Lowering the Garden Wall: Marketplace Leakage and Quality Curation

Ben Casner*

September 11, 2024

Abstract

If a policy change reduces the friction that users face when shifting their transaction off of a market platform once they have used its interface to find one another, this threatens the platform's profit. I show in this paper that a marketplace can profitably respond to such a change by reducing the intensity with which it screens sellers while simultaneously offering a robust refund policy. This reduces the appeal of transacting off the platform and reduces the incentive of sellers to steer consumers to direct transactions. Introducing a competing platform increases seller welfare and eliminates this trade-off by deterring platforms from allowing low-quality sellers to join their marketplaces.

JEL classification:

Keywords: app stores, platform governance, marketplace leakage

*bcasner@ftc.gov The views expressed in this article are those of the author and do not necessarily reflect those of the Federal Trade Commission or any individual Commissioner. I would like to thank Robin Lee, Matt Leisten, Tesary Lin, Marius Schwartz and Doug Smith, as well as participants at the 2024 FTC Internal Research Conference, 2024 IIOC conference (Boston), and 2024 BU Platform Strategy Research Symposium for valuable comments and feedback.

1 Introduction

Many apps in mobile app stores avoid paying commissions by not allowing users to pay for services or sign up for subscriptions directly through the app store's payment system. For example, in Apple's app store ecosystem Spotify allows existing users to sign in to their accounts and stream media, but does not have options to use Apple's payment system to sign up for services. Until recently apps like Spotify did not include links directing users to their firms' online sign up flow because doing so was specifically forbidden by app store. However, in recent years Apple began allowing external links as a response to regulatory pressure, and in the United States the ruling in *Apple vs. Epic Games* specifically required that Apple allow apps on its store to direct users to external payment options (*Epic games, inc. v. apple inc.* 2021).[1]

More generally, there has been a broad push to ban anti-steering clauses, with Korea passing a law in 2021 requiring app stores to allow alternative payment systems, and the recent Digital Markets Act in the EU includes a clause specifically forbidding dominant platforms from including anti-steering provisions. This creates an avenue for platform leakage (Hagiu and Wright 2023) which may create a constraint on the platform's commissions. With direct anti-steering clauses forbidden, the platform may want to engage in less direct methods to reduce the appeal of steering.

As part of apple's response to the Epic ruling's required permission of direct payment links, when a user opens these links, Apple directs them to a warning screen before opening the app's browser sign-up flow (see figure 1).

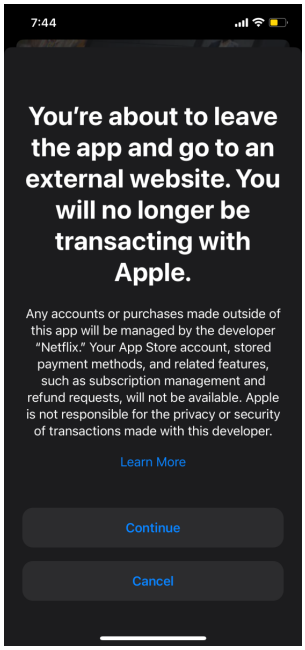


Figure 1: Warning displayed by Apple upon users leaving the app store through an external payment link

This warning screen highlights the fact that users cannot request a refund through the Apple app store, which implies that if the user is dissatisfied or encounters a dishonest developer then

[1] In mid 2024 Apple responded to this ruling by changing its terms of service to only allow the links with restrictive conditions and taking a high commission on those sales. Epic has stated it plans to sue over this change, saying that it does not conform to the intent of the original ruling.

they would be unable to get a refund (or would need to engage in a troublesome chargeback through their credit card company) if the developer refuses to provide one. This warning could be an effective leakage prevention strategy if consumers find the threat credible and are therefore unwilling to be steered toward direct transactions by app developers. This credibility requires that there must be a sufficient presence of unsatisfactory apps whose developers refuse to provide refunds. Which raises the question of whether this decision creates an incentive for Apple to be less diligent in curating the quality of apps on its store so that sellers who have a choice between paying a commission on transactions through the App store or a cheaper direct payment will have incentive to transact through the store.

To answer this question, I develop a model where consumers use search for sellers to transact with via a platform's search portal. The sellers can decide to either transact via the platform, or they steer consumers toward direct transactions. If the sellers allow on-platform purchases then they must pay a commission to the platform, but if they steer consumers to direct transactions then they compensate consumers for a switching cost which represents both the inconvenience to consumers and the difficulty of notifying consumers about the existence of a direct option. The platform's revenue comes from charging an entrance fee to consumers and an ad valorem commission to sellers.

In the benchmark model, the platform charges a commission which makes sellers indifferent about steering and an entrance fee which makes consumers indifferent about participating on the platform. The platform's ability to extract surplus from consumers is more efficient than its ability to extract from sellers, therefore its profit is increasing as consumers retain more of the gains from trade. The main limitation on the platform's ability to extract surplus from sellers is the switching cost, so the platform's profit is increasing in the switching cost.

In Section 5, I introduce low-quality sellers whose products provide negative utility but who are ex ante and interim (from the perspective of the market game) indistinguishable to consumers from high-quality sellers. If a consumer purchases from a low-quality seller via the platform they can seek an ex post refund, in which case the consumer receives 0 utility and the seller receives no revenue from that transaction. Therefore the low-quality seller can only make positive profit by steering consumers to direct transactions. The presence of these low-quality sellers reduces consumers' confidence in the quality of sellers steering to direct transactions, which reduces their willingness to pay and thus reduces the profitability of steering relative to allowing on-platform transactions and hence allows the platform to charge a higher commission. If the proportion of low-quality sellers is high enough consumers will only transact via the platform, and so the platform can extract all surplus from sellers. The platform controls the proportion of low-quality sellers, and the proportion of low-quality sellers that it allows is decreasing in both the switching cost and consumers' search costs. If the switching cost is sufficiently high, then the platform does not need to discourage direct transactions so the platform does not admit low-quality sellers. If the switching cost is low but the search cost is not, then the platform admits some low-quality sellers to discourage sellers from steering, but not so many that consumers are willing to to engage in additional search required to avoid transacting directly with sellers. If the switching and search cost are both low, then the platform admits just enough sellers such that consumers are only willing to transact through the platform.

In Section 6 I assume that regulators have required the platform holder to allow for a competing

marketplace to be installed on the same device. This platform does not charge a separate entrance fee, but does take a commission on seller revenue. I assume that sellers can multi-home freely, while consumers single-home on one platform. This competitive bottleneck setup (Armstrong 2006) leads to intense competition for consumers, forcing the platforms to engage in a high level of screening or risk losing all transaction revenue. The reduced presence of low-quality sellers makes steering consumers more attractive for sellers, which leads to a lower commission equal to the benchmark without low-quality sellers. Further, because competition prevents the platforms from reducing their screening effort (i.e. admitting low-quality sellers), the tradeoffs from policies which reduce switching costs are eliminated.

In Section 7 I describe two extensions to the model. In the first I hold the proportion of low-quality sellers fixed, but allow the consumers to detect low-quality sellers with a positive probability which is controlled by the platform. Similar to the previous section, the platform will reduce consumers' ability to distinguish high- and low-quality sellers if the switching cost is sufficiently low, but unlike the previous section this distinction is binary. If the switching cost is very low and consumers do not retain too much surplus from trade (gross of the entrance fee charged by the platform), then the platform will set the probability sufficiently low such that consumers are unwilling to engage in direct transactions. On the other hand, if the switching cost is relatively high and/or consumers share of gains from trade is sufficiently high, the equilibrium detection probability is independent of the switching cost. This is because the detection probability and switching cost both influence sellers' *per-transaction* revenue in an additive fashion, but do not have an extensive impact on the number of transactions occurring on the platform.

In the second extension I allow for endogenous price-setting by the sellers. I find that including low-quality sellers reduces consumer search intensity similar to Casner (2020) and Eliaz and Spiegler (2011), but the essential logic of the monopoly model is unchanged. If the switching cost is low and search cost is not too high, the platform will admit a positive proportion of low-quality sellers to deter direct transactions.

2 Related Literature

2.1 Disintermediation and Marketplace Leakage

Much of the extant literature on platform disintermediation (which I will henceforth interchangeably refer to as marketplace leakage) has focused on the limitations that disintermediation poses on platforms' pricing power and various policy decisions such as price parity clauses (Wang and Wright 2020), investment in transaction benefits (Lin, Nian, and Foutz 2022), provision of product information to consumers (Bar-Isaac and Shelegia 2023), seller/buyer communication, and intra-platform competition (Hagiu and Wright 2023). Jing (2018) adopts a similar model to mine but is closer to Wang and Wright (2020) in that his model focuses on the interaction between search costs, information externalities between various sales channels, and the effect of price parity clauses. Chang and Miller (2024) study the effects of a South Korean law requiring smart phone platform holders to allow app developers to permit alternative payment options, but find little effect from the policy, which they attribute partly to the small size of the Korean market and partly due to consumer mistrust of direct transactions. One of the closest papers to mine is

Gu and Zhu (2021), who find that as sellers build trust with consumers, they can more easily convince them to transact directly and avoid platform fees, however they examine this empirically in the context of a labor platform with repeated interactions. In contrast my model considers one-shot interactions and the implications of a platform refund policy when consumers have limited information about seller trustworthiness and no way to gain more confidence pre-purchase.

2.2 Quality Curation of Sellers and Platform governance

This article also fits into the literature on quality control in platforms and other non-price policies (Eliaz and Spiegler 2011; Teh 2022; Choi and Jeon 2023). The focus on inclusion of low-quality sellers and their effect on consumer search is similar to that of Casner (2020) and the search obfuscation literature (Ellison and Ellison 2009; Hagiu and Jullien 2011; Heresi 2023). Most of the those articles focus on the impact of platform design decisions on intra-platform competition, while here I focus specifically on the presence of low-quality sellers causing consumers to value the platform's implicit warranty more. Ershov (2024) takes an empirical look at app store marketplacees and finds that congestion effects increase consumer search costs, but his paper does not specifically consider removing low-quality apps and instead focuses on how congestion impacts product discovery while my primary focus is on consumer trust.

Perhaps the closest paper to mine is Etro (2021), he looks at the effects of platform business models on pricing incentives, quality curation of app stores, and the incentives of app stores to acquire developers in order to secure exclusives for their own platforms. However his focus is primarily the comparison of platforms which monetize consumers through an entry fee as opposed to simply taking commissions from sellers and how it affects inter-platform competition. His approach to modeling curation is highly reduced form to maintain tractability and because it is not the focus of his model.

3 Benchmark Model

To analyze the questions posed in the introduction, I develop a model with a monopoly platform, continuum of consumers with mass 1, and a continuum of sellers. The platform sets a percentage commission ξ that it takes from the sellers' gross revenue and a participation price ρ that it charges consumers for access to the platform. The sellers can choose to allow transactions through the platform in which case they receive gross revenue μV where V is the total value of trade that is homogeneous across both sellers and consumers (assuming a taste match, see below), and μ is the proportion of that value captured by sellers. Given the platforms commission this leads to net revenue $(1-\xi)\mu V$ per transaction when trade occurs through the platform. Alternately sellers can disallow on-platform transactions and instead steer consumers to direct transactions. If they do so, this imposes a switching cost t reflecting the difficulty of finding the alternative payment, the additional hassle costs of entering new payment information and any foregone convenience benefits of transacting on the platform. This switching cost is fully borne by the sellers, but sellers pay no commission to the platform for direct transactions, meaning they receive revenue $\mu V - t$ per customer when they steer. Consumers have access to an outside option whose value is normalized to 0, and they must search on the platform to find sellers following a

Wolinsky (1986) style undirected search process. They pay a search cost s for each observation, and with probability σ the seller they observe is a taste match, in which case the value of trade is V, and with probability $1-\sigma$ the seller's product does not match the consumer's taste, in which case the value of trade is 0. The timing of the model is as follows:

1. The platform sets commission ξ and price ρ
2. Sellers make steering decisions.
3. Consumers make participation decisions and search

3.1 Modeling Decisions

Low-quality sellers: The low-quality sellers in this model should be thought of as "bad" products that are unsatisfactory, rather than actively malicious products which would e.g. steal personal information, as the latter would have a far more deleterious impact on the reputation of the platform. We should also think of the sellers in these model as being small firms without a well established reputation. Larger firms (such as Epic Games) would have known quality and also do not need the platform's search functionality for product discovery, unlike the sellers in this model.

Reduced form trade and pricing: The reduced form search and trade formulation in this paper is broadly similar to Casner and Teh (2024). Adopting a reduced form trade formulation instead of allowing endogenous pricing is obviously quite a strong assumption, but it has the benefit of eliminating price effects from the benchmark model. This isolates the warranty effect which is the main focus of the paper instead of mixing that effect with the search obfuscation effect from low quality sellers which is explored more thoroughly elsewhere in the literature.[2] In the appendix I solve a richer model allowing endogenous pricing and show that the the main effect is robust to this extension. See Section 7 for details.

Ad valorem commissions: I choose to focus on ad valorem commissions as this is an institutional feature of the motivating market (and indeed part of the impetus for the *Apple v. Epic* lawsuit). Changing to a flat fee per-transaction would not significantly change the main results as the limitation on the platform's ability to extract seller value is the sellers' ability to steer consumers to direct transactions. If the platform were instead to charge an entry fee for sellers, then it would be able to extract all of the sellers' expected surplus, but this is largely an artifact of the simplifying assumptions in my model and would not hold true in the face of significant seller heterogeneity.

Consumer participation fee ρ: Many of the platforms that form motivating examples for this paper are associated with specific hardware. Therefore it is important to acknowledge that platforms which derive significant revenue from device sales have different incentives than those which are entirely monetized through on-platform interactions (Etro 2021). I include ρ as a stand in for these device sales, so for example in the case of Apple ρ would represent the price of an iPhone. By including this entrance fee from the beginning, I show that my results are robust even in the face of device funded platforms (to use Etro's terminology).

[2] See for example Eliaz and Spiegler (2011) and Casner (2020)

The platform's direct control of α: Allowing the platform to directly chose the proportion of low-quality sellers it admits on the platform is obviously a significant abstraction. A richer model of screening might include a base rate of low-quality sellers arriving at the marketplace, which the platform can then engage in costly effort to screen out. This does not significantly change the results of the model, but does matter for the interpretation. My results that the platform sets higher α in the face of lower t should be thought of as a reduction in the marginal benefit to the platform of engaging in screening effort, rather than the platform intentionally recruiting low-quality sellers to its marketplace.

4 Benchmark Equilibrium

4.1 Consumer Utility and Demand

The expected value to a consumer of searching once is

$$\sigma \times (1-\mu)V + (1-\sigma) \times 0 - s$$

Note that because the switching cost t is borne fully by sellers, the proportion of sellers who steer vs. those who do not does not affect this value of search. We assume the value of consumers' outside option is 0, so consumers will choose to search so long as $\sigma \times (1-\mu)V > s$ (which I henceforth assume to be the case). Further, it is obviously the case that $(1-\mu)V > \sigma \times (1-\mu)V - s$, so a consumer who has searched and found a taste match will stop searching and trade rather than search again. Therefore consumers will stop as soon as they find a match, which means the expected number of searches is $\frac{1}{\sigma}$ and consumers' ex ante expected utility from participating in the platform is

$$(1-\mu)V - \frac{s}{\sigma} - \rho$$

Assume consumers will participate if indifferent. Then given homogeneity of consumers, they will all participate if the following inequality is satisfied

$$\rho \leq (1-\mu)V - \frac{s}{\sigma} \tag{1}$$

4.2 Seller Steering Decisions and Profit

Given a mass 1 of consumers, the number of consumers who will stop at a given seller on their first visit is σ. The remaining $1-\sigma$ consumers will search a second time, so the expected number of consumers who will visit a seller on their second visit is $1-\sigma$, they too will stop with probability σ, so the mass of consumers searching three times is $(1-\sigma)^2$ and so on. This implies that each sellers' demand is $\sum_{j=0}^{\infty}(1-\sigma)^j\sigma = \frac{\sigma}{1-(1-\sigma)} = 1$, and so sellers' profit if they allow transactions through the platform is

$$\pi = (1-\xi)\mu V \tag{2}$$

Because the switching cost is borne entirely by sellers, their profit if they steer consumers to direct transactions is

$$\pi^{steer} = \mu V - t \tag{3}$$

Sellers will prefer on-platform transactions if

$$\begin{aligned} \pi &\geq \pi^{steer} \\ (1-\xi)\mu V &\geq \mu V - t \\ \frac{t}{\mu V} &\geq \xi \end{aligned} \tag{4}$$

Again assuming participation on the platform when indifferent, seller homogeneity implies that all sellers will allow on-platform transactions or all sellers will steer depending on the value of ξ.

4.3 Platform Decisions and Profit

If all consumers participate and no seller steers consumers to direct transactions, then the mass of transactions will be equal to 1, so the platform's profit is given by

$$\Pi = \begin{cases} \rho + \xi\mu V & \xi \leq \frac{t}{\mu V} \text{ and } \rho \leq (1-\mu)V - \frac{s}{\sigma} \\ 0 & \xi > \frac{t}{\mu V} \text{ or } \rho > (1-\mu)V - \frac{s}{\sigma} \end{cases}$$

The platform's profit is increasing in both ρ and ξ so long as neither of them violates their relevant participation inequalities, so both 1 and 4 will bind and in equilibrium

$$\xi = \frac{t}{\mu V} \tag{5}$$

and

$$\rho = (1-\mu)V - \frac{s}{\sigma} \tag{6}$$

Which gives equilibrium platform profit

$$\Pi = (1-\mu)V - \frac{s}{\sigma} + t \tag{7}$$

The following proposition follows directly from inspection of 7

Proposition 1. *Platform profit is decreasing in sellers' ability to capture surplus (μ) and increasing in the switching cost (t).*

The platform is able to capture all consumer surplus, while it only captures part of seller surplus, so as μ increases the platform is less able to capture the value of trade in on its platform. Second as more value is allocated to sellers, their benefit from steering transactions off platform increases, which in turn reduces the proportion of seller surplus that the platform is able to capture. This additional effect cancels out the upward platform-profit effect of seller revenue increasing in μ. The first effect depends on part on homogeneity of consumers, if the platform were only able to capture a small proportion of consumer surplus and the sellers' value of steering consumers were then the marginal effect of shifting surplus allocation from consumers to sellers could instead be positive, but the steering option effect would still remain, and so platform profit would decrease in μ regardless.

The impact of t is more straightforward. Platform profit is increasing in t because this reduces the value of the steering option for sellers, allowing the platform to capture additional surplus while maintaining seller participation. Together, the results of Proposition 1 imply that the platform wants to shift as much surplus as it can to consumers in order to better extract it, and it also wants to make it as difficult as possible for sellers to steer consumers to direct transactions. Calling back to the motivating example of the *Epic v Apple* ruling, we would expect that forcing Apple to allow app developers to link directly to outside payment options (effectively lowering t) would decrease profits from the App store, which would motivate Apple to engage in policies that shift surplus from sellers to consumers, and also could lead to more extreme responses as I will outline in the next section.

5 Preventing Leakage via Low Quality-Sellers

One of the advantages of transacting through the platform instead of directly is that the platform provides a sort of warranty in that consumers who are dissatisfied with their purchase can seek a refund through the platform. One way that a platform could indirectly raise the switching cost in response to an external force imposing direct reductions in t would be to highlight this guarantee. Apple's warning to consumers as highlighted in Figure 1 is an example of exactly this. Of course, this warning will have little effect when consumers are transacting with large companies that have an established reputation, therefore in order for this threat to be credible the platform would need to relax the intensity with which it screens the quality of sellers on its store. By allowing these apps into its previously strictly walled garden, the platform increases consumers' uncertainty about the trustworthiness of direct transactions and which reduces their appeal despite the lack of switching cost.

To model this strategy, I suppose the platform allows a proportion α of lower quality apps onto its marketplace who provide negative utility $-\ell$.[3] The distribution of payoffs from trade is now

$$\begin{cases} -\ell & \text{seller is low-quality} \\ 0 & \text{with probability } 1-\sigma \text{ if the seller is high-quality} \\ (1-\mu)V & \text{with probability } \sigma \text{ if seller is high-quality} \end{cases}$$

However consumers cannot observe their true match value draw until after they have made a purchase. Instead they observe a signal that is equal to their true value if they encounter a high-quality seller, and randomly drawn with probability σ of being positive if they encounter a low-quality seller and their true match value is $-\ell$. Once a consumer has purchased from a seller and observed a null match value, they have the option to seek a refund of the price they paid. This refund will be successful if they purchased through the platform, but will be denied if they purchased directly. Upon receiving a successful refund the consumers' payoff is 0, but if the refund is denied their net payoff remains $-\ell$.[4] Regardless of the outcome of their request

[3] ℓ is a catchall for the negative utility from receiving an unsatisfactory product and no refund. It captures the fact that the low-quality sellers will charge a positive price, but could also be thought of as incorporating bad feelings from the interaction. Note that if prices were endogenous in this model and low-quality sellers charged a different price than high-quality seller, consumers could distinguish between the two seller types easily, so the only equilibrium would involve a pooling strategy.

[4] In principle the post-refund payoff might be negative if monetary relief does not alleviate all of a consumer's

I assume consumers do not search for and purchase a different product after purchasing from an unsatisfactory seller.[5] Given the refund policy, the low-quality sellers only make a positive profit if they make a sale through the direct channel, so they will only offer direct purchases. Incorporating this observation, consumers' search and purchase decision process is shown visually in Figure 2.

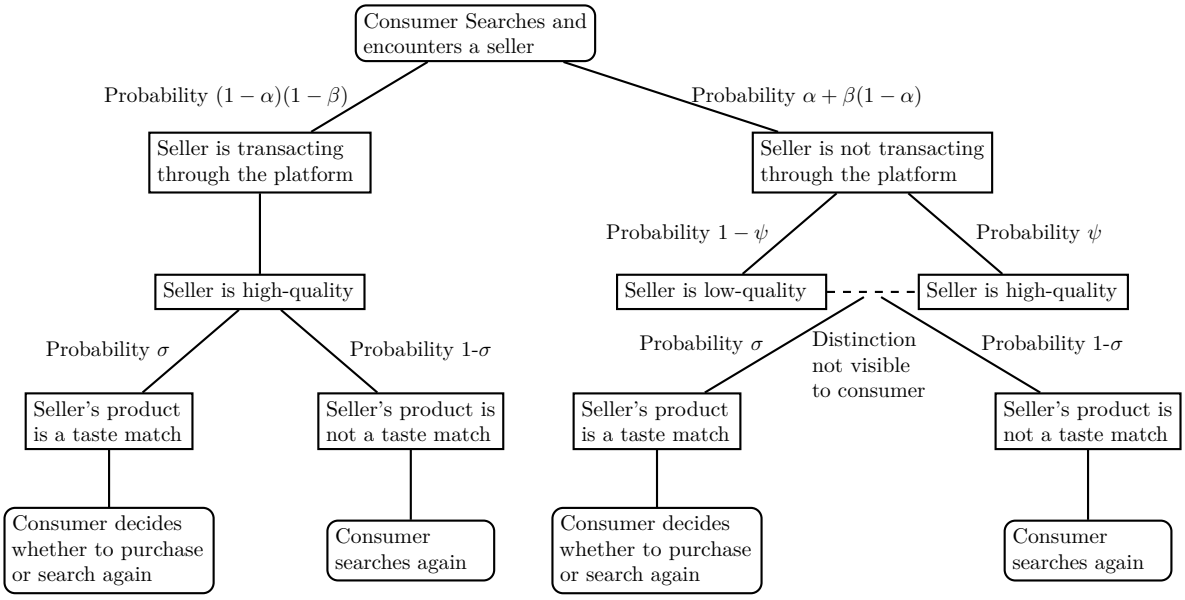


Figure 2: Consumers' search and purchasing process. Right angled box nodes indicate events. Boxes with rounded corners indicate decisions.

If no high-quality seller steers consumers to direct purchases then observing a direct transaction fully reveals a seller as low-quality and so no consumer will purchase directly. The low-quality sellers make 0 profit through both channels so they are indifferent between allowing direct purchases and not.

5.1 Consumer Utility and Demand

Let β represent the proportion of high-quality sellers who steer consumers to direct transactions and α be the proportion of low-quality sellers admitted by the platform. Then the probability a seller is high-quality conditional on a consumer observing that they are steering to direct transactions is

$$\frac{\beta(1-\alpha)}{\beta(1-\alpha)+\alpha} \equiv \psi$$

In this case, the expected value to a consumer of trading directly transacting with a seller (recall that the switching cost t is fully absorbed by sellers) is

$$\psi(1-\mu)V - (1-\psi)\ell$$

dissatisfaction, but I set it to 0 here to avoid unnecessary notation as a more general payoff would not make a difference to my results.

[5] In practice I would expect that some proportion of consumers would search for a replacement product after receiving a refund, or possibly even after failing to receive a refund. Stationarity of the search problem implies that this would increase the value of consumer participation and limit the negative impacts of allowing these low quality sellers on the platform, which would only strengthen the results of this section, so I omit that possibility to aid tractability and exposition.

Given ex ante symmetry of sellers, a consumer cannot benefit from searching again if they are willing to stop at another seller offering direct transactions, therefore the consumer will be willing to stop and accept this offer if the value of doing so is greater than the value of searching and only purchasing from a seller through the platform. Denote the value of searching and only purchasing through the platform by $\mathbb{V}^p$. That value of search is given by the recursive formulation

$$\mathbb{V}^p = \sigma(1-\alpha-\beta(1-\alpha))(1-\mu)V - s + (1-\sigma+\sigma\alpha+\beta(1-\alpha))\mathbb{V}^p$$

Which simplifies to

$$\mathbb{V}^p = (1-\mu)V - \frac{s}{\sigma(1-\alpha)(1-\beta)}$$

So a consumer will never stop and purchase from a seller offering direct transactions so long as

$$(1-\mu)V - \frac{s}{\sigma(1-\alpha)(1-\beta)} \geq \psi(1-\mu)V - (1-\psi)\ell$$

Which, after some algebra, is equivalent to

$$(1-\mu)V + \ell \geq \left(\frac{s}{\sigma(1-\alpha)(1-\beta)}\right)\frac{\beta(1-\alpha)+\alpha}{\alpha} \tag{8}$$

The right hand side is a convex function that approaches infinity as α approaches 0 or 1. The intuition is that if α is close to 0, then it is not worth searching again as the likelihood of encountering a low-quality product is very low, while if α is close to 1 it takes so long to find a seller transacting through the platform that it is not worth searching for one. I assume that V is sufficiently large such that there is some range of α where 8 is satisfied. Following similar steps as above, the value of search when a consumer is willing to purchase from any seller is

$$\mathbb{V}^{all} = (1-\alpha)(1-\mu)V - \alpha\ell - \frac{s}{\sigma} \tag{9}$$

Comparing with $\mathbb{V}^p$, we can easily show that $\mathbb{V}^p \geq \mathbb{V}^{all}$ is the exact same cutoff as 8, so it is rational for a consumer to search only for sellers transacting on the platform when this inequality is satisfied.

Denote the minimum value of α such that 8 is satisfied by $\tilde{\alpha}$.

Denote by $\bar{\alpha}$ the α which solves

$$\sigma(1-\beta)(1-\alpha)(1-\mu)V = s \tag{10}$$

When $\alpha > \tilde{\alpha}$, consumers will only purchase through the platform, so their expected value of a single search is $\mathbb{V}^p$. For α below $\bar{\alpha}$, $\mathbb{V}^p$ which is consumers expected value of participating in the platform gross of ρ is positive. For any $\alpha > \bar{\alpha}$, no consumer will participate in the platform for any $\rho > 0$. When β is not too large, it is simple to show that $\bar{\alpha} > \tilde{\alpha}$.[6] Consumer search behavior can therefore be summarized as follows in Table 1.

The platform has no reason to set $\alpha > \bar{\alpha}$, so I henceforth ignore that case. Consumers' value of

[6] For details see the proof of Proposition 2.

α	Behavior
$\alpha > \bar{\alpha}$	Consumers do not search
$\tilde{\alpha} \leq \alpha \leq \bar{\alpha}$	Consumers search but only buy through the platform
$\alpha < \tilde{\alpha}$	Consumers search and purchase upon seeing any positive value signal

Table 1: Consumer search behavior at different levels of α

participation (gross of ρ) in the two remaining cases is summarized in Table 2

α	Value
$\tilde{\alpha} \leq \alpha \leq \bar{\alpha}$	$(1-\mu)V - \frac{s}{\sigma(1-\alpha)(1-\beta)}$
$\alpha < \tilde{\alpha}$	$(1-\alpha)(1-\mu)V - \alpha\ell - \frac{s}{\sigma}$

Table 2: Consumers' value of participation at different levels of α

When consumers are willing to purchase on seeing any positive value signal, they risk being scammed, however if they only purchase through the platform, they must pay an additional expected search cost. $\tilde{\alpha}$ represents the risk of an unsatisfactory transaction such that the additional search cost is worth avoiding a potentially low-quality seller.

5.2 Seller Steering Decisions and Profit

Given a proportion β of high-quality sellers steering consumers to direct transactions, a seller's profit if they allow on-platform transactions is

$$\pi = \begin{cases} \mu(1-\xi)V & \alpha < \tilde{\alpha} \\ \frac{\mu(1-\xi)V}{(1-\alpha)(1-\beta)} & \tilde{\alpha} \leq \alpha \leq \bar{\alpha} \end{cases}$$

Where the second case arises because consumers only buy through the platform, so the mass 1 of consumers will be divided equally among the mass $(1-\alpha)(1-\beta)$ of sellers allowing on-platform transactions.

For the sellers who do steer to direct transactions, I assume that consumers diminished trust in their quality reduces their payoff, such that the division of expected surplus (gross of switching cost) is the same, even though the total expected gains from trade are smaller.[7] The sellers' profit from steering consumers to direct transactions is then

$$\pi = \begin{cases} \psi\mu V - t & \alpha < \tilde{\alpha} \\ 0 & \tilde{\alpha} \leq \alpha \leq \bar{\alpha} \end{cases}$$

When $\beta = 0$ we have the following result:

Proposition 2. *If $\beta = 0$, no seller will individually deviate to steering consumers to direct transactions for any $\xi < 1$ when $\alpha > 0$.*

[7]This could be microfounded with e.g. a Nash bargaining process with negotiation weights μ and $1-\mu$. However the extension in the appendix with endogenous pricing produces a similar result where increasing α reduces the surplus of steering sellers, so this assumption helps this more tractable model capture the same dynamics. This result is also consistent with the empirical literature, Chang and Miller (2024) find evidence suggesting that app store users in South Korea appear to be less trustful of direct payments than of payments through the App Store.

Unless otherwise noted all proofs are relegated to the appendix. Proposition 2 says that no matter how high ξ gets, no individual seller can profitably begin steering consumers to direct transactions because consumers will have no confidence that the seller is not a low-quality seller, so steering leads to no demand and 0 profit.[8] This result clearly demonstrates the mechanism this model was intended to highlight, but the starkness of the result depends on $\beta = 0$. If consumers believe that a positive mass of high-quality sellers are not allowing on-platform transactions then there exists an interior ξ where sellers can make more profit from selling through the direct channel. Therefore for the rest of this section I assume that a minimum proportion $\underline{\beta}$ of the high quality sellers will steer consumers to the direct channel regardless of the platform's commission.[9]

With this assumption in place, and assuming $\alpha \leq \tilde{\alpha}$ the remaining $1 - \underline{\beta}$ high-quality sellers will transact through the platform so long as

$$\psi\mu V - t \leq \mu(1-\xi)V$$

Solving for ξ

$$\xi \leq 1 - \psi + \frac{t}{\mu V} \tag{11}$$

Comparing to 4, sellers' profit from steering is reduced by the diminished confidence consumers have in the quality of steering sellers, which means they will tolerate a higher commission before steering consumers off the platform. Alternatively, if $\alpha > \tilde{\alpha}$, the platform can set $\xi = 1$ and sellers making an active choice will still not steer as they earn 0 profit with either option.

5.3 Platform Decisions and Profit

The platform only takes commission on transactions that it handles directly. If consumers are purchasing from any seller, then some of those transactions happen off-platform, while if $\alpha \geq \tilde{\alpha}$ the platform is handling all transactions. This gives the following conditional profit function

$$\Pi = \begin{cases} \rho + (1-\alpha)(1-\underline{\beta})\xi\mu V & \alpha < \tilde{\alpha} \\ \rho + \xi\mu V & \tilde{\alpha} \leq \alpha \leq \bar{\alpha} \end{cases} \tag{12}$$

Given consumer and seller behavior derived in the previous subsections, the platform will set ρ such that consumers are just indifferent between participating and not, and ξ such that sellers are just indifferent between steering and not. We can substitute the values of ξ and ρ that achieve this indifference to get

$$\Pi = \begin{cases} (1-\alpha)(1-\mu)V - \alpha\ell - \frac{s}{\sigma} + (1-\alpha)(1-\underline{\beta})\left[1 - \psi + \frac{t}{\mu V}\right]\mu V & \alpha < \tilde{\alpha} \\ (1-\mu)V - \frac{s}{(1-\alpha)(1-\underline{\beta})\sigma} + \mu V & \tilde{\alpha} \leq \alpha \leq \bar{\alpha} \end{cases} \tag{13}$$

For $\alpha < \tilde{\alpha}$, the platform is trading off reduced consumer participation value and more off-platform trades against an increase in ξ. When $\alpha \geq \tilde{\alpha}$ platform profit is strictly decreasing in α as the

[8] Note that this is true even when the consumer has a positive probability of distinguishing a low-quality seller. See the online appendix for details.

[9] There are multiple possible motivations for this assumption. Examples include: strategic trembles, desire to keep sensitive commercial data away from the large platform, or small businesses with taste based aversion to transacting through the platform.

platform captures all seller surplus and increased α only reduces consumer surplus by making search more difficult, therefore the platform will never set $\alpha > \tilde{\alpha}$. If s is sufficiently small, it is obvious from inspection that the platform's profit is greater when $\alpha = \tilde{\alpha}$ than when it is less as all trade will flow through the platform rather than some of the trades happening through direct transactions. Denote the profit maximizing proportion of α set by the platform as α^*.

Proposition 3. *α^* is positive if s is sufficiently small or of μ is not too small and $\frac{1-\underline{\beta}}{\underline{\beta}}\mu V$ is sufficiently large relative to t and ℓ. α^* is weakly decreasing in t.*

If s is sufficiently small, then the direct cost to consumers from being steered by the platform to on-platform transactions is relatively small, so the platform sets $\alpha^* = \tilde{\alpha}$ and captures all of the surplus in the market. In this case the willingness of sellers to switch over to direct transactions is irrelevant as no consumer will be willing to take a direct transaction if offered. Because t is fully borne by sellers, $\tilde{\alpha}$ is constant in t and so the proportion of low-quality sellers does not change with the switching cost.[10] On the other hand, if s is sufficiently large such that $\alpha^* < \tilde{\alpha}$, then the platform is trading off an increased commission against a worsened consumer experience (and hence lower entry fee) and more transactions happening directly. As t increases sellers are less willing to steer consumers to direct transactions, so the marginal benefit of including low-quality sellers decreases and so does α^*. If μ is too small then sellers do not capture enough surplus for the platform to find extracting more from them worth the cost of degrading the consumer experience, and the concomitant lowered consumer fee ρ so $\alpha^* = 0$. Similarly, if ℓ is too large then even if μ is relatively high the reduction in consumer willingess to pay for entry to the marketplace outweighs any increase in commission and $\alpha^* = 0$.[11]

One of the major implications of Proposition 3 is that a policy change which reduces t (such as the decision in *Epic v. Apple*) would create an incentive for the platform to engage in leakage prevention strategies, and consequently one of the impacts we might expect from this decision is either less incentive for Apple to screen the quality of apps on its store and/or efforts such as the warning screen in Figure 1 to increase the salience to consumers of the possibility they might be scammed.

6 Introducing a Competing Platform can Increase the incentive to screen

Turning to a different court ruling, one of Epic's proposals as a remedy for anti-competitive behavior in its win against Google is to require easy installation of app stores to complete with Google Play on Android phones(*Epic games, inc. v. apple inc.* 2023). More broadly, one could conceive of a regulation requiring providers of hardware to allow easy loading of competing software platforms on to their devices. To evaluate this potential policy proposal I introduce a competing platform which operates on the same hardware as the original, but which does not charge an

[10] If we were to relax this assumption, then as t increases consumers would be less willing to engage in direct transactions and $\tilde{\alpha}$ would decrease in t.

[11] We might expect ℓ and μ to be very closely related, as both relate to sellers' ability to extract surplus from the consumers. However, I show that my main results are robust to endogenous pricing in the appendix and so I leave both variables in reduced form here to simplify exposition and avoid taking a stand on the exact relationship between ℓ and μ.

entrance fee. I index the two platforms by $k \in \{0, 1\}$ and WLOG I choose 0 to be the integrated platform which also charges the entrance fee ρ. Both platforms are free to choose their ad-valorem commissions ξ_k and screening intensity α_k. The sellers can make steering decisions independently for each platform, but can multi-home freely. Consumers single-home on one platform, but pay the entrance fee ρ regardless of which platform they choose to join.[12]

The timing of the model is now

1. Platform 0 sets entry fee ρ
2. 0 and 1 set ξ_k and α_k
3. Sellers make steering decisions
4. Consumers decide whether to pay ρ and choose which platform to participate on if they do
5. Participating consumers search and transact.

6.1 Consumer behavior

For the sake of tractability, I assume s is large enough such that $\max\{\alpha_k\} < \tilde{\alpha}$ and consumers will search and stop at their first positive match. A consumer's expected value of participation is therefore $\max\{\mathbb{V}_0^{all}, \mathbb{V}_1^{all}\}$ (for simplicity I will henceforth drop the superscript from the gross value of participation).[13] Consumers will participate on the platform which provides the highest expected value, and I assume that they randomize if indifferent.

6.2 The Seller Steering Decision is Unchanged

Sellers make steering decisions independently for each platform. All sellers participate on both platforms given the lack of multi-homing costs and non-negative profits on each, so given a mass m_k of consumers participating on platform k, each seller will transact $\frac{m_k}{1}$ times, and so the profit from allowing on-platform transactions is

$$\pi_k = m_k \mu (1 - \xi_k) V$$

and from steering

$$\pi_k^{steer} = m_k \left(\psi_k \mu V - t\right)$$

Where ψ_k is defined equivalently to ψ but using α_k and β_k. Comparing π_k and π_k^{steer} to find the ξ_k which makes sellers indifferent yields 11. So the seller problem is essentially unchanged from the monopoly model.

[12] This is the "competitive bottleneck" setup of Armstrong (2006). Sellers have more reason to multi-home than consumers—especially given that consumers in this model have unit demand—so similar to Casner and Teh (2024), even if I were to allow multi-homing by consumers it would not make much difference to my results as consumers would simply search exclusively on the platform where they have the highest expected payoff.

[13] I assume undifferentiated platforms here to maintain comparability with the monopoly framework. If I were to assume, e.g. Hotelling competition with a transport cost z then the comparison to previous sections would not be direct because even a monopolist would face a continuous consumer-side demand function in that context. That said, I expect that my results would generalize so long as the transport cost is not too high.

6.3 Platform Curation and Commissions

Given the assumption that $\max_k \alpha_k < \tilde{\alpha}$, profit for each platform is given by

$$\begin{aligned}\Pi_0 =&(m_0 + m_1)\rho + m_0(1-\alpha_0)(1-\underline{\beta})\xi_0\mu V\\ \Pi_1 =&m_1(1-\alpha_1)(1-\underline{\beta})\xi_1\mu V\end{aligned}$$

Where the mass of participating consumers is given by

$$m_k = \begin{cases} 1 & \mathbb{V}_k > \mathbb{V}_{-k} \\ \frac{1}{2} & \mathbb{V}_k = \mathbb{V}_{-k} \\ 0 & \mathbb{V}_k < \mathbb{V}_{-k} \end{cases}$$

From 9, $\mathbb{V}_k$ is strictly decreasing in α, so for any positive α_k, $-k$ can win the entire market by setting $\alpha_{-k} = \alpha_{-k} - \iota$ where ι is some infinitesimal reduction in the proportion of low-quality sellers. Competition will thus drive both α_k to 0 as the only equilibrium, which from 11 means that

$$\xi_0 = \xi_1 = \frac{t}{\mu V}$$

Since $\alpha = 0$ in equilibrium for both platforms, $\xi_0 = \xi_1$, and $\rho = \mathbb{V}_0 = \mathbb{V}_1$. With the decreased commission sellers' welfare increases under competition compared to the monopoly, and consumers are no worse off, while total welfare increases.[14] Interestingly, sellers are benefiting from competition despite the fact that the multi-homing paradigm is that of a competitive bottleneck. The reason is that the intense competition for consumers increases consumer confidence in the quality of sellers who choose to steer to direct transactions. This increases the value of the outside option for sellers for any given t. It also means that policies which decrease t do not result in higher α unlike in section 5.

7 Extensions

7.1 Changing Consumers' Ability to Detect Low-Quality Sellers Can Prevent Leakage

Another strategy that the platform could consider instead of curating the proportion of low-quality sellers is changing the degree to which they are detectable by consumers. For example, changing the salience of reviews on a product's display page or just showing averages and making the consumer go to more effort to look at detailed written reviews, or changing the strictness with which it polices review manipulation. In Section 5 I abstracted from this process by assuming that consumers had no ability to distinguish high- and low-quality sellers. In section B of the Online Appendix I show that changing consumers' information structure can have a similar effect at preventing leakage and also presents similar tradeoffs for the platform.

In contrast to the base model, I assume that consumers who encounter a seller receive a signal as to that seller's quality. With probability γ the signal accurately reports the seller's quality,

[14] This welfare statement ignores the sellers who are excluded from the market, but given that in this case that group mostly comprises low-quality sellers I do not think this is particularly concerning.

and with probability $1-\gamma$ the consumer receives a signal that the seller is high quality, meaning that consumers face a "fully revealing bad news" information structure.[15]

For this extension I assume the proportion of low-quality sellers α is fixed, and instead allow the platform to vary the signal informativeness γ, and I assume that the platform sets γ before buyers and sellers have a chance to act. The rest of the agents' decision variables and timing are unchanged.

For consumers' value of participation we can go through a similar derivation as for $\mathbb{V}^{all}$ to get

$$\hat{\mathbb{V}}^{all} = \frac{[(1-\alpha)(1-\mu)V - (1-\gamma)\alpha\ell] - \frac{s}{\sigma}}{1-\gamma\alpha}$$

$\hat{\mathbb{V}}^{all}$ is increasing as γ increases because the probability of purchasing from a low-quality seller approaches 0.

The refund policy means that only high-quality sellers will offer transactions through the platform, so the value of searching and transacting only through the platform $\mathbb{V}^{p}$ is unchanged and is constant in γ. There exists a critical value $\bar{\gamma} \geq 0$ such that consumers will stop and purchase at their first match with a positive signal for all $\gamma > \bar{\gamma}$. Further, if $\bar{\gamma} > 0$ then consumers are indifferent between the two search strategies when $\gamma = \bar{\gamma}$ and strictly prefer transacting through the platform when $\gamma < \bar{\gamma}$. I assume α is large enough that $\bar{\gamma} > 0$, such that consumers will only transact directly if $\gamma > \bar{\gamma}$.

The seller problem is not substantially different from Section 5, meaning that the limitations on ξ are essentially unchanged. Proposition 4 (proved in the online appendix) summarizes the platform's optimal strategy with regard to γ.

Proposition 4. *There exists a critical value of t denoted $\bar{t}$ such that $\gamma^* > \bar{\gamma}$ if $t > \bar{t}$ and $\gamma = \bar{\gamma}$ for $t \leq \bar{t}$ so long as μ is not too small. $\bar{t}$ is increasing in α and $\underline{\beta}$. $\gamma^* < 1$ if μ is not too small and ℓ not too large. γ^* is independent of t if $\gamma > \bar{\gamma}$ and increasing in α.*

Similar to Section 5, the platform is trading off reduced value extraction from sellers against an increased participation fee from consumers as γ increases. Interestingly, the profit maximizing level of γ does not depend on t so long as the level of t is large enough such that $\gamma > \bar{\gamma}$. This is because a change in γ does not impact the proportion of sellers steering consumers toward direct transactions and so γ provides a purely infra-marginal change in platform profits which is independent of t.[16] Nevertheless the main intuition from Section 5 that the platform will increase the salience of low-quality sellers in response to a decrease in t does survive if that reduction lowers t below the cutoff $\bar{t}$.

7.2 Allowing Endogenous Pricing Yields Similar Results to the Base Model

In section C of the Online Appendix I allow the sellers to price their product endogenously. Competition on the platform takes place via Wolinsky (1986) style search with match values ϵ

[15] I could consider a more general garbled structure a la Lewis and Sappington (1994), but doing so would be less tractable, and the literature on review manipulation (e.g. Smirnov and Starkov (2022)) has not found that generic information structures lead to drastically different results.

[16] This result likely would not be robust to a more complex model with e.g. a continuum of t. It would also not be true of the model with endogenous pricing as both t and ψ influence sellers' price elasticity with regard to platform policy changes.

following an unbounded log-concave distribution $F(\cdot)$ with associated PDF $f(\cdot)$. The switching cost t is subtracted from consumer utility and sellers compensate consumers for the inconvenience endogenously by lowering their price for direct transactions rather than being assumed to fully absorb it. Low-quality sellers must set the same price as high-quality sellers otherwise consumers could distinguish them easily. Finally I assume that sellers have a marginal cost of production c such that the ad-valorem commission influences their prices. Otherwise the model is essentially the same as in the previous sections. For a given proportion α of low-quality sellers and proportion β of high-quality sellers steering to direct transactions consumers will stop search if they encounter a match value greater than or equal to a reservation value U determined by

$$s = (1-\beta)(1-\alpha)\int_{U+p}^{\infty}[(\epsilon-p)-U]f(\epsilon)d\epsilon+(\beta(1-\alpha)+\alpha)\int_{\frac{U+p^{steer}+t}{\psi}}^{\infty}[\psi\epsilon-p^{steer}-t-U]f(\epsilon)d\epsilon \quad (14)$$

which implies that a consumer will stop and purchase from a seller allowing on-platform transactions and charging price p_j if

$$\epsilon_{ij} - p_j > U$$

which leads to demand

$$1 - F(U + p_j)$$

Equivalently, demand for a steering seller is

$$1 - F(\frac{U + p_j^{steer} + t}{\psi}) \quad (15)$$

.

Plugging these into seller profit functions, we get the on-platform price

$$p_j = \frac{c}{1-\xi} + \frac{1 - F(U + p_j)}{f(U + p_j)} \quad (16)$$

and the steering price

$$p^{steer} = c + \psi\frac{1 - F(\frac{U+p_j^{steer}+t}{\psi})}{f(\frac{U+p_j^{steer}+t}{\psi})} \quad (17)$$

Similar to the base model, the lower ψ, the lower the price that a seller can charge. Unlike the base model, seller profit does not change multiplicatively with ψ, but the difference is not large enough to change my conclusions.

The platform's problem is broadly similar to the base model, with essentially the same tradeoffs. Its profit maximizing curation strategy is summarized in Proposition 5

Proposition 5. *There is a profit maximizing level of α denoted α^* which is decreasing in t. $\alpha^* > 0$ if t and s are sufficiently small.*

The proof for this proposition is located in the online appendix. The main conclusions of the base model are replicated with endogenous pricing. If steering consumers to direct transactions is too easy, it can be profitable for the platform to reduce leakage by reducing consumer trust in the quality of steering sellers by admitting more low-quality sellers on to the platform.

8 Conclusion

I have explored the implications of policies which make it easier for sellers on a market platform to steer consumers toward direct transactions. The more the platform can limit sellers' ability to steer, the more it can charge in commissions as this limitation reduces the value to sellers of transacting outside the platform. If this limitation is taken away, then the platform will seek alternative methods to discourage consumers from transacting directly with sellers.

One avenue available to it is reducing screening efforts and allowing more low-quality sellers to participate while simultaneously offering refunds to consumers that request them. The refund discourages low-quality sellers from offering transactions through the platform, while the presence of low-quality sellers reduces consumers' confidence that a seller offering direct transactions is high-quality. The more low-quality sellers the platform admits, the more this warranty through refunds becomes salient to consumers' decision making and the more the sellers have to discount direct transactions if they are to attract any customers.

It is profitable for the platform to admit these low-quality sellers if it is otherwise easy for sellers to steer consumers toward direct transactions and if the sellers share of surplus is sufficiently high. This happens even though the platform is device funded (using the terminology of Etro (2021)) through an entrance fee paid by consumers. This result suggests that policy makers should be careful in designing regulations forbidding anti-steering policies as the platforms may engage in indirect strategies to to eliminate platform leakage.

One way for policy makers to mitigate these alternative strategies is to encourage platform competition. If the competition for consumers is sufficiently intense, then platforms are deterred from engaging in policies which reduce consumer welfare in order to extract more surplus from sellers.

References

Armstrong, Mark. 2006. Competition in two-sided markets. *The RAND Journal of Economics* 37 (3): 668–691.

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

Bar-Isaac, Heski, and Sandro Shelegia. 2023. Search, showrooming, and retailer variety. *Marketing Science* 42 (2): 251–270.

Casner, Ben. 2020. Seller curation in platforms. *International Journal of Industrial Organization* 72:102659.

Casner, Ben, and Tat-How Teh. 2024. Content-hosting platforms: discovery, membership, or both? *RAND Journal of Economics.*

Chang, Boyoon, and Keaton Miller. 2024. The short-run effects of opening mobile in-app payment systems: evidence from south korea. *Available at SSRN 4891074.*

Choi, Jay Pil, and Doh-Shin Jeon. 2023. Platform design biases in ad-funded two-sided markets. *The RAND Journal of Economics* 54 (2): 240–267.

Eliaz, Kfir, and Ran Spiegler. 2011. A simple model of search engine pricing. *The Economic Journal* 121 (556): F329–F339.

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

Epic games, inc. v. apple inc. 2021, Case No. 4: 20-cv-05640-YGR.

Epic games, inc. v. apple inc. 2023, Case No. 3: 20-cv-05671.

Ershov, Daniel. 2024. Variety-based congestion in online markets: evidence from mobile apps. *American Economic Journal: Microeconomics* 16 (2): 180–203.

Etro, Federico. 2021. Device-funded vs ad-funded platforms. *International Journal of Industrial Organization* 75:102711.

Gu, Grace, and Feng Zhu. 2021. Trust and disintermediation: evidence from an online freelance marketplace. *Management Science* 67 (2): 794–807.

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

Hagiu, Andrei, and Julian Wright. 2023. Marketplace leakage. *Management Science.*

Heresi, José Ignacio. 2023. Platform price parity clauses and consumer obfuscation. *The Journal of Industrial Economics* 71 (1): 291–322.

Jing, Bing. 2018. Showrooming and webrooming: information externalities between online and offline sellers. *Marketing Science* 37 (3): 469–483.

Lewis, Tracy R, and David EM Sappington. 1994. Supplying information to facilitate price discrimination. *International Economic Review,* 309–327.

Lin, Jinan, Tingting Nian, and Natasha Zhang Foutz. 2022. Disintermediation and its mitigation in online two-sided platforms: evidence from airbnb. In *Icis.*

Smirnov, Aleksei, and Egor Starkov. 2022. Bad news turned good: reversal under censorship. *American Economic Journal: Microeconomics* 14 (2): 506–560.

Teh, Tat-How. 2022. Platform governance. *American Economic Journal: Microeconomics* 14 (3): 213–54.

Wang, Chengsi, and Julian Wright. 2020. Search platforms: showrooming and price parity clauses. *RAND Journal of Economics* 51 (1): 32–58.

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

A Omitted Proofs

A.1 Proof of Proposition 2

Proof. When $\beta = 0$, then the right hand side of 8 is strictly increasing in α. Given the assumption that V is sufficiently large that there is some range of α where 8 is satisfied, this implies that the inequality must be satisfied when there are no low-quality sellers, meaning $\tilde{\alpha} = 0$. Therefore, any seller who deviates individually to steering consumers will make zero profit, while for any $\xi <= 1$ profit from on-platform transactions is non-negative. ■

A.2 Proof of Proposition 3

Proof. If s is sufficiently small, then $\alpha^* = \tilde{\alpha}$, which by definition is the maximum of the minimum α such that 8 is satisfied and 0. Because the switching cost is borne fully by sellers, 8 is invariant to t and so α^* is constant in t in this case.

If $0 < \alpha^* < \tilde{\alpha}$, then it is determined by the first order condition derived from taking the derivative of 13 with regard to α:

$$-(1-\mu)V - \ell - (1-\underline{\beta})\left[1 - \psi + \frac{t}{\mu V}\right]\mu V + \frac{\underline{\beta}(1-\alpha)}{[\underline{\beta}(1-\alpha)+\alpha]^2}(1-\underline{\beta})\mu V = 0$$

For any value of α, this derivative is decreasing in t, therefore the α at which it crosses from positive to negative must also decrease, and hence so does α^*. $\alpha^* > 0$ while s is large only if this derivative is positive at $\alpha = 0$, i.e.

$$\frac{1-\underline{\beta}}{\underline{\beta}}\mu V - (1-\underline{\beta})t - (1-\mu)V - \ell > 0$$

This can only be positive if μ is not too small, and $\frac{1-\underline{\beta}}{\underline{\beta}}\mu V$ is sufficiently large relative to t and ℓ. ■

B Online Appendix: Changing Consumers' Ability to Detect Low-Quality Sellers Can Prevent Leakage

Another strategy that the platform could consider instead of curating the proportion of low-quality sellers is changing the degree to which they are detectable by consumers. For example, changing the salience of reviews on a product's display page or just showing averages and making the consumer go to more effort to look at detailed written reviews, or changing the strictness with which it polices review manipulation. In Section 5 I abstracted from this process by assuming that consumers had no ability to distinguish high- and low-quality sellers. Here I show that changing consumers' information structure can have a similar effect at preventing leakage and also presents similar tradeoffs for the platform.

In contrast to the base model, I assume that consumers who encounter a seller receive a signal as to that seller's quality. With probability γ the signal accurately reports the seller's quality, and with probability $1-\gamma$ the consumer receives a signal that the seller is high quality, meaning

that consumers face a "fully revealing bad news" information structure.[17]

For this section I assume the proportion of low-quality sellers α is fixed, and instead allow the platform to vary the signal informativeness γ, and I assume that the platform sets γ before buyers and sellers have a chance to act. The rest of the agents' decision variables and timing are unchanged.

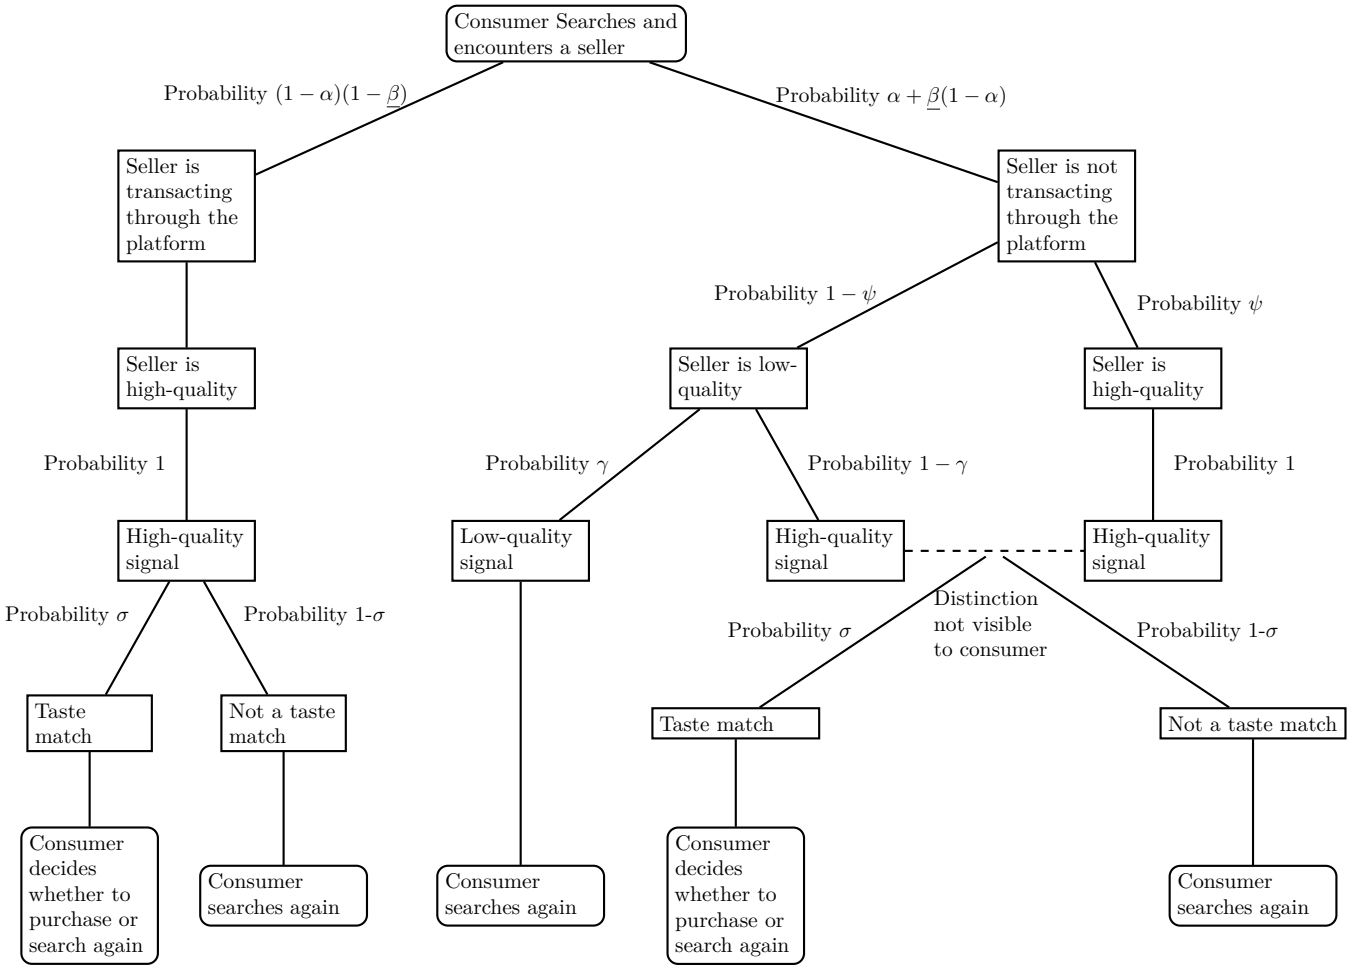


Figure 3: Consumers' search and purchasing process with a probability of detecting low-quality sellers. Right angled box nodes indicate events. Boxes with rounded corners indicate decisions.

B.1 Consumers

Because low-quality sellers are only making a profit on off-platform transactions, any seller offering transactions through the platform is—as in previous sections—known to be high quality, meaning that the change in information does not make a difference to consumers' beliefs. On the other hand, if the consumer encounters a seller offering off-platform transactions and receives a

[17] I could consider a more general garbled structure a la Lewis and Sappington (1994), but doing so would be less tractable, and the literature on review manipulation (e.g. Smirnov and Starkov (2022)) has not found that generic information structures lead to drastically different results.

positive signal, then from Bayes' rule the probability that this seller is high-quality is

$$\begin{aligned}\hat{\psi} \equiv & \frac{\psi}{\psi + (1-\psi)(1-\gamma)} \\ = & \frac{(1-\alpha)\beta}{(1-\alpha)\beta + \alpha(1-\gamma)}\end{aligned}$$

$\hat{\psi}$ is consumers' *signal adjusted* belief that a seller steering transactions off-platform is high-quality. The more informative (higher) γ is, the more consumers trust the high-quality signal. Because consumers never encounter a low-quality signal at a high-quality seller, they will always believe a low-quality signal is accurate.

Taking this into account, we can go through a similar derivation as for $\mathbb{V}^{all}$ to get

$$\hat{\mathbb{V}}^{all} = \frac{[(1-\alpha)(1-\mu)V - (1-\gamma)\alpha\ell] - \frac{s}{\sigma}}{1-\gamma\alpha}$$

$\hat{\mathbb{V}}^{all}$ is increasing as γ increases because the probability of purchasing from a low-quality seller approaches 0.[18] When $\gamma = 1$, the threat of purchasing from low-quality sellers is completely eliminated and the only effect of their presence is as a form of search obfuscation (Ellison and Ellison 2009; Eliaz and Spiegler 2011; Casner 2020).

Because only high-quality sellers will offer transactions through the platform, the value $\mathbb{V}^p$ of searching and only transacting via the platform is unchanged relative to the benchmark. Consumers will therefore search and purchase from the first seller they match and have a positive signal with if $\hat{\mathbb{V}}^{all} \geq \mathbb{V}^p$ and will instead only transact through the platform if that inequality is not satisfied. Expanding, that inequality is equivalent to:

$$\frac{[(1-\alpha)(1-\mu)V - (1-\gamma)\alpha\ell] - \frac{s}{\sigma}}{1-\gamma\alpha} \geq (1-\mu)V - \frac{s}{\sigma(1-\alpha)(1-\beta)}$$

The left hand side is increasing in γ, and approaches $\mathbb{V}^{all}$ from above as γ approaches 0. If α is small enough that $\mathbb{V}^{all} > \mathbb{V}^p$, then $\hat{\mathbb{V}}^{all} \geq \mathbb{V}^p$ for all γ. Note however that with $\gamma = 1$ $\hat{\mathbb{V}}^{all} > \mathbb{V}^p$ even if $\mathbb{V}^{all} < \mathbb{V}^p$. Therefore there exists a critical value $\bar{\gamma} \geq 0$ such that consumers will stop and purchase at their first match with a positive signal for all $\gamma > \bar{\gamma}$. Further, if $\bar{\gamma} > 0$ then consumers are indifferent between the two search strategies when $\gamma = \bar{\gamma}$ and strictly prefer transacting through the platform when $\gamma < \bar{\gamma}$. For the remainder of this section I assume α is large enough that $\bar{\gamma} > 0$, such that consumers will only transact directly if $\gamma > \bar{\gamma}$.

B.2 Sellers

Similar to Section 5, a seller's profit if if they allow transactions through the platform is given by

$$\pi = \begin{cases} \mu(1-\xi)V & \gamma > \bar{\gamma} \\ \frac{\mu(1-\xi)V}{(1-\alpha)(1-\beta)} & \gamma \leq \bar{\gamma} \end{cases}$$

[18] There exist values of s such that it could be decreasing, but those values would require $\frac{s}{\sigma} > (1-\alpha)(1-\mu)V$, which implies that a buyer would never search even if they could perfectly distinguish high- and low-quality sellers

While their profit if they steer to off-platform transactions is

$$\pi = \begin{cases} \hat{\psi}\mu V - t & \gamma > \bar{\gamma} \\ 0 & \gamma \leq \bar{\gamma} \end{cases}$$

The main difference from Section 5—apart from the cutoffs—is that the steering sellers' profit is scaled down by $\hat{\psi} \geq \psi$. Consumers greater ability to distinguish high- and low-quality sellers gives them greater confidence in sellers product when transacting off platform, which increases their willingness to pay. For $\gamma \leq \bar{\gamma}$, sellers will transact on the platform even if $\xi = 1$, while if that inequality is reversed the cutoff γ at which they are indifferent is

$$\xi \leq 1 - \hat{\psi} + \frac{t}{\mu V} \tag{18}$$

Again, the main difference is that γ increases the appeal of direct sales relative to transacting on the platform. Notably, the logic behind Proposition 2 remains valid. If $\beta = 0$ then the probability of having encountered a high-quality seller conditional on having encountered a seller steering to direct transactions is 0, so $\hat{\psi} = 0$ even if consumers know one seller has deviated to steering.[19] Therefore, for consistency I continue to assume that a proportion $\underline{\beta}$ of sellers exogenously choose to steer consumers to direct transactions. Finally, if $\gamma \leq \bar{\gamma}$ then sellers will not endogenously choose to steer to direct transactions even if $\xi = 1$.

B.3 Platform

Once again similar to Section 5, the platform's profit is given by

$$\Pi = \begin{cases} \rho + (1-\alpha)(1-\underline{\beta})\xi\mu V & \gamma > \gamma^* \\ \rho + \xi\mu V & \gamma \leq \gamma^* \end{cases} \tag{19}$$

The platform will set prices such that consumers are indifferent about participation and sellers are indifferent about steering:

$$\Pi = \begin{cases} \frac{[(1-\alpha)(1-\mu)V - (1-\gamma)\alpha\ell] - \frac{s}{\sigma}}{1-\gamma\alpha} + (1-\alpha)(1-\underline{\beta})\left(1 - \hat{\psi} + \frac{t}{\mu V}\right)\mu V & \gamma > \gamma^* \\ (1-\mu)V - \frac{s}{\sigma(1-\alpha)(1-\beta)} + \mu V & \gamma \leq \gamma^* \end{cases} \tag{20}$$

The platform has no reason to set $\gamma < \bar{\gamma}$ as its profit is constant in consumers' ability to distinguish high- and low-quality sellers once they will never transact outside the platform. Denote the platform profit maximizing level of $\gamma$by γ^*. Taking the derivative of Π with regard to γ gives the rest of Proposition 4:

Proposition B.1. *There exists a critical value of t denoted $\bar{t}$ such that $\gamma^* > \bar{\gamma}$ if $t > \bar{t}$ and $\gamma = \bar{\gamma}$ for $t \leq \bar{t}$ so long as μ is not too small. $\bar{t}$ is increasing in α and $\underline{\beta}$. $\gamma^* < 1$ if μ is not too small and ℓ not too large. γ^* is independent of t if $\gamma > \bar{\gamma}$ and increasing in α.*

Proof. From 20, platform profits are increasing and unbounded in t if $\gamma > \bar{\gamma}$ but constant in t otherwise. Therefore there must exist t sufficiently large such that the platform's profit is greater

[19] This remains true for any generic information structure short of fully informative good news.

with $\gamma > \bar{\gamma}$. On the other hand, from the definition of $\bar{\gamma}$, $\hat{\mathbb{V}} = \mathbb{V}^{all}$ and so $\frac{[(1-\alpha)(1-\mu)V-(1-\gamma)\alpha\ell]-\frac{s}{\sigma}}{1-\gamma\alpha} = (1-\mu)V - \frac{s}{\sigma(1-\alpha)(1-\underline{\beta})}$. Meanwhile $(1-\alpha)(1-\underline{\beta})\left(1-\hat{\psi}\right) < 1$, so for t sufficiently small platform profit is strictly greater with $\gamma = \bar{\gamma}$ so long as μ is sufficiently large that the decrease in seller fees is not compensated for and surmounted by an increase in ρ as γ increases.

For any $\gamma > \bar{\gamma}$, the derivative of platform profit is

$$\alpha\left[\underbrace{\frac{(1-\alpha)(1-\mu)V-(1-\gamma)\alpha\ell-\frac{s}{\sigma}}{(1-\gamma\alpha)^2}+\frac{\alpha\ell}{1-\gamma\alpha}}_{\text{Change in }\rho}-\underbrace{\frac{(1-\alpha)^2(1-\underline{\beta})\underline{\beta}\mu V}{[(1-\alpha)\underline{\beta}+\alpha(1-\gamma)]^2}}_{\text{Change in }\xi}\right]$$

This derivative is negative at $\gamma = 1$ if μ is not too small and ℓ not too large, and it is increasing in α, meaning that any interior maximum must increase in α also. ■

Similar to Section 5, the platform is trading off reduced value extraction from sellers against an increased participation fee from consumers as γ increases. Interestingly, the profit maximizing level of γ does not depend on t so long as the level of t is large enough such that $\gamma > \bar{\gamma}$. This is because a change in γ does not impact the proportion of sellers steering consumers toward direct transactions and so γ provides a purely infra-marginal change in platform profits which is independent of t.[20] Nevertheless the main intuition from Section 5 that the platform will increase the salience of low-quality sellers in response to a decrease in t does survive if that reduction lowers t below the cutoff $\bar{t}$.

C Online Appendix: Benchmark Model with Endogenous Seller Pricing.

In this section I extend the benchmark model to allow for sellers to set prices endogenously. Similar to the benchmark, the model includes a monopoly platform, a continuum of consumers with mass 1 who search on the platform and pay a search cost s for each observation, and a continuum of sellers competing via a Wolinsky (1986) search process with full recall. The platform sets a percentage commission ξ that it takes from the developers' gross revenue. The developers set prices p_j in response to the platform's fee and a marginal cost of production c, and they also have the ability to set a price for off-platform transactions denoted p_j^θ. Consumers on the platform can choose to transact directly with sellers rather than on the platform, but they incur a switching cost t if they do so, reflecting the difficulty of finding the alternative payment, the additional hassle costs of entering new payment information and any foregone convenience benefits of transacting on the platform. The timing of the model is as follows:

1. The platform sets commission ξ
2. Sellers set prices
3. Consumers search and make purchasing decisions

[20] This result likely would not be robust to a more complex model with e.g. a continuum of t. It would also not be true of the model with endogenous pricing as both t and ψ influence sellers' price elasticity with regard to platform policy changes.

C.1 Consumer utility and demand

If consumer i searches and observes seller j receives an idiosyncratic match value draw $\epsilon_{ij} \sim F[\underline{\epsilon}, \infty)$ where F is a differentiable log-concave CDF with corresponding PDF $f(\epsilon)$. If they choose to search again they can recall any previous observation at no cost. Their utility from purchasing seller j's product (gross of search costs) is

$$\epsilon_{ij} - p_j$$

The consumer expects that sellers on the platform are charging the symmetric price p. As is standard in the literature I assume passive beliefs such that observing a deviation does not change consumers' expectations. Given a current observation $U \equiv \epsilon_{ij} - p_j$ If the consumer chooses to search once more paying the search cost s then their expected utility from doing so is

$$\int_{\underline{\epsilon}}^{\infty} \max\{U, \epsilon - p\} d\epsilon - s$$

Setting these two utilities equal and simplifying, i is indifferent about searching if

$$s = \int_{U+p}^{\infty} (\epsilon - p - U) d\epsilon$$

which then implies that j faces demand

$$1 - F(U + p_j) \tag{21}$$

Once aware of a seller, the consumer can choose to transact with them directly and pay the direct transaction price p_j^θ instead of the on-platform price p_j, but they also incur the switching cost t when they do so. Their match value does not change if they engage in this showrooming behavior, so a consumer will choose this direct channel only if

$$p_j - p_j^\theta > t$$

This cutoff is common to all consumers, so all consumers participating on the platform who buy from j will transact directly with the seller or they will all purchase through the platform.

C.2 Seller profit and prices

A seller setting price p_j against expected price p and commission ξ (and choosing to set $p_j - p_j^\theta \leq t$) faces demand given by 21 which yields profit

$$\pi_j = ((1 - \xi)p_j - c)\left[1 - F(U + p_j)\right] \tag{22}$$

For the hetero benchmark model, I focus on symmetric pricing, so taking the first order condition for p_j and then imposing symmetry in on platform prices yields

$$p_j = \frac{c}{1 - \xi} + \frac{1 - F(U + p_j)}{f(U + p_j)} \tag{23}$$

Log-concavity ensures sufficiency of the FOCs for profit maximization, this price gives profit

$$\pi_j = (1-\xi)\frac{(1-F(U+p_j))^2}{f(U+p_j)} \tag{24}$$

The reservation match value $U + p_j$ does not depend on the equilibrium price as long as prices on the platform are symmetric, so this price is increasing in ξ while on-platform profit is decreasing. A seller could instead induce all consumers who purchase from it to transact directly (e.g. by setting p_j arbitrarily high) and transacting with all its consumers via a single pooled price p_j^{steer}. In this case its profit is

$$(p_j^{steer} - c)\left[1 - F(U + t + p_j^{steer})\right]$$

Where t enters on-platform demand because requiring consumers to incur the switching cost is equivalent to raising price by that cost. Taking the FOC

$$p_j^{steer} = c + \frac{1 - F(U + t + p_j^{steer})}{f(U + t + p_j^{steer})} \tag{25}$$

Again, log concavity ensures sufficiency for this to be profit maximization, so the seller's profit from an individual deviation to selling off-platform is

$$\pi_j^{steer} = \frac{(1 - F(U + t + p_j^{steer}))^2}{f(U + t + p_j^{steer})} \tag{26}$$

The RHS of 25 is decreasing in t for any value of p_j^{steer}, which implies that the value of p_j^{steer} which solves 25 is decreasing in t and so it must be the case that $t + p_j^{steer}$ is increasing in t and $1 - F(U + t + p_j^{steer})$ is decreasing in t so π_j^{steer} is also decreasing in t. We can apply this same logic to show **??** and **??** are decreasing in t, and essentially identical logic to show that p_j^{steer} and π_j^{steer} are increasing in p, which in turn is increasing in ξ. Further, given the similarity between 25 and 23, p_j is equivalent to p^{steer} with $t = 0$, which implies $p^{steer} + t > p_j$ for any U. This observation on t completes the proof of Lemma C.1.

Lemma C.1. *p_j is increasing in ξ and π_j is decreasing in ξ. Both are constant in t. p_j^{steer} and π_j^{steer} are both decreasing in t and constant in ξ. For any U $p^{steer} + t > p_j$ when $\xi = 0$.*

Note that these are the prices and profits of an *individual* seller deviating to inducing showrooming by consumers. If proportion β of sellers induce direct transactions then similar to Bar-Isaac, Caruana, and Cuñat (2012) the value of an additional search is

$$s = \beta \int_{U+p^{steer}+t}^{\infty} (\epsilon - p^{steer} - t - U) f(\epsilon) d\epsilon + (1-\beta) \int_{U+p}^{\infty} (\epsilon - p - U) f(\epsilon) d\epsilon \tag{27}$$

The demand and pricing equations are unchanged from an individual deviation, except that from Lemma C.1, for any value of U, $U + p^{steer} + t > U + p$, meaning that the RHS of 27 is decreasing in β, and therefore so is U, meaning that prices and profits for both deviating and non-deviating sellers' profits are increasing in β. Lemma C.2 compares the rate of change for the two seller types

Lemma C.2. *In equilibrium $\beta \in \{0, 1\}$*

Proof. Taking the derivative of 23 with regard to U

$$\frac{dp}{dU} = \frac{-(1-F(U+p))f'(U+p) - f(U+p)^2}{f(U+p)^2}(1+\frac{dp}{dU}) \tag{28}$$

solving

$$\frac{dp}{dU} = \frac{-(1-F(U+p))f'(U+p) - f(U+p)^2}{2f(U+p)^2 + (1-F(U+p))f'(U+p)} \tag{29}$$

Now taking the derivative of 24

$$\frac{d\pi_j}{dU} = (1-\xi)(1+\frac{dp}{dU})\frac{-(1-F(U+p))^2 f'(U+p) - 2(1-F(U+p))f(U+p)^2}{f(U+p)^2}$$

Using 29 and solving

$$\frac{d\pi_j}{dU} = -(1-\xi)(1-F(U+p_j)) \tag{30}$$

Following identical logic

$$\frac{d\pi_j^{steer}}{dU} = -(1-F(U+p_j^{steer}+t)) \tag{31}$$

As Lemma C.1 implies $p_j^{steer} + t > p_j$ for $\xi = 0$, $\frac{d\pi_j}{dU} < \frac{d\pi_j^{steer}}{dU}$, while for $\xi \approx 1$ that inequality is reversed. However, if $(1-\xi)(1-F(U+p_j)) \geq (1-F(U+p_j^{steer}+t))$ then log concavity and Lemma C.1 imply that $\pi_j > \pi_j^{steer}$ because

$$\begin{aligned}
\pi_j =& (1-\xi)\frac{(1-F(U+p_j))^2}{f(U+p_j)} \\
\geq& (1-F(U+p_j^{steer}+t))\frac{(1-F(U+p_j))}{f(U+p_j)} \\
>& \frac{(1-F(U+p_j^{steer}+t))^2}{f(U+p_j^{steer}+t)} = \pi_j^{steer}
\end{aligned}$$

Where the first inequality comes from the assumption that $(1-\xi)(1-F(U+p_j)) \geq (1-F(U+p_j^{steer}+t))$, and the second comes from the fact that $1-\xi < 1$ and so it must be the case that $U+p_j^{steer}+t > U+p_j$ if $(1-\xi)(1-F(U+p_j)) \geq (1-F(U+p_j^{steer}+t))$. From the contra-positive $\pi_j < \pi_j^{steer}$ then implies $(1-\xi)(1-F(U+p_j)) < (1-F(U+p_j^{steer}+t))$.

$\frac{d\pi_j}{d\beta} = \frac{d\pi_j}{dU}\frac{dU}{d\beta}$ and we know $\frac{dU}{d\beta} < 0$, so we can simply reverse the signs on $\frac{d\pi_j}{dU}$ and $\frac{d\pi_j^{steer}}{dU}$ above when considering the impact of β on seller profits. This then implies that $\pi_j - \pi_j^{steer}$ is decreasing in β whenever $\pi_j < \pi_j^{steer}$, so in the equilibrium of the market sub-game $\beta = 0$ or $\beta = 1$. ■

The intuition behind this lemma is that an increase in ξ reduces the benefits that sellers who transact on the platform get from a reduction in search intensity. If ξ is sufficiently large that sellers are willing to leave the platform, then it is also large enough that sellers who steer consumers off platform capture a larger proportion of the gain in surplus from the resulting increase in search intensity than do those who allow on-platform transactions. This pattern reinforces the incentives to steer consumers and so all of the sellers will begin steering consumers as soon as ξ is large enough that one seller would find it profitable to do so. This implies that there is a region of ξ where $\beta = 0$ and $\beta = 1$ are both equilibria of the market sub-game, but I assume that sellers coordinate on selling on-platform whenever this is the case.

Given these assumptions there must be some cutoff where all sellers switch from allowing transactions on the platform to steering consumers off-platform. I formalize this in Proposition C.1

Proposition C.1. *There exists a cutoff commission ξ^* such that no seller will deviate to steering consumers off-platform and $\beta = 0$ for $\xi \leq \xi^*$. While $\beta = 1$ for $\xi > \xi^*$. ξ^* is increasing in t*

Proof. From Lemma C.1, when $\xi = 0$, $\pi_j > \pi_j^{steer}$. Further π_j is decreasing in ξ and approaches 0 as $\xi \to 1$, while π_j^{steer} is constant in ξ for $\beta = 0$. Therefore from the intermediate value theorem there must be some value of ξ such that $\pi_j = \pi_j^{steer}$, and $\pi_j < \pi_j^{steer}$ for ξ above this cutoff. This cutoff is β^*. From Lemma C.2 it then must be the case that $\beta = 1$ for $\xi > \xi^*$. ■

C.3 Platform Profit and Commission

Every consumer will make a purchase eventually, so from Proposition C.1 the platform's profit in the hetero benchmark is given by

$$\Pi(\xi) = \begin{cases} \xi \left[\frac{c}{1-\xi} + \frac{1-F(U+p_j)}{f(U+p_j)} \right] & \xi \leq \xi^* \\ 0 & \xi > \xi^* \end{cases}$$

Profit in the first case is positive and increasing in ξ, therefore the platform will set $\xi = \xi^*$. Meaning that ξ in equilibrium is always the commission that leaves sellers indifferent about allowing on-platform transactions when all other sellers do so. i.e.

$$(1-\xi^*)\frac{(1-F(U+p))^2}{f(U+p)} = \frac{(1-F(U+t+p^{steer}))^2}{f(U+t+p^{steer})}$$

Which means ξ is given implicitly by

$$\xi^* = 1 - \frac{\frac{(1-F(U+t+p^{steer}))^2}{f(U+t+p^{steer})}}{\frac{(1-F(U+p))^2}{f(U+p)}} \tag{32}$$

The fraction on the right hand side is decreasing in t, therefore ξ^* is increasing in t. This proves Proposition C.2

Proposition C.2. *Platform profits are increasing in the switching cost.* $\frac{d\Pi}{dt} > 0$

As will become relevant later, this proposition and the commission in 32 remain valid if some proportion $\underline{\beta}$ of sellers choose to steer consumers to direct transactions for exogenous reasons. An increase in $\underline{\beta}$ will lower the equilibrium commission, but will not qualitatively change any of the results in this section.

Intuitively, proposition C.2 says that any policy (such as the Epic vs. Apple ruling) which reduces the friction for sellers who want to steer consumers off platform will reduce the platform's profit. Therefore the platform has an incentive to increase this switching cost as much as possible. This could take the form of welfare enhancing investments in the quality of transactions on the platform, but may also create an incentive to engage in other conduct that is less consumer friendly.

D Preventing Leakage via Low Quality-Sellers

One of the advantages of transacting through the platform instead of directly is that the platform provides a sort of warranty in that consumers who are dissatisfied with their purchase can seek a refund through the platform. One way that a platform could indirectly raise the switching cost in response to an external force eliminating direct options would be to highlight this guarantee. Apple's warning to consumers as highlighted in Figure 1 is an example of exactly this. Of course, this warning will have little effect when consumers are transacting with large companies that have an established reputation, therefore in order for this threat to be credible the platform would need to relax the intensity with which it screens the quality of sellers on its store. By allowing these apps into its previously strictly walled garden, the platform increases consumers' uncertainty about the trustworthiness of direct transactions and which reduces their appeal despite the lower price and lack of switching cost.

To model this strategy, I suppose the platform allows a proportion α of lower quality apps onto its marketplace who provide 0 utility. The distribution of match values is now

$$\begin{cases} 0 & \text{seller is low quality} \\ \epsilon_{ij} \sim F[\underline{\epsilon}, \infty) & \text{seller is high quality} \end{cases}$$

However consumers cannot observe their true match value draw until after they have made a purchase. Instead they observe a signal that is equal to the value of ϵ they receive if their actual match value is positive, and randomly drawn from $F(\cdot)$ if they encounter a low-quality seller and their true match value is 0. Once a consumer has purchased from a seller and observed a null match value, they have the option to seek a refund of the price they paid. This refund will be successful if they purchased through the platform, but will be denied if they purchased directly. Regardless of the outcome of their request I assume consumers do not search for and purchase a different product after purchasing from an unsatisfactory seller.[21] Given the refund policy, the low-quality sellers only make a positive profit if they make a sale through the direct channel, so they will only offer direct purchases. If no honest seller steers consumers to direct purchases then the low-quality sellers make 0 profit through both channels so they are indifferent between allowing direct purchases and not.

For tractability and ease of exposition, I assume $c = 0$ for the entirety of Section D.

D.1 Consumer Search Behavior and Utility

In this paradigm the interim expected utility of an observed prospect is

$$\epsilon_{ij} - p_j$$

[21] In practice I would expect that some proportion of consumers would search for a replacement product after receiving a refund, or possibly even after failing to receive a refund. Stationarity of the search problem implies that this would increase the value of consumer participation and limit the negative impacts of allowing these low quality sellers on the platform, which would only strengthen the results of this section, so I omit that possibility to aid tractability and exposition.

If purchasing through the platform and

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

If purchasing directly, where $\frac{\beta(1-\alpha)}{\beta(1-\alpha)+\alpha} \equiv \psi$ is the probability that a randomly encountered seller is a normal seller conditional on encountering a seller inducing direct transactions. The expected value of an additional search (given a proportion β) of high-quality sellers inducing direct transactions is

$$(1-\beta)(1-\alpha)\int_{\underline{\epsilon}}^{\infty}(\epsilon-p)f(\epsilon)d\epsilon + (\beta(1-\alpha)+\alpha)\int_{\underline{\epsilon}}^{\infty}[\psi\epsilon - p^{steer} - t]f(\epsilon)d\epsilon - s$$

Therefore, given an expected utility in hand U equalizes the incremental value of an additional search with the search cost if

$$s = (1-\beta)(1-\alpha)\int_{U+p}^{\infty}[(\epsilon-p)-U]f(\epsilon)d\epsilon + (\beta(1-\alpha)+\alpha)\int_{\frac{U+p^{steer}+t}{\psi}}^{\infty}[\psi\epsilon - p^{steer} - t - U]f(\epsilon)d\epsilon \quad (33)$$

which implies that a consumer will stop and purchase from a seller allowing on-platform transactions and charging price p_j if

$$\epsilon_{ij} - p_j > U$$

which leads to demand

$$1 - F(U + p_j)$$

Equivalently, demand for a steering seller is

$$1 - F(\frac{U + p_j^{steer} + t}{\psi}) \quad (34)$$

D.2 Seller pricing and profit

34 and the definition of ψ immediately imply

Proposition D.1. *If $\beta = 0$, no individual high quality seller will deviate to steering consumers to direct transactions for any $\xi < 1$.*

Proof. If $\beta = 0$, then $\psi = 0$, meaning that no consumer will purchase from a steering seller as there is 0 probability that the seller they encounter is high-quality. Given that each seller represents an infinitesimal mass, a single high-quality seller deviating to steering consumers does not change this, so that consumer would acheive 0 demand and hence 0 profit compared to a positive profit on the platform. ■

Sellers allowing transactions on the platform price similarly to the on-platform price in section C, except that U is determined by 33. Therefore their profit is

$$(1-\xi)\frac{(1-F(U+p))^2}{f(U+p)} > 0$$

Proposition D.1 says that no matter how high ξ gets, no individual seller can profitably begin steering consumers to direct transactions because consumers will have no confidence that the

seller is not a low-quality seller, so steering leads to no demand and 0 profit. This result clearly demonstrates the mechanism this model was intended to highlight, but the starkness of the result depends on $\beta = 0$. If consumers believe that a positive mass of high-quality sellers are not allowing on-platform transactions then there exists an interior ξ where sellers can make more profit from selling through the direct channel. Therefore for the rest of this section we assume that a minimum proportion $\underline{\beta}$ of the high quality sellers will steer consumers to the direct channel regardless of the platform's commission.[22]

With this stipulation, profit for sellers who steer consumers off the platform is

$$(p_j^{steer} - c)\left(1 - F(\frac{U + p_j^{steer} + t)}{\psi})\right) \tag{35}$$

This problem is identical for the low- and high-quality sellers. Taking the first order condition (and noting that log-concavity of F ensures sufficiency of the FOC) we get price

$$p^{steer} = c + \psi \frac{1 - F(\frac{U+p_j^{steer}+t}{\psi})}{f(\frac{U+p_j^{steer}+t}{\psi})} \tag{36}$$

and profit

$$\pi^{steer} = \psi \frac{\left(1 - F(\frac{U+p_j^{steer}+t}{\psi})\right)^2}{f(\frac{U+p_j^{steer}+t}{\psi})} \tag{37}$$

This profit is greater than 0, so it remains to demonstrate that there is no interior $\beta \in (\underline{\beta}, 1)$ as established in Lemma D.1

Lemma D.1. *$U + p < \frac{U+p^{steer}+t}{\psi}$; U is decreasing in β, α, and ξ; $\frac{U+p^{steer}+t}{\psi}$ is increasing in t and α; and in equilibrium $\beta \in \{\underline{\beta}, 1\}$*

Proof. First, suppose by contradiction that $U + p > \frac{U+p^{steer}+t}{\psi}$. Then from 36 this would imply

$$U + c + \frac{1 - F(U + p)}{f(u + p)} > \frac{U + c + t}{\psi} + \frac{1 - F(\frac{U+p^{steer}+t}{\psi})}{f(\frac{U+p^{steer}+t}{\psi})}$$

but U, c, and t are all positive, $\psi < 1$, $F(\cdot)$ is log concave, so $U + p > \frac{U+p^{steer}+t}{\psi}$ implies $\frac{1-F(U+p)}{f(u+p)} < \frac{1-F(\frac{U+p^{steer}+t}{\psi})}{f(\frac{U+p^{steer}+t}{\psi})}$ which means that the right hand side of the above inequality must be greater than the left, which contradicts the assumption that $U + p > \frac{U+p^{steer}+t}{\psi}$. Next, from the definition of ψ

$$(\beta(1-\alpha)+\alpha)\int_{\frac{U+p^{steer}+t}{\psi}}^{\infty}[\psi\epsilon - p^{steer} - t - U]f(\epsilon)d\epsilon = \beta(1-\alpha)\int_{\frac{U+p^{steer}+t}{\psi}}^{\infty}[\epsilon - \frac{p^{steer} + t + U}{\psi}]f(\epsilon)d\epsilon \tag{38}$$

which combined with the fact that $U + p < \frac{U+p^{steer}+t}{\psi}$ implies that, if we hold U constant, the

[22] This could be justified for various reasons, strategic trembles, desire to keep sensitive commercial data away from the large platform, or small businesses with taste based aversion to transacting through the platform or a concern that high prices on the platform would harm brand image.

right hand side of 33 is decreasing in β, meaning that U must decrease in β to maintain equality with the search cost. The logic for U decreasing in α is essentially identical. The logic for U decreasing in ξ is essentially similar, but focuses on the first integral in 33 decreasing in ξ as p increases.

Now suppose $\frac{U+p^{steer}+t}{\psi}$ were decreasing or constant in t, this would imply that U is decreasing in t, which means the first integral on the RHS of 33 would be increasing, while the second integral is non-decreasing, but this creates a contradiction as 33 could then no longer be satisfied. Therefore $\frac{U+p^{steer}+t}{\psi}$ must be increasing in t. The logic for the fraction increasing in α is essentially identical.

Turning to the second part of the lemma: by identical logic as in the proof of Lemma C.2, we have that

$$\frac{d\pi_j}{d\beta} = -\frac{dU}{d\beta}(1-\xi)\left(1-F(U+p_j)\right)$$

While from 35

$$\frac{d\pi^{steer}}{d\beta} = \frac{\partial \pi^{steer}}{\partial U}\frac{dU}{d\beta} + \frac{\partial \pi^{steer}}{\partial \beta} + \frac{\partial \pi^{steer}}{\partial p_j^{steer}}\frac{dp_j^{steer}}{d\beta}$$

From the envelope theorem $\frac{\partial \pi^{steer}}{\partial p_j^{steer}} = 0$ so we can treat price as a constant when evaluating this derivative. Taking the derivative of 35 with regard to beta shows $\frac{\partial \pi^{steer}}{\partial \beta} > 0$, while $\frac{\partial \pi^{steer}}{\partial U}\frac{dU}{d\beta}$ which from 35 is equal to $-f(\frac{U+p_j^{steer}+t}{\psi})\frac{dU}{d\beta}(p_j^{steer}-c) = -\frac{dU}{d\beta}\psi\left(1-F(\frac{U+p_j^{steer}+t}{\psi})\right)$. Noting that this implies $\frac{d\pi^{steer}}{d\beta} > -\frac{dU}{d\beta}\psi\left(1-F(\frac{U+p_j^{steer}+t}{\psi})\right)$. Noting that $\frac{dU}{d\beta} < 0$ we can then repeat the rest of the logic from Lemma C.2 to show that $\pi_j < \pi^{steer}$ implies $\frac{d\pi_j}{d\beta} < \frac{d\pi^{steer}}{d\beta}$, which by the same logic as in Lemma C.2 completes the proof. ■

The logic behind the lack of interior β is similar to that in the benchmark, except that the self-reinforcing increased profitability of inducing direct transactions is enhanced by the fact that as more sellers deviate to doing so, ψ increases, which increases consumer confidence in the quality of sellers who induce direct transactions but reduces search intensity. We can then state Proposition D.2

Proposition D.2. *There exists a cutoff ξ^* such that $\beta = \underline{\beta}$ for $\xi \leq \xi^*$ and $\beta = 1$ for $\xi > \xi^*$. ξ^* is increasing in α and t*

Proof. Taking the derivative of π_j with regard to ξ.

$$\begin{aligned}\frac{d\pi_j}{d\xi} =& \frac{\partial \pi_j}{\partial U}\frac{dU}{d\xi} + \frac{\partial \pi_j}{\partial \xi} \\ =& -(1-\xi)(1-F(U+p))\frac{dU}{d\xi} - \frac{(1-F(U+p))^2}{f(U+p)}\end{aligned}$$

Note that the envelope theorem once again implies that we can treat p_j as a constant for this derivative. Similarly

$$\begin{aligned}\frac{d\pi^{steer}}{d\xi} =& \frac{\partial \pi^{steer}}{\partial U}\frac{dU}{d\xi} \\ =& -\psi(1-F(\frac{U+p_j^{steer}+t}{\psi}))\frac{dU}{d\xi}\end{aligned}$$

By the same logic as in Lemma D.1, $pi^{steer} > \pi_j$ implies $-\psi(1 - F(\frac{U+p_j^{steer}+t}{\psi}))\frac{dU}{d\xi} > -(1-\xi)(1-F(U+p))\frac{dU}{d\xi}$, and $-\frac{(1-F(U+p))^2}{f(U+p)} < 0$, so $\pi^{steer} - \pi_j$ is increasing in ξ when it is non-negative, so it can cross from negative to positive at most once. Further this crossing point must exist because from Lemma D.1 $U + p < \frac{U+p_j^{steer}+t}{\psi}$, which implies that $\pi_j > \pi^{steer}$ when $\xi = 0$, while $\pi_j \to 0$ as $\xi \to 1$, while π^{steer} is positive and always increasing in ξ. Define this crossing point as ξ^*.

Similarly

$$\begin{aligned}\frac{d\pi^{steer}}{d\alpha} =& \frac{\partial \pi^{steer}}{\partial U}\frac{dU}{d\alpha} + \frac{\partial \pi^{steer}}{\partial \alpha} \\ =& -\psi(1 - F(\frac{U + p_j^{steer} + t}{\psi}))\frac{dU}{d\alpha} + \frac{\partial \pi^{steer}}{\partial \alpha}\end{aligned}$$

and

$$\begin{aligned}\frac{d\pi_j}{d\alpha} =& \frac{\partial \pi_j}{\partial U}\frac{dU}{d\alpha} \\ =& -\psi(1 - F(\frac{U + p_j^{steer} + t}{\psi}))\frac{dU}{d\alpha}\end{aligned}$$

Which by similar logic as above implies $\pi_j - \pi^{steer}$ is increasing in α for any ξ when it is non-negative. Therefore ξ^* is increasing in α. We can repeat this argument again with t to show that ξ^* is increasing in t.

■

D.3 Platform Commission and Profit

The probability that a consumer purchases through the platform is

$$\frac{(1-\underline{\beta})(1-\alpha)(1-F(U+p))}{(1-\underline{\beta})(1-\alpha)(1-F(U+p)) + (\underline{\beta}(1-\alpha)+\alpha)\left(1-F(\frac{U+p_j^{steer}+t}{\psi})\right)}$$

$$\Pi = \begin{cases} \frac{(1-\underline{\beta})(1-\alpha)(1-F(U+p))}{(1-\underline{\beta})(1-\alpha)(1-F(U+p)) + (\underline{\beta}(1-\alpha)+\alpha)\left(1-F(\frac{U+p_j^{steer}+t}{\psi})\right)}\xi p & \xi \le \xi^* \\ 0 & \xi > \xi^* \end{cases}$$

Let

$$\gamma \equiv \frac{(1-\underline{\beta})(1-\alpha)(1-F(U+p))}{(1-\underline{\beta})(1-\alpha)(1-F(U+p)) + (\underline{\beta}(1-\alpha)+\alpha)\left(1-F(\frac{U+p_j^{steer}+t}{\psi})\right)}$$

p is increasing in ξ, so Π is increasing in ξ for $\xi < \xi^*$, which implies that ξ^* equalizes π_j and π^{steer}, so we can find

$$\xi^* = 1 - \psi\frac{\frac{\left(1-F(\frac{U+p_j^{steer}+t}{\psi})\right)^2}{f(\frac{U+p_j^{steer}+t}{\psi})}}{\frac{(1-F(U+p_j))^2}{f(U+p_j)}} \tag{39}$$

and in equilibrium

$$\Pi = \gamma \xi^* p \tag{40}$$

Proposition D.3. *There is a profit maximizing level of α denoted α^* which is decreasing in t. $\alpha^* > 0$ if t and s are sufficiently small.*

Proposition D.3 says that including the low quality sellers is more profitable the lower the switching cost is. In other words, when a policy change lowers the switching cost, the platform will respond by screening the quality of sellers on its marketplace less intensively in order to reduce the appeal of transacting off the platform.

Proof. ξ^* equalizes π_j and π^{steer}, which from the proof of Lemma D.1 implies $\psi\left(1 - F(\frac{U+p_j^{steer}+t}{\psi})\right) = (1-\xi^*)(1-F(U+p))$ so we can find,

$$\begin{aligned}
\gamma =& \frac{(1-\underline{\beta})(1-\alpha)}{(1-\underline{\beta})(1-\alpha) + (\underline{\beta}(1-\alpha)+\alpha)\frac{1-\xi^*}{\psi}} \\
=& \frac{(1-\underline{\beta})(1-\alpha)}{(1-\underline{\beta})(1-\alpha) + (\underline{\beta}(1-\alpha)+\alpha)^2\frac{1-\xi^*}{\underline{\beta}(1-\alpha)}} \\
=& \frac{(1-\underline{\beta})(1-\alpha)^2}{(1-\underline{\beta})(1-\alpha)^2 + (\underline{\beta}(1-\alpha)+\alpha)^2\frac{1-\xi^*}{\underline{\beta}}}
\end{aligned}$$

Which gives

$$\gamma = \frac{(1-\underline{\beta})(1-\alpha)^2}{(1-\underline{\beta})(1-\alpha)^2 + (\underline{\beta}(1-\alpha)+\alpha)^2\frac{1-\xi^*}{\underline{\beta}}} \tag{41}$$

Which goes to 0 as $\alpha \to 1$ so long as ξ does not approach 1 too quickly (which from 39 will be true so long as $\ln(F(\cdot))$ is not too concave (i.e. the rate of change of the hazard ratio $\frac{1-F(x)}{f(x)}$ is not too large).

$\underline{\beta}$ is a constant, ξ and p are increasing in α (the latter because U is decreasing), so the platform's choice of α trades off increased margin against more consumers transacting off platform. The first order condition for α is

$$0 = \frac{d\gamma}{d\alpha}\xi^* p + \gamma\left(\frac{d\xi^*}{d\alpha}p + \xi^*\frac{dp}{d\alpha}\right) \tag{42}$$

The following identities will thus be useful:

From 23 (recall that the pricing problem for sellers allowing on-platform transactions is qualitatively unchanged from the hetero benchmark)

$$\frac{dp}{d\alpha} = -\frac{(1-F(U+p)f'(U+p) + f(U+p)^2}{(1-F(U+p)f'(U+p) + 2f(U+p)^2}\frac{dU}{d\alpha} \tag{43}$$

Which we can then use to find

$$\frac{dU}{d\alpha} + \frac{dp}{d\alpha} = \frac{f(U+p)^2}{(1-F(U+p)f'(U+p) + 2f(U+p)^2}\frac{dU}{d\alpha} \tag{44}$$

Next, from 37 we can restate ξ^*

$$\xi^* = 1 - \frac{\pi^{steer}}{\frac{(1-F(U+p_j))^2}{f(U+p_j)}} \tag{45}$$

.

ξ^* is the only place p^{steer} appears in Π, so 41, 45, and the envelope theorem together imply that the platform will treat p^{steer} as a constant when setting α.[23]

We can state the platform's maximization problem using a Lagrangian multiplier λ

$$\max_{\xi,\alpha} \gamma\xi p + \lambda \left(\xi - 1 + \psi \frac{\frac{\left(1-F(\frac{U+p_j^{steer}+t}{\psi})\right)^2}{f(\frac{U+p_j^{steer}+t}{\psi})}}{\frac{(1-F(U+p_j))^2}{f(U+p_j)}} \right)$$

Taking the FOC with regard to ξ:

$$\begin{aligned} &\frac{(1-\underline{\beta})^2(1-\alpha)^2\frac{\left(\underline{\beta}(1-\alpha)+\alpha\right)}{\underline{\beta}}}{\left[(1-\underline{\beta})(1-\alpha)^2+(\beta(1-\alpha)+\alpha)^2\frac{1-\xi}{\underline{\beta}}\right]}\xi p + \gamma p + \lambda = 0 \\ &p\left[\frac{(1-\underline{\beta})^2(1-\alpha)^2\frac{\left(\underline{\beta}(1-\alpha)+\alpha\right)}{\underline{\beta}}}{\left[(1-\underline{\beta})(1-\alpha)^2+(\beta(1-\alpha)+\alpha)^2\frac{1-\xi}{\underline{\beta}}\right]}\xi + \gamma\right] = -\lambda \end{aligned} \tag{46}$$

Where the first term in the derivative on the first line comes from 41. Now taking the FOC with regard to α

$$\begin{aligned} 0 =& \frac{-2(1-\underline{\beta})(1-\alpha) - \left[-2(1-\underline{\beta})(1-\alpha) + 2(\underline{\beta}(1-\alpha)+\alpha)(1-\underline{\beta})\frac{1-\xi}{\underline{\beta}}\right]}{\left[(1-\underline{\beta})(1-\alpha)^2+(\underline{\beta}(1-\alpha)+\alpha)^2\frac{1-\xi}{\underline{\beta}}\right]^2}\xi p \\ &-\gamma\xi\frac{f(U+p)^2+(1-F(U+p))f'(U+p)}{2f(U+p)^2+(1-F(U+p))f'(U+p)}\frac{dU}{d\alpha} \\ &+\lambda\Bigg(\frac{\frac{\partial\psi}{\partial\alpha}}{\frac{(1-F(U+p))}{f(U+p)}}\Bigg[\frac{(1-F(\frac{U+p^{steer}+t}{\psi})^2)}{f(\frac{U+p^{steer}+t}{\psi})} \\ &+\frac{U+p^{steer}+t}{\psi}\frac{2(1-F(\frac{U+p^{steer}+t}{\psi}))f(\frac{U+p^{steer}+t}{\psi})^2+(1-F(\frac{U+p^{steer}+t}{\psi}))^2f'(\frac{U+p^{steer}+t}{\psi})}{f(\frac{U+p^{steer}+t}{\psi})^2}\Bigg]\Bigg) \end{aligned}$$

Pulling out common terms and substituting in 46

[23] Deviations by the infinitesimal sellers don't affect search behavior, which combined with Wolinsky search causing each seller to act as a local monopolist means we can apply the envelope theorem despite the fact that sellers are not actual monopolists.

$$\left[\gamma+\frac{(1-\underline{\beta})(1-\alpha)^2\frac{(\underline{\beta}(1-\alpha)+\alpha)^2}{\underline{\beta}}}{\left((1-\underline{\beta})(1-\alpha)^2+(\underline{\beta}(1-\alpha)+\alpha)^2\frac{1-\xi}{\underline{\beta}}\right)^2}\right]\frac{\frac{1-F(\frac{U+p_j^{steer}+t}{\psi})}{f(\frac{U+p_j^{steer}+t}{\psi})}}{1-F(U+p)}$$

$$\times\left(\left(\frac{\underline{\beta}}{\underline{\beta}(1-\alpha)+\alpha}\right)\left[(1-F(\frac{U+p_j^{steer}+t}{\psi}))+\frac{U+p_j^{steer}+t}{\psi}\frac{2f(\frac{U+p_j^{steer}+t}{\psi})+(1-F(\frac{U+p_j^{steer}+t}{\psi})f'(\frac{U+p_j^{steer}+t}{\psi}))}{f(\frac{U+p_j^{steer}+t}{\psi})}\right]\right.$$

$$\left.+\frac{dU}{d\alpha}\left[\frac{2f(\frac{U+p_j^{steer}+t}{\psi})^2+(1-F(\frac{U+p_j^{steer}+t}{\psi}))f'(\frac{U+p_j^{steer}+t}{\psi})}{f(\frac{U+p_j^{steer}+t}{\psi})}-\frac{\psi(1-F(\frac{U+p_j^{steer}+t}{\psi}))\frac{f(U+p)^2+(1-F(U+p))f'(U+p)}{2f(U+p)^2+(1-F(U+p))f'(U+p)}\frac{dU}{d\alpha}}{\frac{1-F(U+p)}{f(U+p)}}\right]\right)$$

$$=\xi\left[\frac{2(1-\underline{\beta})(1-\alpha)+\left[2(\underline{\beta}(1-\alpha)+\alpha)(1-\underline{\beta})\frac{1-\xi}{\underline{\beta}}-2(1-\underline{\beta})(1-\alpha)\right]}{\left[(1-\underline{\beta})(1-\alpha)^2+(\underline{\beta}(1-\alpha)+\alpha)^2\frac{1-\xi}{\underline{\beta}}\right]^2}p-\gamma\frac{f(U+p)^2+(1-F(U+p))f'(U+p)}{2f(U+p)^2+(1-F(U+p))f'(U+p)}\frac{dU}{d\alpha}\right]$$

The derivative is positive when the LHS is greater than the RHS. When $t=0$ and $\alpha=0$, $\xi=0$, meaning that the RHS of this equation is 0, while the LHS is positive. Therefore $\alpha>0$ for t below some cutoff. ■