting low-quality sellers. Therefore consumers' gains from the recommendation may be mitigated by an increase in the proportion of low-quality sellers. The ability to recommend a seller can even cause the platform to admit low-quality sellers under parameter sets where it would not absent the ability to recommend.

Competition between platforms has an ambiguous effect on the equilibrium proportion of low-quality sellers relative to the monopoly model. Similar to the results of Hagiu and Jullien (2014), I show that entry can *increase* the proportion of low-quality sellers. Comparing a duopoly to a model with only one platform, each consumer will be affiliated with a platform which is weakly preferred to the option they would choose in a monopoly setting. The proportion of low-quality sellers increases with entry if this ability to choose a more preferred option means that consumers are more loyal to that option than in the monopoly setting. This happens when elasticity of consumer participation in a platform's market with regard to the proportion of low-quality sellers decreases as a result of entry. If instead of modeling entry we think of an increase in competition as consumers viewing the platforms as closer substitutes, (modeled in Section 5 as a reduction in transportation costs along a Hotelling line) then this elasticity increases as the platforms become more substitutable, which reduces the equilibrium proportion of low-quality sellers.

2 Related Literature

This article primarily relates to three broad strands of the literature: search obfuscation, noisy search/search diversion, and search design in platforms. The idea behind search obfuscation starts with the observation that increasing search costs often lead to higher prices and larger profits, with the extreme case being Diamond's (1971) result that firms will price as local monopolies if consumers can be completely prevented from searching. Gabaix and Laibson (2006) show that firms can take advantage of myopic consumers by shrouding information. Ellison and Ellison (2009) and Ellison and Wolitzky (2012) demonstrate that coordination on increasing search costs can survive even with sophisticated consumers. Although as Ellison and Wolitzky point out, it is not clear why firms would not simply coordinate on price if they can coordinate on increasing search costs[4]. Armstrong and Zhou (2015) explain the phenomenon of "exploding offers" in both labor and consumer search markets as an attempt by firms to deter search. The obfuscation in my model is not due to seller level behavior but instead a measure taken by the market designer to regain some of the price control it sacrifices by not selling directly to consumers[5]. The obfuscation acts as a coordination device between sellers and the platform similar to Lubensky (2017). The ability of the platform to mitigate competition between sellers means that I also contribute to the literature on negative within group effects on platforms (Weyl 2010; Belleflamme and Peitz 2019).

The search obfuscation in my model is mechanically similar to the noisy search in Yang (2013). Encountering a low-quality seller is roughly equivalent to an irrelevant search in that model, which means that reducing the proportion of low-quality sellers is similarly equivalent to increasing the precision of search. In that paper the precision of search is not a strategic variable, whereas I allow the platform in my model to choose the proportion of low quality sellers.

There have been other articles which examine motivations for platforms to reduce search quality. However, the benefits to the platform in my model from introducing

[4] de Roos (2018) shows that limited product comparability can aid price collusion. Therefore, price collusion and obfuscation together may be easier to justify than either alone.

[5] I do not address the question of why the business chooses to operate as a platform instead of selling directly. See Hagiu and Wright (2015) for a discussion of the tradeoffs between selling as a platform and direct sales.

noise to the search process are distinctly unlike those from the "search diversion" described by Hagiu and Jullien (2011) and White (2013). With search diversion, the platform garbles the search process in order to expose consumers to a product which is separate from the product they wish to find, but which gives the platform a higher profit margin. The platform's tradeoff from search diversion is similar to the tradeoff in my model in that garbling search reduces the appeal of the platform to consumers, but the benefit to a platform from search diversion takes the form of increased trading volume for high margin products. Whereas in my model the benefit to the platform instead arises from higher market prices for all products.

In terms of search design in platforms, Dinerstein et al. (2018) show that search design has a significant impact on seller pricing decisions in platform markets. They describe how a change in eBay's search design which lowered search costs and increased comparability of search results reduced price dispersion and equilibrium prices. Similarly, Nosko and Tadelis (2015) demonstrate that the quality of sellers on a platform has a significant impact on consumer retention. In the theoretical literature, Wang and Wright (2018) create a model with a platform search environment broadly similar to mine, but they focus on contracting problems between the platform and sellers.

Entry by a competing platform has an ambiguous effect on equilibrium obfuscation in my model, but increased intensity of competition between already active platforms will always lead to less obfuscation. This finding matches Hagiu and Jullien's (2014) result that weak competition can lead to more search diversion than would exist in a monopoly market, whereas strong competition deters it. It is also similar to Hagiu and Hałaburda's (2014) result that a competitive market leads to more information revelation by platforms.

Finally, I show that the platform can credibly recommend high-quality sellers even though these sellers give the platform more profit. My finding that the recommendation can improve consumer welfare has empirical support from Chen and Yao's (2016) result that following platform recommended search order can significantly reduce search costs with little loss in expected match utility. The platform's credibility in my model comes from the fact that the recommendation helps the platform mitigate the lemon effect without reducing the search obfuscation effect. This result differs significantly from Yea's (2018) cheap talk model where recommendation is only credible if all of the platform's messages are revenue equivalent in equilibrium. Che and Hörner (2018) also explore platform recommendation systems, but their recommender is a benevolent planner maximizing social welfare rather than a profit maximizing sales platform.

3 The Monopoly Model

The agents in the monopoly model consist of a monopoly platform, a unit mass of consumers, each with unit demand, a continuum of high-quality sellers, and a continuum of low-quality sellers, both with strictly positive total mass[6].

Consumers decide whether to shop on the platform, those who do not participate in the platform's market have access to an outside option. The population distribution of the utility for this outside option follows the continuous and concave CDF $Q(\cdot)$. Let α be the proportion of low quality sellers in the platform's market, and $V(\alpha)$ the expected value to a consumer of participating in the platform's market given α. Consumers will participate in the platform's market if $V(\alpha)$ is greater than their value of the outside option, so the mass of participating consumers is $Q\left(V(\alpha)\right)$.

[6] In equilibrium all sellers charge the same price, so the platform only cares about the mass of transactions. Unit demand implies that consumers only consider the proportion of high vs. low-quality sellers on the platform and its effect on the search process. The model does not include an entry or operating cost, so the scaling effect from relative mass of sellers to consumers on seller profits does not change any seller decisions. For all of these reasons, the exact mass of sellers in the game or on the platform does not matter, but for the sake of notational simplicity I assume that the mass of sellers on the platform is 1.

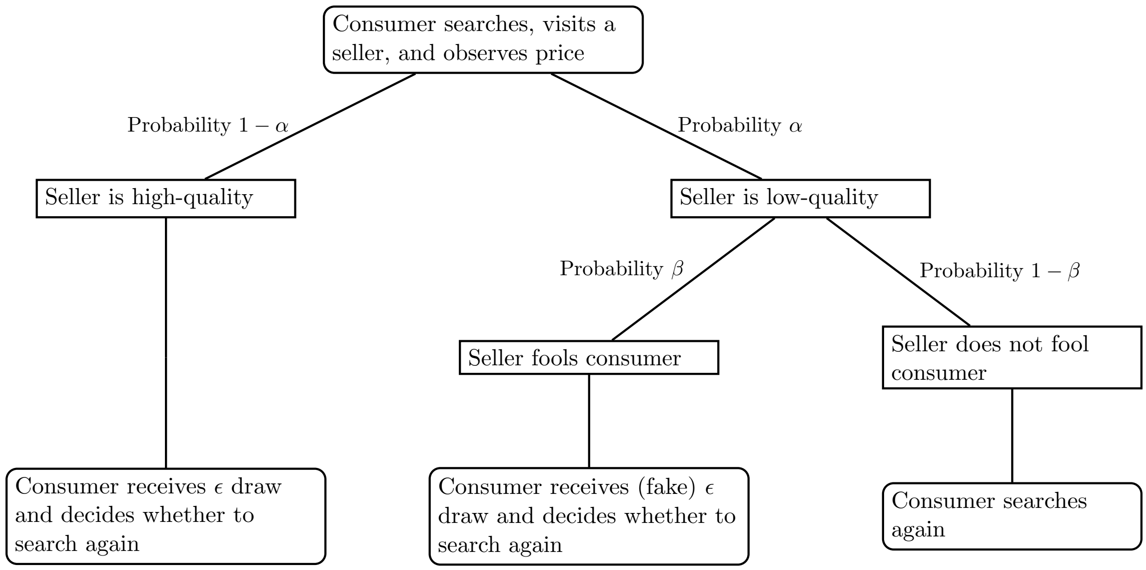

Figure 1: The process by which consumers learn about seller quality and their match value. Rounded corners indicate consumer decisions, right angled corners indicate random events.

Those consumers who do decide to participate face a search problem similar to Wolinsky (1986) or Anderson and Renault (1999). Consumer i's utility when purchasing from seller j charging price p_j is

$$-p_j + \epsilon_{ij} \tag{1}$$

ϵ_{ij} is a consumer-seller specific match utility term. Consumers cannot observe p or ϵ before visiting a seller and can only purchase from sellers they have visited. For high-quality sellers ϵ is distributed according to the log-concave distribution function $f(\epsilon)$ with support $[\underline{\epsilon}, \bar{\epsilon}]$, $\underline{\epsilon} \geq 0$ and $\bar{\epsilon} \in (\underline{\epsilon}, \infty)$[7].

For low-quality sellers $\epsilon = 0$, so consumers never purchase from a seller they know to be of low-quality. Low-quality sellers attempt to imitate high-quality sellers and successfully fool a visiting consumer with exogenous probability $\beta > 0$. Thus, upon

[7]Strictly speaking the assumption that $\underline{\epsilon} \geq 0$ isn't necessary for my results, but I include it to prevent confusion regarding the distinction between low- and high-quality sellers.

visiting a low-quality seller, consumers do not purchase with probability $1 - \beta$ and with probability β the consumer incorrectly believes the low-quality seller to be a high-quality seller. Upon successfully deceiving a consumer the low-quality seller is indistinguishable from a high-quality seller, but this does not necessarily mean that the consumer will purchase from that seller. A consumer who is fooled by a low-quality seller receives a fake match utility signal which follows $f(\epsilon)$ and represents the extent to which the sales tactics which successfully deceive a consumer appeal to that consumer. Consumers cannot distinguish between a true match utility draw from a high-quality seller and this fake match utility signal from the successful low-quality sellers. The consumer purchases from a low-quality seller only if that low-quality seller successfully fools the consumer *and* the perceived match utility is sufficiently high. This process is shown as a flowchart in Figure 1.

The platform receives an exogenous proportion of sellers' revenue and sets α to maximize profits. In order to focus on the primary mechanisms in my model I assume that the platform can perfectly and costlessly determine seller quality and so does not face any sort of screening cost when changing the proportion of low-quality sellers; see Section 6 for a discussion of screening costs and their implications for the interpretation of this model. I assume that α is publicly observable[8]. Sellers' only decision is what price to charge given α.

3.1 Equilibrium

As is standard in the search literature I use symmetric weak perfect Bayesian equilibrium for my solution concept and assume that searchers' beliefs about other sellers' prices are not affected by observing one seller deviate from equilibrium. The timing of

[8]Public observability here represents the reputation costs to the platform of allowing low quality sellers.

the model proceeds as follows:

1. The platform sets the proportion of high-quality and low-quality sellers and this proportion is observable to consumers.

2. Consumers commit to either participating in the platform's market or remaining with the outside option[9].

3. Sellers admitted by the platform set prices simultaneously.

4. Consumers participating in the platform's market search among sellers and make purchasing decisions.

Consumer Search Behavior

Consumers are aware of the presence of low-quality sellers, so if the proportion of low-quality sellers on the platform is α, then a consumer who has searched and found a prospect $-p_j + \epsilon_{ij}$ will evaluate the value of this current option as

$$\frac{1-\alpha}{\beta\alpha+1-\alpha}(-p_j+\epsilon_{ij}) - \frac{\beta\alpha}{\beta\alpha+1-\alpha}p_j = \frac{1-\alpha}{\beta\alpha+1-\alpha}\epsilon_{ij} - p_j$$

The term $\frac{1-\alpha}{\beta\alpha+1-\alpha}$ is the mathematical representation of the lemon effect, so I call it the *lemon coefficient*. It represents the fact that, conditional on a consumer evaluating a product as high-quality, the probability that the product they are observing is actually high quality is the probability $1-\alpha$ that they encountered a high-quality seller, divided by the probability $\beta\alpha + 1 - \alpha$ of the event that they evaluate a product as high quality.

[9] Commitment is not a binding assumption in the monopoly benchmark model since the search environment is stationary, so if consumers find participation initially worthwhile they will always find it worthwhile. The main intuition should still hold if search is not stationary. See Burdett and Vishwanath (1988) or Casner (2018) for a discussion of the effects of a non-stationary search environment.

This coefficient is decreasing in β because the more able the low-quality sellers are to imitate the high-quality sellers, the less the consumer will believe that the product they are considering is actually high-quality.

Search is undirected, so a searching consumer will randomly select a new seller to visit[10]. If a consumer chooses to search for another prospect, they find a low-quality seller with probability α, and conditional on visiting a low-quality seller they correctly determine that this seller is low-quality with probability $1-\beta$, so with probability $\alpha(1-\beta)$ the consumer must search at least once more. Those consumers who search twice need to search a third time with probability $\alpha(1-\beta)$ and so on. Let c denote the (strictly positive) cost of searching, the expected cost of finding an additional prospect is then

$$\begin{aligned}\sum_{k=1}^{\infty}[\alpha(1-\beta)]^k c &= \frac{c}{1-\alpha(1-\beta)} \\ &= \frac{c}{\beta\alpha+1-\alpha}\end{aligned}$$

Denote by p the price that sellers set in a symmetric equilibrium and ϵ' the hypothetical match utility draw at a new prospect. Consumers expect that all unobserved sellers set price p so the expected payoff from finding another prospect is

$$\int_{\underline{\epsilon}}^{\bar{\epsilon}} \max\left[\frac{1-\alpha}{\beta\alpha+1-\alpha}\epsilon_{ij}-p_j, \frac{1-\alpha}{\beta\alpha+1-\alpha}\epsilon'-p\right] f(\epsilon')d\epsilon' - \frac{c}{\beta\alpha+1-\alpha}$$

A searcher can recall past prospects freely, so the new prospect is preferred only if $\epsilon' > \epsilon_{ij} + \frac{(\beta\alpha+1-\alpha)(p-p_j)}{(1-\alpha)}$[11]. Setting the utility of searching equal to the utility of staying

[10] High-quality sellers in this environment face a symmetric problem and will hence set symmetric prices. Low-quality sellers only make a profit if they successfully imitate a high-quality seller, so they must set the same price as the high quality sellers. Since all sellers are ex ante identical, consumers have no incentive to engage in ordered search if doing so requires even the slightest effort.

[11] I use this assumption of free recall because it matches the previous literature and it slightly eases derivation. In equilibrium no consumer exercises recall so there are no qualitative differences between allowing recall vs. no recall for this model.

with the current prospect and simplifying, consumers are indifferent between staying with the current prospect and searching again if

$$c = \int_{\epsilon_{ij} + \frac{(\beta\alpha+1-\alpha)(p-p_j)}{(1-\alpha)}}^{\bar{\epsilon}} \left[(1-\alpha)\left(\epsilon' - \epsilon_{ij}\right) + (\beta\alpha + 1 - \alpha)(p_j - p)\right] f(\epsilon')d\epsilon'$$

Rearranging and simplifying

$$\Longrightarrow c = (1-\alpha) \int_{\epsilon_{ij} + \frac{(\beta\alpha+1-\alpha)(p-p_j)}{(1-\alpha)}}^{\bar{\epsilon}} \left[\epsilon' - \epsilon_{ij} - \frac{\beta\alpha + 1 - \alpha}{(1-\alpha)}(p - p_j)\right] f(\epsilon')d\epsilon' \tag{2}$$

Define

price, so the stopping rule is given by

$$c = (1-\alpha) \int_{U}^{\bar{\epsilon}} \left(\epsilon' - U\right) f(\epsilon')d\epsilon' \tag{3}$$

and consumers will stop at a seller if they draw $\epsilon > U$. U is consumers' reservation match value and depends on α only through the latter's effect on Equation (3).

Lemma 1. *In a symmetric equilibrium the reservation match value U is decreasing in the proportion of low quality sellers α.*

Proof. All proofs are relegated to the appendix. ■

The right hand side of Equation (3) is decreasing in U. From the assumption of symmetry in prices, U does not change with α except through its effect on the stopping rule, so the RHS of Equation (3) is also decreasing in α. If α increases, then U must decrease in order to maintain equality. Intuitively, as the cost of searching increases relative to the marginal value, consumers are willing to stop after observing a lower match value. The possibility of encountering a low-quality seller then having to search again increases the effective search cost and hence reduces consumer search intensity. This decrease in search persistence as the proportion of low-quality sellers increases is the mathematical expression of the search obfuscation effect defined in Section 1.

It is worth noting that β has no effect on search behavior once a consumer is participating in the platform's market. This is because the probability that the consumer is being scammed is the same across all prospects. Thus, while the lemon effect reduces the consumers' perceived valuation of each prospect, it reduces their valuation *equally* across all prospects, and so does not influence search behavior. This result relies on ex-ante symmetry across sellers and does not hold for the recommended seller in Section 4.

Seller pricing

Sellers' only decision once they join a platform is what price to charge. Since consumer beliefs about other sellers' prices are not affected by observing deviations, they use the stopping rule derived above. Note that U is derived given consumers' expectations over prices. The precise implication of this stopping rule is that consumers purchase any observed prospect whose net expected utility exceeds $\frac{1-\alpha}{\beta\alpha+1-\alpha}U - p$[12]. A consumer will stop at a seller who sets $p_j > p$ only if the match utility draw is high enough so that the net expected utility from purchasing exceeds this cutoff. Thus the probability that a given consumer who visits apparently high-quality seller j charging price p_j buys from that seller is the probability that $\frac{1-\alpha}{\beta\alpha+1-\alpha}\epsilon_{ij} - p_j > \frac{1-\alpha}{\beta\alpha+1-\alpha}U - p$, or $1 - F\left(U + \frac{\beta\alpha+1-\alpha}{(1-\alpha)}(p_j - p)\right)$.

Let γ denote the probability that a consumer stops after visiting an arbitrary

[12]While it would be more intuitive to derive equilibrium prices if the stopping rule were expressed in terms of this net expected utility (e.g. Bar-Isaac, Caruana, and Cuñat (2012) do so), the gains from doing so would be more than offset by unwieldy expressions later on.

seller[13]. With probability $1-\gamma$ consumers leave the first seller they visit and so a mass $Q\left(V(\alpha)\right)(1-\gamma)$ will evaluate at least two prospects, $Q\left(V(\alpha)\right)(1-\gamma)^2$ three prospects and so on. Taking this to the limit as in Anderson and Renault (1999) or Bar-Isaac, Caruana, and Cuñat (2012), the mass of consumers who visit each seller is $\frac{Q(V(\alpha))}{\gamma}$ and this mass will not be affected by an individual sellers' deviation since each seller is infinitesimal compared to the size of the market[14]. The platform takes an exogenous percentage ξ of seller revenue, and the marginal cost of production is 0. high-quality seller profit is then

$$\pi = p_j \frac{Q\left(V(\alpha)\right)(1-\xi)}{\gamma}\left[1 - F\left(U + \frac{\beta\alpha + 1 - \alpha}{(1-\alpha)}(p_j - p)\right)\right]$$

and low-quality seller profit is $\beta\pi$. It is a matter of simple calculus to find that the profit maximizing price in both cases is determined implicitly by

$$p_j = \frac{(1-\alpha)}{\beta\alpha + 1 - \alpha}\frac{1 - F\left(U + \frac{\beta\alpha+1-\alpha}{(1-\alpha)}(p_j - p)\right)}{f\left(U + \frac{\beta\alpha+1-\alpha}{(1-\alpha)}(p_j - p)\right)} \tag{4}$$

Symmetry then implies $p_j = p$ so

$$p = \frac{(1-\alpha)}{\beta\alpha + 1 - \alpha}\frac{1 - F(U)}{f(U)} \tag{5}$$

Log concavity of $f(\cdot)$ ensures sufficiency of the first order condition[15]. The lemon coefficient appears in the price as well as consumers' evaluation of their match utility.

[13] In equilibrium this probability is $(\alpha\beta + 1 - \alpha)(1 - F(U))$, but since that probability is based on other sellers' prices I introduce γ here to emphasize that seller j's price is disciplined only by the proportion of consumers who stop after visiting j.

[14] Technically the mass should be this fraction divided by the total mass of sellers, but as mentioned in Section 3 I assume this latter mass to be 1 since the ratio of buyers to sellers is irrelevant for all agents' decisions.

[15] For more detail, see Bagnoli and Bergstrom (2005)

As consumers' faith in the product they are purchasing decreases, sellers must lower their prices to compensate. On the other hand, the search obfuscation effect reduces the appeal of moving on to a different seller, which reduces U and pushes equilibrium prices upward. If the obfuscation effect dominates the lemon effect, which happens when low-quality sellers are not too adept at imitating high-quality sellers (β small), then softened competition causes prices to increase with the proportion of low-quality sellers α. This logic is formalized in Lemma 2

Lemma 2. *Equilibrium prices are increasing in the proportion of low-quality sellers if their probability of fooling any given consumer is sufficiently low. Formally: at any given level of α, p is increasing in α for β sufficiently small.*

Consumer Surplus and Platform Profit

A consumer who participates on the market will stop after observing $\epsilon > U$. With probability $\frac{1-\alpha}{\beta\alpha+1-\alpha}$ this observation comes from a high-quality seller, so the expected match utility is $\frac{1-\alpha}{\beta\alpha+1-\alpha}\int_U^{\bar{\epsilon}} \epsilon \frac{f(\epsilon)}{1-F(U)} d\epsilon$. The expected number of prospects observed is $\frac{1}{1-F(U)}$, with $\frac{1}{\beta\alpha+1-\alpha}$ searches required to find a prospect in expectation, so the expected search cost is $\frac{c}{(1-F(U))(\beta\alpha+1-\alpha)}$. Finally, all sellers will charge the same price p, so

$$V(\alpha) = \frac{1-\alpha}{\beta\alpha+1-\alpha}\int_U^{\bar{\epsilon}} \epsilon \frac{f(\epsilon)}{1-F(U)} d\epsilon - \frac{c}{(1-F(U))(\beta\alpha+1-\alpha)} - p$$

Using Equation (3) and canceling like terms

$$V(\alpha) = \frac{(1-\alpha)\int_U^{\bar{\epsilon}} U f(\epsilon) d\epsilon}{(1-F(U))(\beta\alpha+1-\alpha)} - p \tag{6}$$

From Equation (4) and collecting like terms

$$V(\alpha) = \frac{(1-\alpha)}{\beta\alpha + 1 - \alpha}\left(U - \frac{1 - F(U)}{f(U)}\right) \tag{7}$$

Both the lemon effect and search obfuscation reduce $V(\alpha)$. The term $U - \frac{1-F(U)}{f(U)}$ is the consumer surplus retained by consumers who purchase from high-quality sellers. This surplus is not affected by the lemon effect, but because search obfuscation reduces U, consumers will receive less surplus as a result of reduced search intensity as the proportion of low-quality sellers increases. Receiving positive surplus is conditional on purchasing from a high-quality seller, so the expected value of participating in the platform's market is scaled down by the lemon coefficient. The more likely consumers are to be ripped off by a low-quality seller, the lower their expected payoff.

This clearly illustrates the difference between the obfuscation and lemon effects. The obfuscation effect allows the sellers to capture more of the expected surplus available in the market. This reduces participation but can be beneficial to the platform because it receives a share of the increase in sellers' profits. On the other hand, the lemon effect reduces the total amount of expected surplus, which benefits no one.

Every consumer who shops on the platform will purchase at the symmetric equilibrium price p. The platform operates costlessly, so the platform's profit is

$$\Pi = \xi p Q(V(\alpha))$$

Because ξ is an exogenous scalar, maximizing Π is equivalent to maximizing $\frac{\Pi}{\xi} = pQ(V(\alpha))$, so to simplify notation I suppress ξ unless it is relevant.

Proposition 1. *The platform admits a positive proportion of low-quality sellers if* *the search cost is sufficiently low and low-quality sellers are not too adept at imitating high-quality sellers:* $\alpha > 0$ *in equilibrium if both conditions 1 and 2 hold.*

1. β *is sufficiently small*

2. c *is sufficiently small and* $\max_\epsilon f'(\epsilon)$ *is sufficiently close to 0*

Intuitively, Proposition 1 says that the platform will admit low-quality sellers if the competition between sellers on the platform is sufficiently intense (so that the increased prices from softening competition compensate for the reduction in consumer participation from search obfuscation) and the loss in consumer confidence stemming from the possibility of being fooled by the low-quality sellers is not too large (so that the the reduction in consumer participation from the lemon effect does not eliminate the benefits to the platform from search obfuscation). The condition on $f'(\epsilon)$ ensures that the net effect of softening competition overcomes the reduction in demand from the search obfuscation effect when the proportion of low-quality sellers $\alpha = 0$ and competition without search obfuscation is sufficiently intense[16].

Example

Suppose $\epsilon \sim U[0,1]$ for the high-quality sellers and $Q(V(\alpha)) = V(\alpha)$ (equivalent to assuming that the platform is a monopoly at one end of a Hotelling line with transport cost 1). Then we can solve Equation (3) explicitly to find that $U = 1 - \sqrt{\frac{c}{1-\alpha}}$.

[16] Requiring $f'(\cdot)$ small is in fact stronger than necessary. Strictly speaking the necessary condition is $\lim_{U \to \bar{\epsilon}} \frac{Q'(V(0))}{Q(V(0))} \frac{\left(1 + \frac{f(U)^2 + f'(U)(1-F(U))}{f(U)^2}\right)}{\frac{f(U)^2 + f'(U)(1-F(U))}{f(U)^2}} < \infty$, but this condition is unwieldy and does not add to the intuition of the model.

Plugging this into Equation (4) and Equation (7) we get that

$$p = \frac{1-\alpha}{\beta\alpha + 1 - \alpha}\sqrt{\frac{c}{1-\alpha}}$$
$$V(\alpha) = \frac{1-\alpha}{\beta\alpha + 1 - \alpha}\left(1 - 2\sqrt{\frac{c}{1-\alpha}}\right)$$

Consumer participation is decreasing in both c and α, and for $\frac{c}{1-\alpha} > \frac{1}{4}$ no consumers find it worthwhile to participate in the platform's market at all. We can combine these findings to give platform profit

$$\Pi = \frac{\xi}{t}\left(\frac{1-\alpha}{\beta\alpha + 1 - \alpha}\right)^2 \sqrt{\frac{c}{1-\alpha}}\left(1 - 2\sqrt{\frac{c}{1-\alpha}}\right)$$

Figure 2 shows numerical calculations for the derivative of platform profit with regard to the proportion of low-quality sellers at $\alpha = 0$ for a range of parameter values. We can see that the marginal profit from admitting low-quality sellers is monotonically decreasing in both β and c. The platform will only admit low-quality sellers if this marginal profit is positive, which holds only if both c and β are small (the lower left corner of the graph). For $c > 0.06$ or $\beta > 0.2$ the platform will never admit a positive proportion of low-quality sellers.

It may seem like the platform only admits $\alpha > 0$ in a relatively small set of the parameter space. However, note that for $c = 0.25$ no consumers would participate in the platform's market even at $\alpha = 0$, and market platforms tend to have search costs that are quite low relative to the purchase price (Blake, Nosko, and Tadelis 2016), so it is entirely plausible that most platforms would operate close to the horizontal axis.

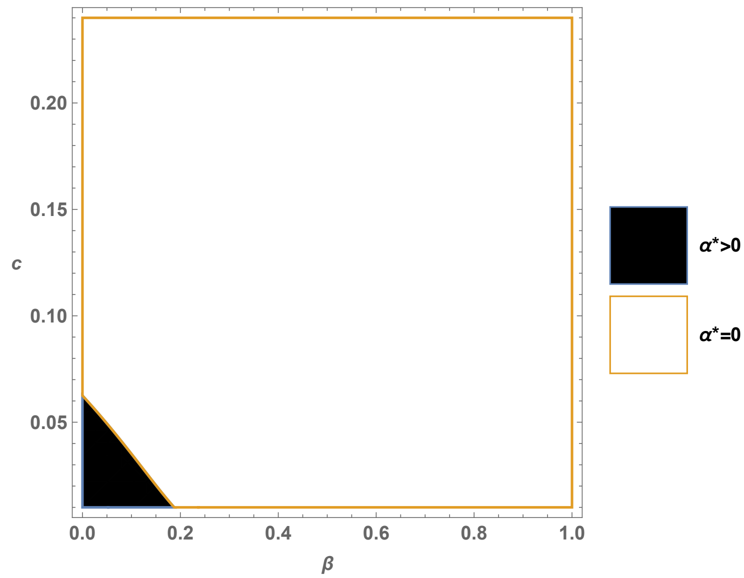

Figure 2: A parameter space showing the values where the platform will admit a positive proportion of low-quality sellers in this numerical example. The platform admits a positive proportion of low-quality sellers if the derivative of profit with regard to α is positive at $\alpha = 0$. When parameters lie in the shaded region, the platform sets $\alpha > 0$ in equilibrium, but it admits no low-quality sellers in the white region.

Equilibrium Curation Decision

Assuming that conditions 1 and 2 of Proposition 1 are satisfied is equivalent to assuming that the platform will set $\alpha \in (0, 1)$, because if the platform sets $\alpha = 1$ then no consumer will participate in the platform's market, so the only possible corner solution is at $\alpha = 0$, which is precluded by Proposition 1. Define an *interior equilibrium* as any equilibrium with $\alpha \in (0, 1)$. The following first order condition for the platform's curation decision (derived from setting the derivative of platform profits equal to 0 and rearranging) is a necessary condition in any interior equilibrium.

$$\frac{\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}}{p\left(1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)} = \frac{Q'(V(\alpha))}{Q(V(\alpha))} + \beta B(\alpha) \tag{8}$$

Where

$$B(\alpha) = \frac{(1-F(U))}{(1-\alpha)\int_U^{\bar{\epsilon}}(\epsilon - U)f(\epsilon)d\epsilon} \frac{Q(V(\alpha))\frac{1-F(U)}{f(U)} + pQ'(V(\alpha))\left(U - \frac{1-F(U)}{f(U)}\right)}{Q(V(\alpha))\frac{1-F(U)}{f(U)}\left(1 + \frac{f(U)^2 + f'(U)(1-F(U))}{f(U)^2}\right)} \tag{9}$$

Roughly speaking, the left hand side of Equation (8) represents the elasticity of market prices with regard to the proportion of low-quality sellers, and the first term on the right is the elasticity of consumers' participation in the platform. $B(\alpha)$ is the distortion in the platform's market caused by the lemon effect. It is the result of β and consumers' expectations about β. It increases the marginal costs of admitting more low-quality sellers, and so reduces α in equilibrium, but since any interior equilibrium requires β to be small, the effect of $B(\alpha)$ must be negligible when the platform is admitting positive α. The following condition is stronger than needed to give sufficiency of the first order condition, but is nevertheless not terribly restrictive[17]:

Assumption 1. *The variance of $f(\cdot)$ is sufficiently large so that Π is single peaked in α*

Let α^* denote the proportion of low-quality sellers in equilibrium.

Lemma 3. *Under Assumption 1, α^* is determined implicitly and uniquely by Equation (8) for any interior equilibrium.*

Lemma 3 allows me to use Equation (8) to evaluate the comparative statics of α^*.

Proposition 2. *Under Assumption 1, for any interior equilibrium*

1. *The proportion of low-quality sellers decreases in search cost: α^* is decreasing in c.*

[17]See the proof of Lemma 3 for details of why this assumption is sufficient for single peakedness of platform profit in the curation decision.

2. *If $Q(V(\alpha))$ is linear (as in the hotelling example) then the proportion of low-quality sellers decreases as they become more adept at imitating the high-quality sellers: α^* is decreasing in β for linear $Q(V(\alpha))$.*

The obfuscation effect from the low-quality sellers effectively increases the search cost, but comes at the cost of lost consumer participation and lower consumer confidence. Increasing c does not contribute to the lemon effect, but does increase prices, so if the search cost increases, then the platform has less incentive to increase the effective cost further via low quality sellers. Increasing the search cost does not contribute to the lemon effect so directly increasing this cost seems to be a preferable option for the platform rather than obfuscating search via low-quality sellers. However, one can easily imagine that an action which is so obviously hostile to consumers could result in a public relations nightmare and potentially even regulatory scrutiny. Even if consumers and regulators were to remain complacent about an increase in search costs, many platforms have a number of different markets which share a common search environment. For example, the market for high-end headphones on Amazon is likely much less competitive than the market for exercise supplements. Thus, allowing more low-quality sellers in the supplement market would allow Amazon to increase the search cost in that market without also increasing search costs in the headphone market. Additionally, if screening costs are a factor (see Section 6) then obfuscation via low-quality sellers may be more cost effective than direct obfuscation via increased search costs.

Potential low-quality sellers who do not get to participate in the platform are worse off with a higher search cost, but the platform and all sellers (both low- and high-quality) participating in the platform are better off. Sellers are able to charge a higher price because the proportion of low-quality sellers is lower which reduces the lemon effect and increases total profits. Interestingly, the consumers can be better off with a higher search cost as well. The probability that they encounter and then trade with a low-quality seller goes down, and the effect of this increased confidence relative to an equilibrium with lower search cost increases market participation.

As β increases, the lemon effect becomes stronger compared to the search obfuscation effect, so the marginal cost of admitting additional low-quality sellers increases and α^* decreases. Linearity of $Q(\cdot)$ ensures that the effects of consumers' decreased willingness to pay are outweighed by their decreased market participation. If $Q(\cdot)$ is highly concave, then $Q'(\cdot)$ might be small enough relative to $Q(\cdot)$ that the price reduction from an increase in β can be stronger than the effect on demand and an increase in β can *increase* α^* because the platform would rather serve fewer consumers at a higher market price. However, no matter what the effect on α^*, both the platform and sellers are worse off as β increases because it causes consumers' confidence to decrease more quickly as α increases. As with c, the effect of β on consumer welfare is ambiguous. Directly, it reduces welfare by making consumers more vulnerable to the low-quality sellers, but if the reduction in α^* is large enough then consumers can be made better off in equilibrium because the number of low-quality sellers they encounter decreases.

4 Recommending a Seller

Recommending certain sellers to consumers seems like behavior contradictory to increasing search costs by admitting low-quality sellers. However, if the platform can credibly recommend a high-quality seller to consumers then this behavior is in fact a *complementary* strategy to admitting low-quality sellers.

I model the recommendation process by having the platform choose a high-quality seller to be *prominent* in the sense of Armstrong, Vickers, and Zhou (2009). All consumers visit the recommended seller first at no search cost, but if they choose to move on and visit other sellers then their search process is undirected sequential search as in the previous section. As will become apparent below, the platform is strictly better off recommending a high-quality seller than a low-quality seller, so consumers view this recommendation as credible and believe with probability 1 that the recommended seller is high-quality. Denote the recommended seller by R, in this case consumer i's evaluation of the utility from the recommended seller given match value draw ϵ_{iR} and price p_R is simply

$$-p_R + \epsilon_{iR}$$

The value of searching for another prospect is quite similar to the basic monopoly model

$$\int_{\underline{\epsilon}}^{\bar{\epsilon}} \max\left[-p_R + \epsilon_{iR}, \frac{1-\alpha}{\beta\alpha + 1 - \alpha}\epsilon' - p\right] f(\epsilon')d\epsilon' - \frac{c}{\beta\alpha + 1 - \alpha}$$

Where p is the equilibrium price of the non-recommended sellers. The new prospect is preferable if

$$\epsilon' > \frac{\beta\alpha + 1 - \alpha}{1 - \alpha}\left(\epsilon_{iR} + p - p_R\right)$$

By logic almost identical to the previous section, the stopping rule is then given by

$$c = (1-\alpha)\int_{\frac{\beta\alpha+1-\alpha}{1-\alpha}(\epsilon_{iR}+p-p_R)}^{\bar{\epsilon}} \left[\epsilon' - \frac{\beta\alpha + 1 - \alpha}{1-\alpha}\left(\epsilon_{iR} + p - p_R\right)\right] f(\epsilon')d\epsilon' \tag{10}$$

Define U_R as the ϵ_{iR} which solves Equation (10). Then

$$c = (1-\alpha)\int_{\frac{\beta\alpha+1-\alpha}{1-\alpha}(U_R+p-p_R)}^{\bar{\epsilon}} \left[\epsilon' - \frac{\beta\alpha+1-\alpha}{1-\alpha}(U_R+p-p_R)\right] f(\epsilon')d\epsilon' \tag{11}$$

As in Section 3, U_R as derived above is based on the recommended seller charging the expected price p_R. Equation (11) implies that consumers will stop at the recommended seller if the net utility they receive from doing so exceeds $U_R - p_R$. Suppose the recommended seller deviates to price p'_R, then consumers stop at the recommended seller if $\epsilon_{iR} - p'_R > U_R - p_R$ and the probability that these consumers stay after visiting is $1 - F(U_R + p'_R - p_R)$. The mass of consumers participating in the platform's market is $Q(V(\alpha))$, so the recommended seller's profit is given by

$$\pi_R = p'_R(1-\xi)Q(V(\alpha))(1 - F(U_R + p'_R - p_R)) \tag{12}$$

The recommended seller's decision does not change participation decisions by consumers because they commit to the platform before observing sellers' actual prices[18]. The profit maximizing price is thus given by

$$p'_R = \frac{1 - F(U_R + p'_R - p_R)}{f(U_R + p'_R - p_R)}$$

In equilibrium consumers' expectations about the price are correct so

$$p_R = \frac{1 - F(U_R)}{f(U_R)} \tag{13}$$

Note the absence of the lemon coefficient. Consumers believe with certainty that the

[18]the commitment assumption has more effect than in the baseline because the search environment is no longer completely stationary. Allowing a subset of consumers to leave the platform after visiting the recommended seller would yield broadly similar results with similar intuition.

recommended seller is high-quality, so α influences the recommended seller's price only through U_R. Consumers who leave U_R and visit other sellers no longer consider the recommended seller, so the other sellers compete and price as in the previous section, and consumers use the stopping rule defined by Equation (3)[19]. Consumers' greater confidence in the recommended seller's product gives the recommended seller a competitive advantage which is formalized in Lemma 4.

Lemma 4. *In the equilibrium with a recommended seller*

1. *Consumers visiting sellers after the recommended seller have the same reservation match value U as in Section 3 and these other sellers charge the same price p.*

2. *The net expected utility above which the consumer stops searching is the same at the recommended seller and the other sellers: $U_R - p_R = \frac{1-\alpha}{\beta\alpha+1-\alpha}U - p$.*

3. *The recommended seller sets a higher price than the other sellers if the proportion of low-quality sellers is positive: $p_R > p$ if $\alpha > 0$.*

4. *The recommended seller sets an identical price to the other sellers and the reservation match values are identical if there are no low-quality sellers: $U_R = U$ and $p_R = p$ if $\alpha = 0$.*

Define $V^R(\alpha)$ as the expected value to consumers of participating in the platform's market when there is a recommended seller. Solving explicitly for this value

$$V^R(\alpha) = \left[\int_{U_R}^{\bar{\epsilon}} \epsilon \frac{f(\epsilon)}{1 - F(U_R)} d(\epsilon) - p_R\right] (1 - F(U_R)) + F(U_R) \frac{(1-\alpha)}{\beta\alpha + 1 - \alpha} \left(U - \frac{1 - F(U)}{f(U)}\right) \quad (14)$$

[19]This result is not surprising. Armstrong, Vickers, and Zhou (2009) find that having a prominent seller makes no difference to equilibrium prices with a continuum of sellers. The price difference for the recommended seller in my model comes from the lemon effect.

Corollary 1. *Consumer's value of participation in the market is greater with a recommended seller than in the equilibrium where all sellers are symmetric:* $V^R(\alpha) > V(\alpha)$ *for all* $\alpha > 0$.

The consumer value increases with the recommendation for two reasons: First, the consumers observe the recommended seller with no search cost[20]. Second, while the recommended seller's price is higher than the symmetric equilibrium price, the recommended seller does not capture all of the surplus gained from mitigating the lemon effect, with the remainder going to consumers.

An immediate consequence of Corollary 1 is that consumers ex ante utility is higher at the recommended seller despite its higher price, so visiting the recommended seller first is incentive compatible. If consumers never visit the recommended seller then the search problem is stationary. If consumers do not wish to visit the recommended seller in the first search, then visiting it cannot be more appealing than visiting a random seller in any subsequent search, implying that consumers would never visit the recommended seller. In that case, the value of participation would be exactly equal to participation in the equilibrium with no recommended seller, but by Corollary 1 they can do better by visiting the recommended seller in the first search. Therefore $(1 - F(U_R))$ of the consumers who visit the platform will stop at the recommended seller and pay p_R, the rest will purchase from one of the other sellers and pay p. The platform's profit with a recommended seller is thus

$$\Pi^R = \xi Q(V^R(\alpha)) \left[(1 - F(U_R))p_R + F(U_R)p\right] \tag{15}$$

From Lemma 4 $p_R > p$, so $(1 - F(U_R))p_R + F(U_R)p > p$, and from Corollary 1

[20] The assumption that the prominent seller is observed at no cost can be significantly relaxed and the result will still hold. However it seems natural that search costs would be significantly reduced when seeing the recommended seller and imposing this assumption drastically simplifies the analysis.

$Q(V^R(\alpha)) > Q(V(\alpha))$ for any proportion of low-quality sellers. The platform's profits must be strictly higher with the recommended seller at any fixed α. Given that the platform can still pick any $\alpha \in [0, 1]$ with the recommended seller, its profit must be greater with the ability to recommend. This proves Proposition 3.

Proposition 3. *The platform's profits are strictly higher in the equilibrium of the model where it can recommend a seller than in the model with no recommended seller.*

Also note that if the platform recommends a low quality seller, then share of the platform's profits coming from the recommended seller are scaled down by at least β even if consumers' ex ante beliefs do not change, so the platform is strictly better off recommending a high quality seller.

Proposition 4 summarizes comparisons of the marginal effect of the low-quality sellers between the symmetric monopoly model and the model with the recommended seller. Let α^{R*} denote the equilibrium proportion of low-quality sellers in the model with the recommended seller.

Proposition 4. *For any interior equilibrium:*

1. *If the proportion of low-quality sellers is sufficiently small, then the reservation match value at the recommended seller is smaller than at the other sellers and it is also more responsive to changes in the proportion of low-quality sellers:* $\frac{\partial U_R}{\partial \alpha} < \frac{\partial U}{\partial \alpha}$ *for small* α *and* $U_R < U$ *for small* $\alpha > 0$.

2. *The equilibrium proportion of low-quality sellers is higher with the recommended seller if demand is relatively less sensitive to the curation decision than price: Under Assumption 1* $\alpha^{R*} > \alpha^*$ *if* $Q(\cdot)$ *is sufficiently concave*[21].

[21] Assumption 1 is still sufficient for single-peakedness of the platform profit function.

Part 1 implies that if the equilibrium proportion of low quality sellers is small without a recommended seller, then it will be weakly higher with a recommended seller. The most interesting case being the parameter ranges where the platform would fully screen with no recommended seller but admits a small proportion of low quality sellers if it can make a recommendation.

Corollary 2. *The platform will admit more recommended sellers when it can recommend a seller than when it cannot if the search cost is sufficiently high or the low-quality sellers are sufficiently adept at fooling consumers:* $\alpha^{R*} \geq \alpha^*$ *if* β *and/or* c *sufficiently large.*

α^{R*} can be smaller than α^* when c or β are small because of the higher average price being paid by consumers. Price is more responsive to α with the recommended seller, but because the average price paid by consumers is higher the revenue lost from consumers moving to the outside option is also greater. Part 2 of Proposition 4 says that for $Q(\cdot)$ sufficiently concave the increased ability to raise prices has more effect than the revenue loss from consumers leaving the platform.

Example

Continuing the example from the symmetric equilibrium, we can use Lemma 4 to find

$$U_R - p_R = \frac{1-\alpha}{\beta\alpha + 1 - \alpha}\left(1 - \sqrt{\frac{c}{1-\alpha}}\right) - p$$

Using p from the symmetric equilibrium example and the fact that $p_R = \frac{1-F(U_R)}{f(U_R)} = 1 - U_R$ we can solve this equation to find

$$U_R = \frac{1}{2}\frac{1-\alpha}{\beta\alpha + 1 - \alpha}\left(\frac{\beta\alpha + 2(1-\alpha)}{1-\alpha} - 2\sqrt{\frac{c}{1-\alpha}}\right)$$

$U_R < U$ as $\frac{1-\alpha}{\beta\alpha+1-\alpha}\frac{\beta\alpha+2(1-\alpha)}{1-\alpha} < 2$. The lemon coefficient reduces the appeal of the generic sellers, and so the reservation match value above which consumers purchase from the recommended seller is lower than for the other sellers. Plugging these solutions into Equations (14) and (15), I again compute the derivative of platform profits at $\alpha = 0$. The results are displayed in Figure 3, and they show that admitting low-quality sellers is much more profitable for the platform in the model with the recommended seller.

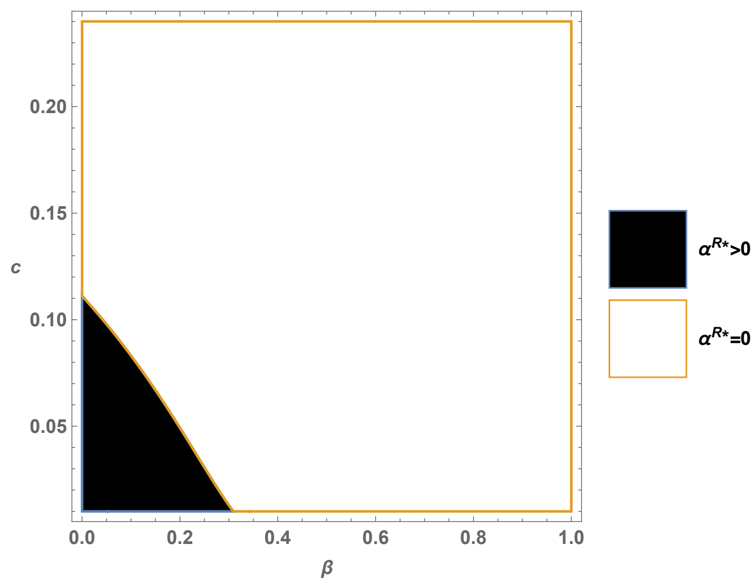

Figure 3: A parameter space showing when the platform will admit a positive proprortion of low-quality sellers when it has the ability to recommend a high-quality seller in this numerical example. When parameters lie in the shaded region, the platform sets $\alpha > 0$ in equilibrium, but it admits no low-quality sellers in the white region.

Comparing these results to Figure 2, the set of parameters where the platform admits low-quality sellers in equilibrium is roughly doubled with a recommended seller. For example: at $c = 0.07$ and $\beta = 0.1$ the platform would not admit any low-quality sellers without a recommended seller, but the marginal profit from the initial low-quality sellers is comfortably positive when a recommendation is possible.

5 Competing Platforms: An illustrative example

The results for competing platforms strongly echo those of Hagiu and Jullien (2014), so to simplify exposition I continue the uniform distrubition example from previous sections in a duopoly environment with two platforms at either end of a unit length Hotelling line. Because I am interested in competition I focus on situations where a monopoly platform would have equilibrium demand greater than $\frac{1}{2}$, which ensures full market coverage in the duopoly model. Further, as the purpose of this section is just to illustrate the principles at work in the duopoly model I ignore potential boundary concerns such as the case where both platforms have demand exactly equal to $\frac{1}{2}$ but they behave as if they are not competing[22]. See Shen, Yang, and Ye (2016) for an example of a model where this latter situation can occur.

The two platforms (denoted 1 and 2) compete in a duopoly setting with single homing consumers and multi-homing sellers. This means that the platforms compete for consumers but not sellers. For this section I focus on symmetric equilibrium where all sellers on platform i set the same price given α_i, and where $\alpha_1 = \alpha_2$. With the additional platform, the timing of the model changes slightly:

1. The platforms set α_1 and α_2 simultaneously.
2. Consumers commit to participating in one of the platforms' markets.
3. Sellers set prices simultaneously on each platform.
4. Consumers participating a platform's market search among sellers and make purchasing decisions.

Because consumers single home, once they have chosen a platform they will act exactly as in Section 3, substituting α_i for α. Similarly, because seller pricing behavior

[22] A more general solution for the duopoly model is available in Appendix B

reacts to the consumers' stopping rule and multi-homing implies that the profit maximization problems on each platform are effectively independent. Prices will therefore be identical (for a given α) to those in Section 3. The search environment is once again stationary so the commitment assumption is not binding, consumers who find platform i more appealing for their first search will always find it to be the most appealing option. I assume a transport cost t, so a consumer at point x on the Hotelling line receives expected surplus

$$V(\alpha_1) - tx \qquad \text{from platform 1}$$
$$V(\alpha_2) - t(1-x) \qquad \text{from platform 2}$$

The consumer who is indifferent between platforms has the position on the line which sets these two values exactly equal, and so by standard Hotelling derivations the mass of consumers participating in platform i's market (denoted $Q_i(V(\alpha_i), V(\alpha_{-i}))$) is given by

$$Q_i(V(\alpha_i), V(\alpha_{-i})) = \frac{1}{2} + \frac{V(\alpha_i) - V(\alpha_{-i})}{2t}$$

Let p_i be the market price on platform i, then platform profit Π_i is given by

$$\Pi_i = \xi p_i Q_i \left(V(\alpha_i), V(\alpha_{-i}) \right)$$

Taking the first order condition with regard to α_i

$$\frac{1}{2p_i} = \frac{\frac{\partial Q_i(V(\alpha_i), V(\alpha_{-i}))}{\partial V(\alpha_i)}}{Q_i(V(\alpha_i), V(\alpha_{-i}))} + \beta B_i(\alpha_i, \alpha_{-i})$$

Where

$$B_i(\alpha_i, \alpha_{-i}) = \sqrt{\frac{1-\alpha}{c}} \frac{Q_i(V(\alpha_i), V(\alpha_{-i}))\sqrt{\frac{c}{1-\alpha}} + p_i \frac{\partial Q_i(V(\alpha_i), V(\alpha_{-i}))}{\partial V(\alpha_i)} \left(1 - 2\sqrt{\frac{c}{1-\alpha}}\right)}{2Q_i(V(\alpha_i), V(\alpha_{-i}))\sqrt{\frac{c}{1-\alpha}}}$$

And using the definition of the participation function

$$\frac{1}{2p_i} = \frac{\frac{1}{2t}}{\frac{1}{2} + \frac{V(\alpha_i) - V(\alpha_{-i})}{2t}} + \beta B_i(\alpha_i, \alpha_{-i})$$

In equilibrium $\alpha_i = \alpha_{-i}$, so we can further simplify to get

$$\frac{1}{2p_i} = \frac{1}{t} + \beta B_i(\alpha_i, \alpha_{-i}) \tag{16}$$

Compare this to the first order condition for a monopoly firm

$$\frac{1}{2p} = \frac{1}{V(\alpha)} + \beta B(\alpha) \tag{17}$$

Now suppose that the duopoly firms are using the monopoly curation decision (i.e. $\alpha_i = \alpha_{-i} = \alpha^*$). The left hand side of Equations (16) and (17) will then be equal. The symmetric equilibrium curation decision in the duopoly will have more low-quality sellers than the monopoly if and only if the right hand side of Equation (16) is less than the right hand side of Equation (17). While not completely obvious, it is fairly straightforward to find that a necessary and sufficient condition for this statement to hold is for consumer participation in a platform's market to be more elastic with regard to $V(\alpha)$ in the monopoly than the duopoly:

$$\frac{1}{V(\alpha^*)} > \frac{1}{t}$$
$$\implies \frac{V(\alpha^*)}{t} < 1$$

But this inequality is equivalent to the statement that the monopolist's equilibrium demand is interior. In other words, if $\frac{V(\alpha^*)}{t} \in (\frac{1}{2}, 1)$ then the competing platforms will always allow more low-quality sellers than the monopolist. This mirrors the frequently seen result that increasing the number of firms can lower quality investment under price competition by reducing the marginal demand which additional quality investment brings to a firm. For any *fixed* α, the consumers who choose platform 1 rather than the outside option in the monopoly model, but who choose platform 2 in the duopoly prefer the duopoly because they are able to choose a platform which is more closely aligned to their tastes. For this reason, it is harder for platform 1 to attract consumers away from platform 2 than would be for platform 1 to get consumers to switch away from the outside option if it were a monopoly. Because stealing business from the other platform is less profitable than expanding the market to consumers who are not shopping, the duopoly platforms are not willing to lower market prices to attract new consumers as much as the monopoly platform would, so the equilibrium proportion of low-quality sellers is higher in the duopoly model.

However, this is not quite equivalent to saying that increased competition increases α. While t doesn't impact the curation decision in the monopoly case, if the duopoly platforms are viewed as closer substitutes (t decreases) then a greater proportion of consumers will switch to platform $-i$ as α_i increases, meaning that the profitability of business stealing increases and the proportion of low-quality sellers in the duopoly model decreases as the degree of competition between the two platforms intensifies. Thus, while entry does not discipline the platforms' search obfuscation behavior, increased substitutability between platforms does.

The generalized condition for the proportion of low-quality sellers to increase in the duopoly model is

$$\frac{\frac{\partial Q_i(V(\alpha^*),V(\alpha^*))}{\partial V(\alpha_i)}}{Q_i(V(\alpha^*),V(\alpha^*))} < \frac{Q'(V(\alpha^*))}{Q(V(\alpha^*))}$$

That is, the equilibrium proportion of low-quality sellers increases with the addition of a second platform if and only if demand is less elastic in the duopoly platform case. This closely mirrors Hagiu and Jullien's (2014) result that competing platforms will engage in more search diversion if competition between them is relatively weak, but not so weak that they are local monopolies.

6 Discussion

In this section I discuss several topics which do not merit a full extension of the model but which nevertheless have potentially important implications for my results.

Screening Costs

In the baseline model I assume that the platform can freely set the proportion of high- and low-quality sellers in order to more clearly demonstrate the search obfuscation mechanism. This is obviously not the case for real world platforms and so the reader might reasonably question how to interpret the results of this model. To answer that question, consider a situation where there is a base rate of low-quality sellers and that the platform must pay a cost to find and expel low quality sellers.

In Figure 4 I show an example of a situation that could then occur. The base rate of low quality sellers if the platform doesn't screen is approximately $\frac{5}{6}$[23]. If the platform had no screening costs then it would set $\alpha^* = \frac{1}{2}$, but because screening out

[23]These numbers are not based on any data, but are instead chosen to optimize readability of the figure.

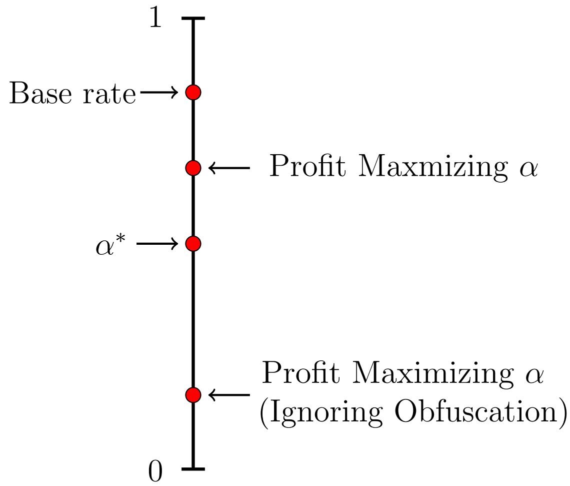

Figure 4: The proportion of low-quality sellers with no screening, at the profit maximizing level, if there were 0 screening cost, and if the platform treated prices as independent of the proportion of low quality sellers.

the low-quality sellers is costly, its profit maximizing proportion is $\frac{2}{3}$. However, if the platform ignored the search obfuscating effect of the low-quality sellers (i.e. if it assumed that getting rid of these sellers would have no effect on price) then it would instead screen to a proportion of $\frac{1}{6}$. An outside observer who fails to take obfuscation into account might then wonder why the platform sets $\alpha = \frac{2}{3}$ instead of $\frac{1}{6}$. If instead α^* were greater than the base rate then we might not see any screening at all, as screening out high-quality sellers might create a significant backlash in public sentiment or even legal action against the platform.

This is a relatively simple conception of screening costs, but it illustrates the point that the phenomena I demonstrate would show up in real world behavior as inaction on the part of the platform. An empirical use for this model could be to explain gaps between predicted optimal levels of screening and actual screening in highly competitive platform markets.

Review Scoring Services

One of the more surprising results of this model is that review scoring websites such as reviewmeta.com may be encouraging search obfuscation. The reason for this is that they counteract the lemon effect by reducing low-quality sellers' ability to imitate the high quality sellers (reducing β). However, they do not eliminate the search obfuscation effect at all, and they may even exacerbate it somewhat because of the extra effort consumers must go through to evaluate products. The cost of admitting low-quality sellers is lower than it would be without these services and the platform has less incentive to reduce α. It is possible that the welfare improvement from reducing the lemon effect might overcome the negative effects of any increase in the proportion of low quality sellers. In that case the net effect of these review scoring sites may still be an increase in consumer welfare despite the undesirable increase in search obfuscation.

Cost of Returns

This model ignores the fact that unsatisfied consumers might wish to return the products they have purchased. Assuming that returns are costly to the platform (as in the case with Amazon), then the main effect of adding returns would be to magnify the negative impact of β on platform profits. Consumers who purchase a product only to later return it not only do not contribute to profits, but increase costs. However, if β is sufficiently low then the volume of returns should be low enough that the increased profits from search obfuscation would still overcome the combined negative effects of the lemon effect and return costs. However, the equilibrium proportion of low-quality sellers would be lower than if returns were not a factor[24].

[24] Customers often complain that Amazon makes return products unnecessarily difficult, either by making the process obtuse or requiring them to visit one of Amazon's physical locations. It may be

The impact of returns on platform profit may be mitigated or even positive if consumers engage in the search process again after returning a low-quality product. While even a return policy that fully reimbursed the purchase price could not eliminate the negative impact of being scammed on consumers as they will still have expended search costs on purchasing a dud product, it would mitigate their welfare loss. Consequently, consumers value of participation $V(\alpha)$ and therefore $Q(V(\alpha))$ would increase at any α with the introduction of a return policy. Furthermore, the expected value of searching for a high-quality product is still positive for the consumers who have returned duds, so they will once again participate in the platform's market and purchase from a seller on the platform, partially counteracting the return costs. See Hinnosaar and Kawai (2018) for a more thorough discussion of sellers' optimal behavior when returns are a factor.

7 Conclusion

This project explores the effects of low-quality sellers in platform markets and why platforms might be reluctant to curate the selection of sellers who list in their markets. Platforms choose the proportion of high and low quality sellers on their markets with no screening cost. Sellers and consumers then participate in a market search game similar to Anderson and Renault (1999) or Wolinsky (1986). The impact of low-quality sellers can be broken down into a search obfuscation effect and a lemon effect. The search-obfuscation effect reduces consumers' willingness to search and so softens competition between sellers on the platform. The lemon effect reflects consumers' belief that the product they are purchasing might be a low-quality product successfully imitating a high-quality product, which reduces their willingness to pay. Both effects

the case that this is an intentional strategy intended to mitigate the costs of accepting returns.

reduce consumers' participation in the market, but only the former is beneficial to the platform, so if the low-quality sellers are too adept at fooling consumers then the platform would prefer to admit no low-quality consumers. However, if the obfuscation effect dominates the lemon effect and the market is sufficiently competitive then the platform-profit maximizing proportion of low-quality consumers is positive *despite the lack of screening cost*.

If the platform can recommend a high-quality seller, then that seller will not be subject to the lemon effect and so charges a higher price. Recommending a high-quality seller is more profitable for the platform than simply obfuscating search as in the baseline model. The ability to recommend may cause the platform to admit low-quality sellers when it would not if it could not recommend a seller. This recommended seller does not capture all of the benefits of the recommendation ,so consumer welfare is higher with a recommended seller at any proportion of low-quality sellers. However, the net effect of recommendation may be negative for consumers if the ability to recommend causes the platform to admit more low-quality sellers.

Platform entry has an ambiguous effect on the level of curation in equilibrium, but increased intensity of competition between platforms who have already entered will always lead to a decrease in the proportion of low-quality sellers. Therefore an interested policy-maker seeking to induce more screening in a market where there are already multiple platforms should look for ways to increase the substitutability of these platforms.

References

Anderson, Simon P, Andre De Palma, and Yurii Nesterov. 1995. Oligopolistic competition and the optimal provision of products. *Econometrica: Journal of the Econometric Society:* 1281–1301.

Anderson, Simon P., and Regis Renault. 1999. Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model. *RAND Journal of Economics* 30, no. 4 (): 719–735.

Armstrong, Mark, John Vickers, and Jidong Zhou. 2009. Prominence and consumer search. *The RAND Journal of Economics* 40 (2): 209–233.

Armstrong, Mark, and Jidong Zhou. 2015. Search deterrence. *The Review of Economic Studies* 83 (1): 26–57.

Bagnoli, Mark, and Ted Bergstrom. 2005. Log-concave probability and its applications. *Economic theory* 26 (2): 445–469.

Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat. 2012. Search, design, and market structure. *American Economic Review* 102 (2): 1140–60.

Belleflamme, Paul, and Martin Peitz. 2019. Managing competition on a two-sided platform. *Journal of Economics & Management Strategy* 28 (1): 5–22.

Blake, Thomas, Chris Nosko, and Steven Tadelis. 2016. Returns to consumer search: evidence from ebay. In *Proceedings of the 2016 acm conference on economics and computation,* 531–545. ACM.

Burdett, Kenneth, and Tara Vishwanath. 1988. Declining reservation wages and learning. *The Review of Economic Studies* 55 (4): 655–665.

Casner, Ben. 2018. Learning while shopping: an experimental investigation into the effects of learning on recall in consumer search. Working Paper.

Che, Yeon-Koo, and Johannes Hörner. 2018. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics* 133 (2): 871–925.

Chen, Yuxin, and Song Yao. 2016. Sequential search with refinement: model and application with click-stream data. *Management Science* 63 (12): 4345–4365.

de Roos, Nicolas. 2018. Collusion with limited product comparability. *The RAND Journal of Economics* 49 (3): 481–503.

Diamond, Peter. 1971. A model of price adjustment. *Journal of Economic Theory* 3 (2): 156–168.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2018. Consumer price search and platform design in internet commerce. *American Economic Review* 108 (7): 1820–59.

Ellison, Glenn, and Sara Fisher Ellison. 2009. Search, obfuscation, and price elasticities on the internet. *Econometrica* 77 (2): 427–452.

Ellison, Glenn, and Alexander Wolitzky. 2012. A search cost model of obfuscation. *The RAND Journal of Economics* 43 (3): 417–441.

Fudenberg, Drew, and Jean Tirole. 1991. *Game theory.* Cambridge, MA: MIT Press.

Gabaix, Xavier, and David Laibson. 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics* 121 (2): 505–540.

Hagiu, Andrei, and Hanna Hałaburda. 2014. Information and two-sided platform profits. *International Journal of Industrial Organization* 34:25–35.

Hagiu, Andrei, and Bruno Jullien. 2011. Why do intermediaries divert search? *The RAND Journal of Economics* 42 (2): 337–362.

———. 2014. Search diversion and platform competition. *International Journal of Industrial Organization* 33:48–60.

Hagiu, Andrei, and Julian Wright. 2015. Multi-sided platforms. *International Journal of Industrial Organization* 43:162–174.

Hinnosaar, Toomas, and Keiichi Kawai. 2018. Robust pricing with refunds. Working Paper.

Lubensky, Dmitry. 2017. A model of recommended retail prices. *The RAND Journal of Economics* 48 (2): 358–386.

Nosko, Chris, and Steven Tadelis. 2015. *The limits of reputation in platform markets: an empirical analysis and field experiment.* Technical report. Working Paper. National Bureau of Economic Research.

Shen, Jian, Huanxing Yang, and Lixin Ye. 2016. Competitive nonlinear pricing and contract variety. *The Journal of Industrial Economics* 64 (1): 64–108.

Wang, Chengsi, and Julian Wright. 2018. Search platforms: showrooming and price parity clauses. Working Paper.

Weyl, E Glen. 2010. A price theory of multi-sided platforms. *American Economic Review* 100 (4): 1642–72.

White, Alexander. 2013. Search engines: left side quality versus right side profits. *International Journal of Industrial Organization* 31 (6): 690–701.

Wolinsky, Asher. 1986. True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics* 101 (3): 493–511.

Yang, Huanxing. 2013. Targeted search and the long tail effect. *RAND Journal of Economics* 44 (4): 733–756.

Yea, Sangjun. 2018. Quality disclosure in online marketplaces. Working Paper.

A Omitted Proofs

Proof of Lemma 1

Because U is determined by Equation (3), the derivatives of both sides with regard to α must be equal.

$$\begin{aligned} 0 &= -\int_U^{\bar{\epsilon}} (\epsilon' - U) f(\epsilon')d\epsilon' - (1-\alpha)\frac{\partial U}{\partial \alpha}\int_U^{\bar{\epsilon}} f(\epsilon')d\epsilon' \\ \Longrightarrow \frac{\partial U}{\partial \alpha} &= -\frac{\int_U^{\bar{\epsilon}} (\epsilon' - U) f(\epsilon')d\epsilon'}{(1-\alpha)(1-F(U))} \end{aligned} \tag{18}$$

since $\int_U^{\bar{\epsilon}} f(\epsilon')d\epsilon' = (1 - F(U))$. All of the elements of the fraction on the right hand side are positive, so the negative of the fraction must be negative and $\frac{\partial U}{\partial \alpha} < 0$. ■

Proof of Lemma 2

Directly taking the derivative of Equation (4)

$$\begin{aligned}\frac{\partial p}{\partial \alpha} &= \frac{-(\beta\alpha+1-\alpha)+(1-\beta)(1-\alpha)}{(\beta\alpha+1-\alpha)^2}\frac{1-F(U)}{f(U)} + \frac{\partial U}{\partial \alpha}\frac{(1-\alpha)}{\beta\alpha+1-\alpha}\frac{-f(U)^2-f'(U)(1-F(U))}{f(U)^2} \\ &= \frac{-\beta}{(\beta\alpha+1-\alpha)^2}\frac{1-F(U)}{f(U)} + \frac{\partial U}{\partial \alpha}\frac{(1-\alpha)}{\beta\alpha+1-\alpha}\frac{-f(U)^2-f'(U)(1-F(U))}{f(U)^2}\end{aligned}$$

In a symmetric equilibrium, U does not depend on β, so as $\beta \to 0$ the first term in the parentheses on the right hand side approaches 0 and the second term approaches $\frac{\partial U}{\partial \alpha}\frac{-f(U)^2-f'(U)(1-F(U))}{f(U)^2}$. From Lemma 1 $\frac{\partial U}{\partial \alpha} < 0$, and Anderson, De Palma, and Nesterov (1995) demonstrate that $\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2} > 0$ for a log concave distribution, so $\frac{\partial U}{\partial \alpha}\frac{-f(U)^2-f'(U)(1-F(U))}{f(U)^2} > 0$. Therefore $\lim_{\beta\to 0}\frac{\partial p}{\partial \alpha} > 0$ and by inspection $\frac{\partial p}{\partial \alpha}$ is continuous in β so it must be the case that $\frac{\partial p}{\partial \alpha} > 0$ for β sufficiently small. ■

Proof of Proposition 1

Taking the derivative of Π with regard to α

$$\begin{aligned}\frac{\partial \Pi}{\partial \alpha} =& \frac{\partial p}{\partial \alpha}Q(V(\alpha)) + pQ'(V(\alpha))V'(\alpha) \\ =& \left(\frac{-\beta}{(\beta\alpha+1-\alpha)^2}\frac{1-F(U)}{f(U)} + \frac{\partial U}{\partial \alpha}\frac{(1-\alpha)}{\beta\alpha+1-\alpha}\frac{-f(U)^2-f'(U)(1-F(U))}{f(U)^2}\right)Q(V(\alpha)) \\ &+ pQ'(V(\alpha))\left(\frac{-\beta}{(\beta\alpha+1-\alpha)^2}\left(U-\frac{1-F(U)}{f(U)}\right) + \frac{(1-\alpha)}{\beta\alpha+1-\alpha}\frac{\partial U}{\partial \alpha}\left(1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)\right) \\ =& \Bigg[\frac{-\beta}{(\beta\alpha+1-\alpha)^2}\left(Q(V(\alpha))\frac{1-F(U)}{f(U)} + pQ'(V(\alpha))\left(U-\frac{1-F(U)}{f(U)}\right)\right) \\ &+ \frac{(1-\alpha)}{\beta\alpha+1-\alpha}\frac{\partial U}{\partial \alpha}\left(pQ'(V(\alpha))\left[1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right] - Q(V(\alpha))\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)\Bigg]\end{aligned} \tag{19}$$

The first term in the brackets represents the change in profits due to the lemon effect. The lemon effect pushes prices down and drives consumers away from platform, so this term is always negative. The second term represents the change in profits resulting from search obfuscation; this effect can be either positive or negative depending on whether the increased prices make up for reduced participation. When $\alpha = 0$ the second term does not depend on β at all. If this second term is positive then for β sufficiently small the entire derivative must be positive and platform profit is increasing in α at $\alpha = 0$, this gives condition 1.

The second term is positive at $\alpha = 0$ if

$$pQ'(V(0))\left[1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right] < Q(V(0))\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}$$

Rearranging terms

$$\frac{Q'(V(0))}{Q(V(0))}\frac{\left(1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)}{\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}} < \frac{1}{p} \tag{20}$$

As c decreases, both U and $V(\alpha)$ increase and p decreases. As U approaches $\bar{\epsilon}$, p gets arbitrarily close to 0. If the left hand side has a finite limit as as U approaches $\bar{\epsilon}$ then the inequality must hold for c sufficiently small. This condition is true if $f'(U)$ is sufficiently small, which gives condition 2.

■

Proof of Lemma 3

Sufficiency of the first order condition follows immediately if Π is single peaked in α. Single-peakedness of Π in α is equivalent to

$$\frac{\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}}{p\left(1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)} - \frac{Q'(V(\alpha))}{Q(V(\alpha))} - \beta B(\alpha)$$

decreasing in α. Since $Q(\cdot)$ is concave and $V(\alpha)$ is decreasing in α, $\frac{Q'(V(\alpha))}{Q(V(\alpha))}$ must be increasing in α. β must be small enough under for Lemma 2 to apply, otherwise the platform would not admit any low-quality sellers since Π is strictly decreasing in α if p is not increasing. Therefore $\frac{1}{p}$ must decrease in α, so the only possible sources of multiple maxima are $\frac{\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}}{\left(1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)}$ and $\beta B(\alpha)$. If the variance of $f(\cdot)$ is sufficiently large (so that $f(\cdot)$ is small) then the effect of $\frac{1}{1-\alpha}$ in the first term of $B(\alpha)$ dominates and $B(\alpha)$ is increasing in α. Finally,

$$\frac{\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}}{\left(1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)} = \frac{1+\frac{f'(U)(1-F(U))}{f(U)^2}}{2+\frac{f'(U)(1-F(U))}{f(U)^2}}$$

is the ratio of the change in price as U decreases to the change in $V(\alpha)$ as U decreases. This ratio can increase or decrease in U, but if $f'(U)$ is sufficiently close to 0 (e.g. with a uniform distribution or high variance normal distribution) then the effects of $\frac{1}{p}$ and $\frac{Q'(V(\alpha))}{Q(V(\alpha))}$ dominate the other terms and profit is single-peaked in α.

■

Proof of Proposition 2

Proof of part 1: First note that Π is increasing in α if the left hand side of Equation (8) minus the right hand side is positive.

As c increases, U decreases and p increases, meaning that $V(\alpha)$ decreases. Since the direct effect of α on p decreases prices via the lemon coefficient, Assumption 1 implies that

$$\frac{\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}}{p\left(1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)} - \frac{Q'(V(\alpha))}{Q(V(\alpha))} - \beta B(\alpha)$$

is increasing in U, so if c decreases then marginal profits in α decrease for all values of α. Concavity of the objective function then further implies that the α which solves

Equation (8) must decrease, so α^* decreases.

Proof of part 2 If $Q(V(\alpha)) = \frac{V(\alpha)}{t}\ t > 0$, then from Equation (19)

$$\begin{aligned}\frac{\partial \Pi}{\partial \alpha} =& \frac{1}{t}\Bigg[\frac{-2\beta(1-\alpha)}{(\beta\alpha+1-\alpha)^3}\left(\frac{1-F(U)}{f(U)}\left(U-\frac{1-F(U)}{f(U)}\right)\right) \\ &+\left(\frac{(1-\alpha)}{\beta\alpha+1-\alpha}\right)^2\frac{\partial U}{\partial \alpha}\left(\frac{1-F(U)}{f(U)}\left[1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right]-\left(U-\frac{1-F(U)}{f(U)}\right)\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)\Bigg] \\ =& \frac{(1-\alpha)}{t(\beta\alpha+1-\alpha)^2}\Bigg[\frac{-2\beta}{\beta\alpha+1-\alpha}\left(\frac{1-F(U)}{f(U)}\left(U-\frac{1-F(U)}{f(U)}\right)\right) \\ &+(1-\alpha)\frac{\partial U}{\partial \alpha}\left(\frac{1-F(U)}{f(U)}\left[1+\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right]-\left(U-\frac{1-F(U)}{f(U)}\right)\frac{f(U)^2+f'(U)(1-F(U))}{f(U)^2}\right)\Bigg]\end{aligned}$$

The terms outside of the brackets do not matter for determining the effect of β on α^*. The first term in brackets is negative and decreasing in β, the second term in the brackets is positive in any interior equilibrium and does not vary with β. From Assumption 1 the sum inside the brackets must be decreasing in α, so if the first term becomes more negative then α^* must decrease.

■

Proof of Lemma 4

Consumers only move on from the recommended seller if ϵ_{iR} is low, which means that any consumer who visits a different seller will never go back to the recommended seller. Deriving the stopping rule for these consumers then follows precisely the same steps as in Section 3. Because consumers are using the Section 3 stopping rule, the generic sellers will charge the Section 3 price.

For the rest of the lemma, begin by comparing Equation (3) and Equation (11), it must be the case that

$$\frac{\beta\alpha+1-\alpha}{1-\alpha}\left(U_R+p-p_R\right)=U$$
$$\implies U_R=\frac{1-\alpha}{\beta\alpha+1-\alpha}U+p_R-p$$

This equation can be rearranged to give part 2. Suppose $p_R \leq p$, then $p_R - p \leq 0$ and $U_R < U$ as $\frac{(1-\alpha)}{\beta\alpha+1-\alpha} < 1$ when $\alpha > 0$ and $\beta > 0$. But $f(\cdot)$ is log concave so $\frac{1-F(x)}{f(x)}$ is decreasing in x and

$$p_R = \frac{1-F(U_R)}{f(U_R)} > \frac{1-F\left(U\right)}{f\left(U\right)} > \frac{\left(1-\alpha\right)}{\beta\alpha+1-\alpha}\frac{1-F\left(U\right)}{f\left(U\right)} = p$$

Which contradicts the assumption that $p_R \leq p$ and proves part 3. The proof of part 4 uses nearly identical logic to show that an inequality in either direction produces a contradiction, noting that $\frac{(1-\alpha)}{\beta\alpha+1-\alpha} = 1$ when $\alpha = 0$.

■

Proof of Corollary 1

With probability $1-F(U_R)$ consumers purchase from the recommended seller and receive expected utility $\int_{U_R}^{\bar{\epsilon}} \epsilon \frac{f(\epsilon)}{1-F(U_R)} d(\epsilon) - p_R$. With probability $F(U_R)$, consumers move on from the recommended seller, in which case their search problem is identical to the equilibrium without a recommended seller and their expected payoff is $\frac{(1-\alpha)}{\beta\alpha+1-\alpha}\left(U - \frac{1-F(U)}{f(U)}\right)$, so

$$V(\alpha) = \left[\int_{U_R}^{\bar{\epsilon}} \epsilon \frac{f(\epsilon)}{1-F(U_R)} d(\epsilon) - p_R\right] (1-F(U_R)) + F(U_R)\frac{(1-\alpha)}{\beta\alpha+1-\alpha}\left(U - \frac{1-F(U)}{f(U)}\right) \tag{21}$$

Focusing on the first term

$$\begin{aligned}\int_{U_R}^{\bar{\epsilon}} \epsilon \frac{f(\epsilon)}{1-F(U_R)} d(\epsilon) - p_R =& \int_{U_R}^{\bar{\epsilon}} (\epsilon - p_R) \frac{f(\epsilon)}{1-F(U_R)} d(\epsilon) \\ >& \int_{U_R}^{\bar{\epsilon}} (U_R - p_R) \frac{f(\epsilon)}{1-F(U_R)} d(\epsilon) \\ =& U_R - p_R \\ =& \frac{(1-\alpha)}{\beta\alpha + 1 - \alpha} \left(U - \frac{1-F(U)}{f(U)} \right)\end{aligned}$$

Where the last equality comes from Lemma 4. Plugging this inequality into the value function

$$\begin{aligned}V(\alpha) >& ((1-F(U_R)) + F(U_R)) \frac{(1-\alpha)}{\beta\alpha + 1 - \alpha} \left(U - \frac{1-F(U)}{f(U)} \right) \\ =& \frac{(1-\alpha)}{\beta\alpha + 1 - \alpha} \left(U - \frac{1-F(U)}{f(U)} \right)\end{aligned} \tag{22}$$

Since U is identical to the reservation utility from the equilibrium without the recommended seller, Equation (22) implies that consumers must be strictly better off with the recommended seller.■

Proof of Proposition 4

Part 1: Since part 2 of Lemma 4 holds for all α, the derivatives of both sides of the equation must be equal

$$\frac{\partial U_R}{\partial \alpha} \left(1 + \frac{f(U_R)^2 + f'(U_R)(1-F(U_R))}{f(U_R)^2} \right) = \frac{1-\alpha}{\beta\alpha + 1 - \alpha} \frac{\partial U}{\partial \alpha} \left(1 + \frac{f(U)^2 + f'(U)(1-F(U))}{f(U)^2} \right) - \frac{\beta}{(\beta\alpha + 1 - \alpha)^2} \left(U - \frac{1-F(U)}{f(U)} \right)$$

Rearranging

$$\frac{\partial U_R}{\partial \alpha} = \frac{1-\alpha}{\beta\alpha + 1 - \alpha} \frac{\partial U}{\partial \alpha} \frac{\left(1 + \frac{f(U)^2 + f'(U)(1-F(U))}{f(U)^2} \right)}{\left(1 + \frac{f(U_R)^2 + f'(U_R)(1-F(U_R))}{f(U_R)^2} \right)} - \frac{\beta}{(\beta\alpha + 1 - \alpha)^2} \frac{\left(U - \frac{1-F(U)}{f(U)} \right)}{\left(1 + \frac{f(U_R)^2 + f'(U_R)(1-F(U_R))}{f(U_R)^2} \right)} \tag{23}$$

Since $U_R = U$ and $\frac{1-\alpha}{\beta\alpha+1-\alpha} = 1$ for $\alpha = 0$, this then implies that $\frac{\partial U_R}{\partial \alpha} < \frac{\partial U}{\partial \alpha}$ for α sufficiently close to 0, which in turn implies that $U_R < U$ for α positive but near 0.

Part 2: Taking the derivative of platform profit

$$\begin{aligned}\frac{\partial \Pi_R}{\partial \alpha} =& Q'(V(\alpha))V'(\alpha)\left[(1-F(U_R)p_R + F(U_R)p\right] \\ &+ Q(V(\alpha))\left[(1-F(U_R))\frac{\partial p_R}{\partial \alpha} + F(U_R)\frac{\partial p}{\partial \alpha} - \frac{\partial U_R}{\partial \alpha} f(U_R)(p_R - p)\right]\end{aligned} \tag{24}$$

Evaluating the derivative of the recommended seller's profit

$$\frac{\partial p_R}{\partial \alpha} = -\frac{\partial U_R}{\partial \alpha}\frac{f(U_R)^2 + f'(U_R)(1-F(U_R))}{f(U_R)^2}$$

From Equation (23) this derivative must be positive. From part 1 of this proposition for small α it is larger than the change in p with α, but the comparison is ambiguous for large α. Additionally, because $p_R > p$ for $\alpha > 0$, the marginal costs of increasing α (the first term on the right hand side of Equation (24)) of can potentially be larger when the platform recommends a seller as well. However, from Corollary 1, $Q'(V(\alpha))$ must be closer to 0 in the equilibrium with the recommended seller because $V(\alpha)$ is larger. For $Q(\cdot)$ sufficiently concave, this effect dominates and $\frac{\partial \Pi_R}{\partial \alpha} > \frac{\partial \Pi}{\partial \alpha}$ for all α, so from Assumption 1, we must have $\alpha^{R*} > \alpha^*$.■

A.1 Proof of Corollary 2

From Proposition 1 and Proposition 2, for β or c sufficiently large, α^* is close to 0. But then from Proposition 4 and Corollary 1 a change in α must result in a larger increase in seller profits must and the smaller decrease in participation when the platform can recommend a seller. It follows immediately that any interior $\alpha^{R*} > \alpha^*$. The equality in the corollary refers to the possibility that c or β may be large enough that the platform does not admit low-quality sellers in either model.

■

B The generalized duopoly model

Suppose that instead of a single monopoly platform, two platforms (denoted 1 and 2) compete in a duopoly setting with single homing consumers and multi-homing sellers. I use the subscript 1 to denote the variables relevant to platform 1 and 2 to those on platform 2. For this section I focus on symmetric equilibrium where all sellers on platform i set the same price given α_i, and where $\alpha_1 = \alpha_2$. With the additional platform, the timing of the model changes slightly:

1. The platforms set α_1 and α_2 simultaneously.
2. Consumers commit to participating in one of the platforms' markets or remaining with the outside option.
3. Sellers set prices simultaneously on each platform.
4. Consumers participating a platform's market search among sellers and make purchasing decisions.

Prices are not observable to consumers so multi-homing sellers set the prices independently on each platform by best responding to consumer search behavior. Similarly, as consumers commit to a platform before searching, the proportion of low-quality sellers on the other platform does not influence their search behavior. For

$i = 1, 2$ the derivations from the monopoly model apply:

$$p_i = \frac{(1-\alpha_i)}{\beta\alpha_i + 1 - \alpha_i} \frac{1 - F(U_i)}{f(U_i)}$$
$$V(\alpha_i) = \frac{(1-\alpha_i)}{\beta\alpha_i + 1 - \alpha_i} \left(U_i - \frac{1 - F(U_i)}{f(U_i)} \right)$$

Where U_i is determined by Equation (3), substituting α_i for α. The mass of consumers who commit to platform i is given by the continuous and differentiable function $Q_i(V(\alpha_1), V(\alpha_2))$, where $Q_i(V(\alpha_i), V(\alpha_{-i}))$ is increasing and concave in $V(\alpha_i)$ and decreasing in $V(\alpha_{-i})$. Furthermore, I assume that

$$\frac{dQ_i(x,x)}{dx} = \frac{dQ_{-i}(x,x)}{dx} \geq 0 \text{ for } i = 1, 2$$

This assumption implies that neither platform has an inherent competitive advantage. If both platforms provide the same expected value, and this value increases while remaining symmetric, then any change in demand must come from the outside option. Behavior is identical to the monopoly once consumers have committed to a platform, so the major implication of introducing platform competition is that consumer participation in a platform's market will depend on both platforms' curation decisions.

Lemma 5. *In the symmetric duopoly model, for β and c sufficiently small and if* $\max f'(\cdot)$ *sufficiently small then $\alpha_i = \alpha_{-i} = 0$ cannot be an equilibrium.*

Proof. Following the same derivation steps as in the proof of Proposition 1, a platform's profit is increasing if β is sufficiently small and at $\alpha_i = \alpha_{-i} = 0$

$$\frac{\frac{\partial Q_i(V(0),V(0))}{\partial V(\alpha_i)}}{Q_i(V(0),V(0))} \frac{\left(1+\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}\right)}{\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}} < \frac{1}{p_i}$$

By precisely the same logic as in the proof of Proposition 1, if

$$\lim_{U_i \to \bar{\epsilon}} \frac{\frac{\partial Q_i(V(0),V(0))}{\partial V(\alpha_i)}}{Q_i(V(0),V(0))} \frac{\left(1+\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}\right)}{\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}} < \infty$$

then the above inequality must hold for c sufficiently small. ■

The logic of the proof of Lemma 5 is similar to that of Proposition 1, and the intuition is similar as well. If the the platform markets are sufficiently competitive, and the lemon effect not too strong, then each platform will have an incentive to obfuscate search even if the other platform has no low-quality sellers.

Assumption 2. *β and c are sufficiently small so that $\alpha_i = \alpha_{-i} = 0$ is not an equilibrium.*

By nearly the same reasoning as in the monopoly model, platform profit is guaranteed to be concave for $f(\cdot)$ sufficiently small across its range, so Assumption 3 gives sufficiency of the first order condition for platform $i's$ curation decision.

Assumption 3. *The variance of $f(\cdot)$ is sufficiently large so that* $\frac{\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}}{p_i\left(1+\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}\right)} - \frac{\frac{\partial Q_i(V(\alpha_i),V(\alpha_{-i}))}{\partial V(\alpha_i)}}{Q_i(V(\alpha_i),V(\alpha_{-i}))} - \beta B_i(\alpha_i, \alpha_{-i})$ *is decreasing in α_i*

Lemma 6 describes equilibrium curation levels in this environment. Let α_D^* denote the symmetric equilibrium curation decision in the duopoly model, then

Lemma 6. *Under Assumption 2 and Assumption 3, if $Q_i(V(\alpha_i), V(\alpha_{-i}))$ is continuously differentiable then an equilibrium with symmetric $\alpha > 0$ exists and for $i = 1, 2$, α_D^* is determined implicitly by*

$$\frac{\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}}{p_i\left(1+\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}\right)} = \frac{\frac{\partial Q_i(V(\alpha_D^*),V(\alpha_D^*))}{\partial V(\alpha_i)}}{Q_i(V(\alpha_D^*),V(\alpha_D^*))} + \beta B_i(\alpha_D^*, \alpha_D^*) \tag{25}$$

Where

$$B_i(\alpha_D^*, \alpha_D^*) = \frac{(1-F(U_i))}{(1-\alpha)\int_{U_i}^{\bar{\epsilon}}(\epsilon - U_i)f(\epsilon)d\epsilon} \frac{Q_i(V(\alpha_D^*),V(\alpha_D^*))\frac{1-F(U_i)}{f(U_i)} + p_i \frac{\partial Q_i(V(\alpha_D^*),V(\alpha_D^*))}{\partial V(\alpha_i)}\left(U_i - \frac{1-F(U_i)}{f(U_i)}\right)}{Q_i(V(\alpha_D^*),V(\alpha_D^*))\frac{1-F(U_i)}{f(U_i)}\left(1+\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}\right)}$$

Proof. $B_i(\alpha_D^*, \alpha_D^*)$ approaches ∞ as α_i approaches 1, so no platform will ever set $\alpha_i = 1$. Assumption 2 ensures that $\alpha_i = \alpha_{-i} = 0$ is never an equilibrium. Assumption 3, Assumption 2 and continuous differentiability of $Q_i(V(\alpha_i), V(\alpha_{-i}))$ together ensure that a solution to Equation (25) in α_i exists for some range of $\alpha_{-i} > 0$ and that this solution is both necessary and sufficient for profit maximization if the optimal response to α_{-i} is positive. Since $V(\alpha)$ is continuous in α and $\alpha_i, \alpha_{-i} \in [0, 1]$, the payoff functions are continuous and (by assumption) concave in the platform's own strategies, and the strategy spaces are compact. Therefore by theorem 1.2 in Fudenberg and Tirole (1991) an equilibrium in pure strategies exists and by Assumption 2 it must have at least one $\alpha > 0$.

Furthermore, there must be an equilibrium which has symmetric $\alpha > 0$. Define $\alpha(\alpha_{-i})$ as platform i's reaction function. Since $Q_i(V(\alpha_i), V(\alpha_{-i}))$ is continuously differentiable, the solution to Equation (25) must vary with α_{-i} continuously. $\alpha(\alpha_{-i})$ is the solution to Equation (25) if this solution is non-negative or 0 if it is not, so $\alpha(\alpha_{-i})$ must be a continuous function. As the domain of $\alpha(\cdot)$ is $[0, 1]$, and the range is contained in this compact, convex set, Brouwer's fixed point theorem implies that there exists α_{-i} such that $\alpha(\alpha_{-i}) = \alpha_{-i}$. Furthermore, from Assumption 2 this α_{-i} must be positive, and so the symmetric solution is determined by Equation (25). ■

Lemma 6 states that under relatively mild conditions a symmetric equilibrium exists where both platforms admit a positive number of low-quality sellers if their retail markets would be highly competitive without them. This does not rule out the possibility of an asymmetric equilibrium, but I leave analysis of such equilibria for future work. Proposition 5 compares α^*, the equilibrium curation decision in the monopoly model, to α^*_D, the equilibrium decision in the duopoly.

Proposition 5. *Under Assumptions 1 to 3, if the equilibrium in the multi-platform model is symmetric then* $\alpha^*_D > \alpha^*$ *if and only if*

$$\frac{\frac{\partial Q_i(V(\alpha^*),V(\alpha^*))}{\partial V(\alpha_i)}}{Q_i(V(\alpha^*),V(\alpha^*))} < \frac{Q'(V(\alpha^*))}{Q(V(\alpha^*))}$$

That is, the equilibrium proportion of low-quality sellers increases with the addition of a second platform if and only if demand is less elastic in the duopoly platform case.

Proof. Assumptions 1 to 3 ensure sufficiency and necessity of the first order conditions. The conclusion comes directly from comparing Equation (25) to Equation (8) and the assumption of concavity in both cases. The presence of $B_i(\alpha^*_D, \alpha^*_D)$ and $B(\alpha^*)$ means that the comparison isn't immediately obvious. However, it is easy to show that

$$\begin{aligned} B_i(\alpha^*_D, \alpha^*_D) =& \frac{(1-F(U_i))}{(1-\alpha)\int_{U_i}^{\bar{\epsilon}}(\epsilon - U_i)f(\epsilon)d\epsilon}\left(\frac{1}{1+\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}}\right. \\ &\left. + \frac{\frac{\partial Q_i(V(\alpha^*_D),V(\alpha^*_D))}{\partial V(\alpha_i)}}{Q_i(V(\alpha^*_D),V(\alpha^*_D))}\frac{p_i\left(U_i - \frac{1-F(U_i)}{f(U_i)}\right)}{\frac{1-F(U_i)}{f(U_i)}\left(1+\frac{f(U_i)^2+f'(U_i)(1-F(U_i))}{f(U_i)^2}\right)}\right) \end{aligned}$$

and

$$B(\alpha^*) = \frac{(1 - F(U))}{(1-\alpha)\int_U^{\bar{\epsilon}}(\epsilon - U)f(\epsilon)d\epsilon}\left(\frac{1}{1 + \frac{f(U)^2 + f'(U)(1-F(U))}{f(U)^2}} + \frac{Q'(V(\alpha^*))}{Q(V(\alpha^*))}\frac{p\left(U - \frac{1-F(U)}{f(U)}\right)}{\frac{1-F(U)}{f(U)}\left(1 + \frac{f(U)^2 + f'(U)(1-F(U))}{f(U)^2}\right)}\right)$$

Given that both $B(\alpha^*)$ and $B_i(\alpha_D^*, \alpha_D^*)$ are positive, the only difference between the two first order conditions is the comparison of $\frac{Q'(V(\alpha^*))}{Q(V(\alpha^*))}$ to $\frac{\frac{\partial Q_i(V(\alpha_i),V(\alpha_{-i}))}{\partial V(\alpha_i)}}{Q_i(V(\alpha_i),V(\alpha_{-i}))}$. If the latter is smaller at $\alpha_i = \alpha_{-i} = \alpha^*$, then by Assumption 3 α_D^* must be larger than α^* for Equation (25) to hold. ■

Because the effects of α on price are identical in the two models, any differences in the curation decision must come from differences in the elasticity of the participation function. It would be unreasonable to expect $Q(V(\alpha^*)) < Q_i(V(\alpha^*), V(\alpha^*))$, however it is entirely possible that consumer participation might be less elastic in the equilibrium of a competitive market. The larger set of choices available to consumers in a duopoly environment means that consumers are more likely to have a strong preference for the platform they choose in equilibrium, (or a strong aversion to both options if they choose the outside option) so a consumer will attract fewer new consumers if it lowers α. This effect is strengthened if $\frac{\partial^2 Q_i(V(\alpha_D^*),V(\alpha_D^*))}{\partial V(\alpha_i)\partial V(\alpha_{-i})} > 0$ since this implies that as platform $-i$ allows more low-quality sellers, platform i loses fewer consumers when it increases α_i.