Optimal discoverability on platforms*

Andrei Hagiu[†] and Julian Wright[‡]

November 2023

Abstract

Choosing how easy to make it for buyers to discover new sellers is a key design decision for platforms. On the one hand, enabling more discoverability generates more transactions and can be more attractive for sellers because they anticipate being discovered by new buyers. On the other hand, discoverability can make sellers more reluctant to participate because they anticipate their existing buyers will discover and purchase from other sellers. We model this fundamental tradeoff and study how a platform's optimal level of discoverability depends on various factors: the degree of substitutability between the sellers' products, the standalone value of the platform's tools, the size of the platform's installed base of buyers, the nature of the platform's fees including its ability to charge differential fees depending on whether a transaction is with a seller's existing buyers or not, the number of sellers, and asymmetries between sellers' sizes.

1 Introduction

A key issue for the design of platforms is how much discoverability to enable. Some platforms are primarily aimed at providing tools for sellers to serve their existing buyers, and offer no or limited ability for buyers to discover new sellers or content providers they did not know about (e.g. Shopify, Substack, Teachable). Others are buyer-focused, and in addition to seller tools, offer search tools and recommendations that make it easier for buyers to discover new sellers or content providers (Amazon, Medium, Udemy).

We study a platform's optimal choice of how much it wants to enable such discoverability. Enabling more discoverability generates a fundamental tradeoff for platforms. On the one hand, it creates more transactions by inducing buyers to purchase from new sellers, thereby

*We would like to thank Chiara Farronato, Shane Greenstein, Jan Krämer, Sam Ransbotham, and Bobby Zhou as well as other participants at Harvard Business School, the Boston University Online Seminar for Digital Businesses and the Platform Strategy Research Symposium for valuable comments and suggestions, as well as David Ronayne and Tat-How Teh for helpful discussions. The paper has also benefited from the insightful feedback of three anonymous referees.

[†]Boston University Questrom School of Business, E-mail: ahagiu@bu.edu

[‡]Department of Economics, National University of Singapore, Singapore, E-mail: jwright@nus.edu.sg

increasing the platform's revenue for any given transaction fee it sets. This can also potentially increase the transaction fees sellers are willing to pay because it allows them to be found by new buyers. On the other hand, it also commoditizes sellers, by making it easier for a seller's otherwise captive buyers to find and purchase from competing sellers. This means some sellers may be reluctant to participate on platforms that enable too much discoverability, potentially decreasing the transaction fees sellers are willing to pay. Indeed, it is this fear of commoditization at the hands of large platforms that has created an opportunity for new platforms to emerge that promise to only enable limited discoverability in order to attract sellers.[1]

To analyze this tradeoff, we build a model in which a platform offers seller tools and attracts sellers, each of whom brings its own initially captive buyers. By its design choices (e.g. how easy it is for buyers to search and compare across the listed sellers), the platform determines what fraction of these buyers also see other sellers that are participating on the platform. The platform also chooses a transaction fee charged to sellers. Taking into account its optimal choice of this transaction fee, we find that the platform's optimal level of discoverability is higher when (i) sellers' products are less substitutable, (ii) the tools the platform offers to sellers are more valuable, (iii) the platform has a larger installed base of buyers, (iv) the number of available sellers that can be brought onto the platform is higher, and (v) sellers are less asymmetric in terms of the number of buyers they bring to the platform. We also show that when the platform can charge differential fees, it would want to charge a lower fee for transactions from a seller's initial captive buyers and a higher fee for transactions from buyers that discover the seller through the platform. In this case, partial discoverability is optimal over a broader range of parameters. And when there are more than two asymmetric sellers, the platform may find it optimal to "drop" the largest sellers and focus instead on the smaller sellers by setting a higher level of discoverability.

We view part of our key contribution as placing platforms on a continuum between B2B input suppliers at one end and marketplaces at the other end, where the continuum is defined by the level of discoverability. To highlight this continuum, it may be useful for readers to have a few examples in mind as they go through our model and results. These examples also help illustrate that platforms do not necessarily want to enable maximum discoverability.

- Both Shopify and Amazon.com attract third-party e-commerce sellers. Both offer tools which allow these sellers to have a web presence. On Amazon.com, sellers are exposed to intense competition amongst themselves (focusing on cases where Amazon does not sell its own first-party products). Buyers can easily compare sellers against each other

[1] For some examples of such platforms, see Hagiu and Wright (2021).

and sort by price. Moreover, Amazon often selects a single seller to recommend to buyers from among the competing sellers via its Buy Box, and sellers compete to win the Buy Box. By contrast, Shopify has been very deliberate in not creating a similar marketplace that enables full discovery for buyers, and in emphasizing to sellers that they can maintain full control over their buyers and that these buyers would not be "shopped around" to other sellers. Its buyer-facing Shop.app (launched in 2020) offers some limited discoverability, but for now it is still mainly focused on enhancing the transactions between sellers and their existing buyers.

- Both Teachable and Udemy connect instructors that offer courses on a wide range of topics with learners. Udemy is essentially a version of the Amazon marketplace for online courses, where learners can browse and discover instructors and courses, complete with a recommendation system based on the learners' interests and courses they have previously taken. By contrast, Teachable solely focuses on providing instructors the tools they need to offer their courses online, without any marketplace enabling discovery of courses for learners. In 2019, Teachable did experiment with Discover, a dedicated sub-domain created for students to browse, preview or enroll in courses from Teachable instructors. Teachable instructors could opt in to appear on the discovery site. Importantly, Teachable went out of its way to make it clear to instructors that Discover was not a full-fledged marketplace (like Udemy and others) which commoditizes instructors.[2] However, since late 2022 the Discover site is no longer active: Teachable seems to have decided to abandon the Discover experiment.

- Both Medium and Substack are platforms connecting independent writers with readers. While it does offer tools for writers to publish their writings online, Medium is very focused on making it easier for readers to discover posts and writers, and it rewards writers when their articles are read by readers. By contrast, Substack is primarily focused on providing writers the tools they need to create and manage newsletters: each author must build their own reader audience, without much help from Substack. Recently, Substack has created a centralized website and app where readers can read and bookmark posts and authors they like, as well as potentially discover new ones. However, most authors still obtain the majority of their audience through their own efforts (e.g. social media, etc.).

There are many other examples, either of platforms making contrasting choices with respect to the amount of discoverability they enable, or of platforms that changed their

[2] See https://teachable.com/blog/discover-by-teachable

positioning over time. For instance, Open Table started off as a pure provider of software tools to restaurants for keeping track of reservations (i.e., a B2B software provider), and only later added the consumer-facing website that enables discoverability, thus becoming a marketplace (i.e., a "proper" platform).

It is important to emphasize that companies that solely offer tools to suppliers and no discoverability are not actually platforms in the standard sense. This is why our analysis can also be interpreted as showing conditions under which a company would want to be closer to a proper platform or closer to a standard business-to-business (B2B) software provider. To keep things as simple as possible, we do not model buyers' decision to join the platform—we assume buyers implicitly affiliate with the platform when the seller they are initially captive to joins the platform. At that point, these buyers are able to use the platform to discover other sellers if the platform enables discovery, so they benefit when more sellers join the platform. And in this case, sellers benefit when more buyers are brought onto the platform by other sellers. This creates an indirect network effects among sellers as soon as there is some discovery on the platform: holding fees and prices constant, each seller benefits when there are more sellers joining the platform, because they bring more buyers that can discover the initial seller.

After fully analyzing the choice of discoverability by a monopoly platform, we also explore what happens when there are competing platforms, showing that competition between symmetric platforms reduces the level of discoverability platforms choose. Moreover, we illustrate how discoverability can be an endogenous way for platforms to differentiate, with one platform offering maximum discoverability and attracting smaller sellers, and the rival platform focusing only on seller tools (i.e. no discoverability) to attract larger sellers.

1.1 Related literature

The paper fits within the burgeoning literature on platform design, with other authors exploring how platforms optimally design consumer search (Hagiu and Jullien, 2011; White, 2013; de Corniere, 2016; Dukes and Liu, 2016; Casner, 2020; Jiang and Zou, 2020; Teh, 2022; and Zhong, 2023), their media content or product recommendations (Barach et al, 2020; Casner and Teh, 2023; Zhou and Zou, 2023), and their reputation system (Shi et al, 2023). Teh (2022) and Choi and Jeon (2023) explore how such optimal design choices are affected by platform business models and how this can lead to misalignment with welfare objectives. Hagiu and Wright (2023) explore how different design choices can be used by platforms to limit leakage or disintermediation. From this literature, our paper is closest to those works focused on search design, and in particular papers showing that the platform may want to add frictions to consumer search in order to relax seller competition, thereby allowing the platform to extract more revenue from sellers (if it charges an ad-valorem or fixed fee). Meanwhile, our focus is on how the extent of discoverability affects sellers' participation incentives given sellers can always sell to their buyers directly without the platform. Indeed, the key distinction relative to the existing literature is that in our setting each seller brings its own captive buyers, so discovery arises from one seller's buyers discovering another seller via the platform, whereas in the existing literature the platform directly attracts buyers which then discover one or more third-party sellers. As a result, and in contrast to the existing literature, the platform in our setting may optimally choose no discoverability at all. This is despite the fact we focus on the platform charging per-transaction fees, which in standard settings imply the platform will choose its design to maximize the volume of transactions, for example, by making search as easy as possible (see Teh, 2022).

Our paper is also related to the four key strategies that can be used to turn product firms into platforms (Hagiu and Altman, 2017). Previous work has only analyzed one of these strategies formally. Specifically, Hagiu et al. (2020) explore the possibility of a multiproduct firm becoming a platform by inviting rivals to sell products or services to the buyers of its core product. The current paper considers one of the other key strategies proposed in Hagiu and Altman, which is "reaching out to customers' customers". Here the firm in question is initially a B2B software provider. Its original customers are sellers who buy its software tools, and their customers are the initial set of buyers they each have. By offering discovery for these buyers, the B2B software provider turns itself into a platform with network effects that helps each seller's buyers discover other sellers.

Our paper is part of an emerging literature that focuses on the downside of participating on platforms from the perspective of sellers. In our paper, the downside is a form of commoditization: sellers bring their buyers to the platform, which then allows those buyers to discover other, possibly competing, sellers. In a sense, by participating on platforms, sellers can lose control over their relationship with their own buyers, something which has been widely discussed in the popular press (see, Hagiu and Wright, 2021 for a discussion). Other related work exploring the downside of participating on platforms include recent work on possible imitation and self-preferencing by hybrid marketplaces like those offered by Amazon and Apple (Anderson and Bedre-Defolie, 2022; Hagiu et al., 2022; and Madsen and Vellodi, 2022). In a similar vein, Mayya and Li (2022) show empirically how participating on food-delivery platforms may commoditize restaurants.

Finally, our paper relates to a strand of the literature on platforms which also models sellers as having some captive buyers and therefore also has the property of dispersed seller pricing in equilibrium. The classic work is Baye and Morgan (2001). The source of captivity is quite different, however. The platform in their setting charges sellers a fixed fee to list, which given sellers are homogeneous, leads to some chance buyers will only be able to buy from the seller they are initially captive to. In our setting, the platform does not want to set a fixed fee (indeed it would subsidize sellers' participation if it could). Closer to our setting, Ronayne (2021) considers the case with transaction fees and assumes, like us, some buyers are captive to sellers for exogenous reasons. However, unlike our setting, there is no intra-platform discoverability: such buyers see all offers on the platform.

2 Baseline model with two sellers

We start with a simple baseline model, which we will later extend in various directions. Suppose there are two symmetric sellers, each of which sells a product that buyers value at v. Both sellers have marginal cost $c < v$. Each seller $i = 1, 2$ starts with a measure one of captive buyers (buyers who only know about the seller). In our baseline model, there are no other buyers, so all buyers (measure two in total) are initially captive to their respective sellers.

The platform offers B2B tools for participating sellers and possibly discoverability, but not any product of its own. We formalize the B2B tools for sellers (which we will call "tools" for brevity) as a reduction in sellers' marginal costs by b. This could be software and other infrastructure that more efficiently handles payments, delivery, customer service, record-keeping and receipts, and so on. If this is the only thing the platform does, the platform can be thought of as just a B2B software-as-a-service company, though with some abuse of terminology, we will still refer to it as a "platform" for convenience.

When both sellers participate, the platform can choose to make a fraction x of buyers aware of both sellers (i.e. make them discover the seller they were not initially aware of), so that the remaining fraction $1 - x$ of buyers are still only aware of the seller they were initially captive to. If $x > 0$, we say that the platform offers discoverability. With probability θ buyers view sellers' products as perfect substitutes, while with probability $1 - \theta$ buyers view them as independent. Of course, this only matters for buyers who are aware of both products, in which case provided both sellers participate, with probability θ they buy only one product and with probability $1 - \theta$ they buy both. We assume the draw of whether a buyer is aware of both sellers is independent of the draw of whether the buyer views sellers' products as substitutes or independent.

Finally, the platform charges each participating seller a non-negative per transaction fee of f, so the effective marginal cost for a seller on the platform is $c + f - b$.

The timing is natural. In period 1 the platform chooses its level of discoverability x and transaction fee f.[3] In period 2, after observing x and f, each seller decides whether to join the platform. Each seller's outside option is to sell to its captive buyers in its direct channel, where it does not benefit from b. Once a seller has decided to participate on the platform in period 2, we assume that it abandons its direct channel, so in period 3 buyers can no longer purchase from it in its direct channel. Then in period 3, each seller sets its price, and demands and payoffs are realized.

For certain choices of f and x, it is possible that there are multiple equilibria in sellers' decisions in period 2, one in which all sellers join given they expect the other seller(s) to join, and one where no sellers join given they expect the other seller(s) not to join. In such cases, we select the equilibrium in which all sellers join, which is sometimes referred to as "favorable beliefs" on the part of sellers.

2.1 Discussion of model assumptions

Some comments about our modelling assumptions are in order. First, our focus on each seller bringing its own buyers distinguishes our framework from standard two-sided platform settings where the platform attracts buyers and sellers separately. Specifically, we do not model buyers' decision to join the platform—instead, we assume buyers implicitly affiliate when the seller they are initially captive to joins the platform. This assumption is made for simplicity. Our model still exhibits indirect network effects between sellers: when joining the platform, each seller benefits from the presence of the other seller via the new buyers that can discover it (provided the platform offers a positive level of discoverability). And in Section 4.1 we explore the case in which the platform also brings some buyers of its own, and discuss how our analysis can easily be extended by allowing buyers to affiliate with the platform based on how much discovery it offers.

Provided each seller brings an equal measure of buyers, the assumption that each seller brings measure one of captive buyers is without loss of generality. We will consider the case in which the two sellers bring different measures of buyers in Section 4.4.1.

Second, to interpret x, one could think of a more elaborate setting in which buyers are heterogeneous in their search cost. There is some cutoff level such that all buyers with search cost below the cutoff discover the rival seller (this is the fraction x) and all those with search cost above the cutoff (i.e. $1 - x$ of buyers) do not search, i.e. they just know their original seller.[4] In this context, one can interpret the platform's choice of x as representing its ability

[3]For a monopoly platform, it is irrelevant whether discoverability x is set before or after (or the same time as) its transaction fee f. The timing of these decisions will matter when we consider competing platforms.

[4]When there are more than two sellers, this same interpretation works if we consider buyers engaging in simultaneous search, so they search all sellers provided their search cost is below some cutoff.

to shift everyone's search cost up or down by its design of the search process. Examples of a platform's design choices that affect x include how prominent they make buyer search, the ability to search based on price or to do side-by-side comparisons, and whether the platform recommends a particular seller to buyers based on price and other factors (e.g. Amazon's Buy Box).

Furthermore, our model is compatible with an alternative interpretation of x, which moves some of the heterogeneity from buyers to sellers. Namely, instead of viewing x as the fraction of buyers aware of both sellers, we could assume that x is the probability with which any given seller is exposed to all buyers. Thus, with probability $1 - x$ each seller remains only exposed to its initially captive buyers. And symmetrically for the other seller. In Online Appendix A.1, we show that this alternative interpretation leads to exactly the same analysis and results as the one provided in Section 3. Regardless of which interpretation of x one prefers, in all cases the platform in our setting only provides information about sellers, rather than trying to persuade buyers to buy from one seller or another (i.e. to change their preferences).

Third, there are two features of the model that are necessary to obtain an interesting tradeoff when choosing discoverability. Specifically, discoverability must result in more transactions in total, but it must also make some transactions contested. Our stark formulation of buyer demand has these properties: with some probability, buyers are interested in both products, so view them as independent, and with the complementary probability they are just interested in one product, so they view competing products as identical. A more realistic but less tractable setting is to have buyers always interested in both products and with downward-sloping demand for each, so that when they are exposed to both products, they buy more in total than if they are just exposed to one. How much more depends on the degree of substitutability buyers perceive between the products. We will show the robustness of our main results to this alternative formulation.

Fourth, we assume sellers each set a single price, so we rule out price discrimination across different types of buyers (i.e. those that are initially captive to the seller, and with some discoverability, those that are initially captive to the other seller(s)). This reflects that sellers may find it difficult to distinguish between buyers. Indeed, buyers could be able to disguise their identity to obtain the more attractive offer in case sellers tried to price discriminate.[5]

Fifth, on timing, we are implicitly assuming that the platform commits to its choice of x and sellers can observe its choice. This captures that it is harder for the platform to change

[5]In Online Appendix A.2, we extend the baseline analysis to the case sellers can effectively price discriminate between their "own" initially captive buyers and those coming from the rival seller after discovery via the platform. This does not change the optimal level of discoverability in a systematic way.

its design choices (e.g. due to technological commitments in designing its search, as well as possibly brand or reputation concerns) than it is for sellers to list (or delist). Without commitment to x, since the sellers' platform participation decisions would be treated as fixed, the platform would always choose maximum discoverability ($x = 1$) given that doing so maximizes the number of transactions facilitated.[6]

The assumption that sellers abandon their direct channel once they decide to participate on the platform is made to simplify the analysis so that we do not need to worry about sellers trying to induce consumers to switch and buy in the direct channel. It can also be justified by viewing each seller's decision to participate on the platform or to develop its direct channel in period 2 as a long-term commitment—so the seller commits to and invests in one channel or the other. Nevertheless, in Online Appendix A.4 we show that the optimal level of discoverability derived below for the baseline model remains unchanged if we allow each seller to maintain its direct channel and buyers to costlessly switch between buying through the platform and buying from each seller's direct channel.[7]

Finally, the adoption of favorable beliefs for equilibrium selection in period 2 has the nice property that it selects the unique equilibrium in period 2 that would arise if instead we had assumed sellers choose sequentially whether to participate on the platform (regardless of the order in which they choose).[8] This equivalence remains true throughout our analysis of Sections 3 and 4, including when there are more than two sellers or when sellers are asymmetric. This reflects that, given a transaction fee and a level of discoverability, the platform always wants to get as many sellers on board as possible, the platform can only set a single transaction fee to do so, and a seller's payoff from not joining doesn't depend on the participation decisions of other sellers.

[6]We have redone our analysis of the baseline setting without commitment in Section A.3 of the Online Appendix.

[7]Ronayne & Taylor (2022) consider a setting where two sellers can sell both through a platform and through their respective direct channels, and some buyers are captive to each of the three possible channels. However, in contrast to our setting, all buyers on the platform, regardless of whether they are captive or not, can always observe both sellers' offerings on the platform, so there is no choice of discoverability by the platform.

[8]In Section A.4 of the Online Appendix we redo the baseline analysis in case sellers hold "unfavorable beliefs," so that they coordinate on the equilibrium in which none of them join whenever that equilibrium exists. As we show there, while the platform's profit is lower in the face of unfavorable beliefs, the baseline characterization of optimal discoverability remains unchanged.

3 Analysis and results

If neither seller joins the platform, then each makes profits

$$v - c. \tag{1}$$

If only one seller joins the platform, that seller's marginal cost is $c - (b - f)$, instead of c for the non-joining seller. Each seller just faces its captive buyers and prices at v. Thus, the joining seller's profit is

$$v + b - f - c$$

while the profit of the non-joining seller is still $v - c$.

Finally, consider the case both sellers join the platform. Given the sellers are symmetric and face equal fees, to determine each seller's expected profit, we just need to determine the measure of captive buyers each seller has after the platform's choice of x. This reflects that given sellers have some fraction of captives and some fraction of buyers who compare the two identical sellers, prices are determined by a mixed strategy pricing equilibrium. In such an equilibrium, each seller's expected profit will equal the profit it can guarantee by just selling to its captives.[9]

Seller i's captive buyers are now made up of seller i's initial captives that did not discover seller j (measure $1 - x$), seller i's captives that discovered seller j but view the two sellers' products as independent (measure $x(1-\theta)$) and seller j's initial captives that discovered seller i but view the two sellers' products as independent (measure $x(1-\theta)$). Thus, each seller's expected profit is

$$(v + b - f - c)(1 - x + 2x(1-\theta)). \tag{2}$$

This is increasing in the extent of discovery x if and only if $\theta < \frac{1}{2}$, i.e. if and only if the two sellers' products are not too substitutable. This makes sense: sellers want to join a platform that induces discovery only if the other participating sellers are not too close substitutes.

The condition for both sellers joining the platform to be an equilibrium is that (2) is no less than (1), or equivalently

$$f \leq b + (v - c)\frac{x(1-2\theta)}{1 + x(1-2\theta)}. \tag{3}$$

[9]This is a special case of the more general result from Proposition 1 in Myatt and Ronayne (2023) which characterizes the expected profits of two or more sellers in a mixed strategy pricing equilibrium allowing sellers to be asymmetric (either in their costs or in their measure of captives). We summarize their more general characterization in Online Appendix A.5.

When $x = 0$, this constraint reduces to $f \leq b$. Without discoverability, there are no interactions between the sellers and no network effects, so the platform just provides tools: each seller adopts it if and only if it offers more value than it charges.

When $x > 0$, if the two sellers' products are not too substitutable ($\theta < \frac{1}{2}$), then the platform can charge $f > b$ and still get both sellers to join given we have assumed sellers coordinate on the equilibrium in which they both join (i.e. they hold "favorable beliefs").[10] In this case, the maximum fee the platform can charge to get both sellers to join is increasing in the amount of discoverability x. On the other hand, when $x > 0$ and the products are sufficiently substitutable ($\theta \geq 1/2$), the platform must charge $f < b$ if it wants both sellers to join. Furthermore, more discoverability now decreases the maximum fee the platform can charge to get both sellers to join.

The platform's demand when it attracts both sellers consists of the $2(1-x)$ buyers who are informed of only one product (and who buy that product only), the $2x(1-\theta)$ buyers who are informed of both products and view them as independent (they buy both), and the $2x\theta$ buyers who are informed of both products and view them as substitutes (they buy one product only). Thus, the platform's profit when both sellers join is

$$\Pi(f, x) = f\left(2(1-x) + 4x(1-\theta) + 2x\theta\right) = 2f\left(1 + x(1-\theta)\right). \tag{4}$$

Clearly, the platform's profit is always increasing in the extent of discovery, holding f and the participation of both sellers fixed. This is natural: discovery expands the number of transactions on the platform.

Substituting in the maximum fee the platform can charge while ensuring the sellers still participate from (3) and defining

$$\mu = \frac{b}{v-c}$$

as the ratio of the value provided by tools to the value provided by the underlying product being sold, we obtain the platform's maximum profit as a function of x only[11]:

$$\Pi(x) = 2\left(\mu + \frac{x(1-2\theta)}{1+x(1-2\theta)}\right)\left(1 + x(1-\theta)\right)(v-c). \tag{5}$$

[10] If the sellers' beliefs were unfavorable, then the platform would face the additional constraint $f \leq b$ (which is binding only if $\theta < \frac{1}{2}$), since otherwise the sellers would coordinate on the equilibrium with neither joining. Similarly, if the sellers maintained their respective direct channels even after joining the platform, and buyers could costlessly switch between purchasing in either channel, then the platform would face the same additional constraint $f \leq b$. Nevertheless, in Online Appendix A.4, we show that the optimal level of discoverability remains the same in these two cases as in Proposition 1 below.

[11] It is straightforward to confirm that the platform always prefers to have both sellers join. Indeed, the platform's profit with one seller joining is half of what it could get with both sellers joining and setting $x = 0$, which is always an option it could choose when it induces both sellers to join.

When products are not too substitutable ($\theta < \frac{1}{2}$), since both the platform and the participating sellers benefit from discovery, Π is increasing in x and the platform will naturally set $x^* = 1$, the maximum amount of discovery. However, when products are more substitutable ($\theta > \frac{1}{2}$), the platform faces a trade-off when choosing the amount of discovery x: on the one hand, a higher x increases the number of transactions, but on the other hand it lowers the participating sellers' profits, so it also lowers the maximum transaction fee f that the platform can extract from the sellers. This can lead to the optimal level of discovery to be set less than one. Relegating the remaining analysis to the appendix, we obtain the following proposition.

Proposition 1. *Suppose each seller starts with a measure one of captive buyers. The platform finds it optimal to induce both sellers to join. The optimal level of discoverability x^* is decreasing in the level of substitutability between products (θ) and increasing in the value offered by the platform's tools for sellers (μ). The platform's optimal level of discovery is given by*

$$x^* = \begin{cases} 1 & if \quad 0 < \theta \leq \theta_1(\mu) \\ \frac{1-\sqrt{\frac{\theta}{(\mu+1)(1-\theta)}}}{2\theta-1} & if \quad \theta_1(\mu) \leq \theta \leq \frac{\mu+1}{\mu+2} \\ 0 & if \quad \theta \geq \frac{\mu+1}{\mu+2} \end{cases}, \tag{6}$$

where $\theta_1(\mu) \in \left(\frac{1}{2}, \frac{\mu+1}{\mu+2}\right)$ is the unique solution in θ to $\frac{\theta}{(1-\theta)^3} = 4(\mu+1)$.

The proposition fully characterizes the platform's optimal choice of discoverability, which is just a function of the underlying parameters θ and μ. A greater level of substitutability between products (i.e. higher θ) induces the platform to choose a lower level of discoverability x^*, because discoverability leads to more intense competition between the sellers, and so makes it harder to attract the two sellers to join the platform. Meanwhile, an increase in the value offered by the platform's tools for sellers (i.e. higher μ) means the platform can charge a higher fee per transaction while keeping sellers willing to participate. This in turn make it more profitable to increase the number of transactions enabled, which the platform does by increasing discoverability. This is why x^* is increasing in μ. In short, the platform's investment in tools and provision of discoverability are mutually reinforcing. Note, however, that even if $\mu = 0$, the platform will set $x = 1$ and derive positive profits if and only if $\theta < \frac{1}{2}$. In other words, provided the sellers' products are not too substitutable, the platform can create positive value for sellers via discoverability and extract positive profits.[12]

[12]More generally, in Online Appendix A.6 we show that if $\theta < \frac{1}{2}$, the platform can derive positive profits even when its tools are worth less than seller tools that are available competitively in the outside market.

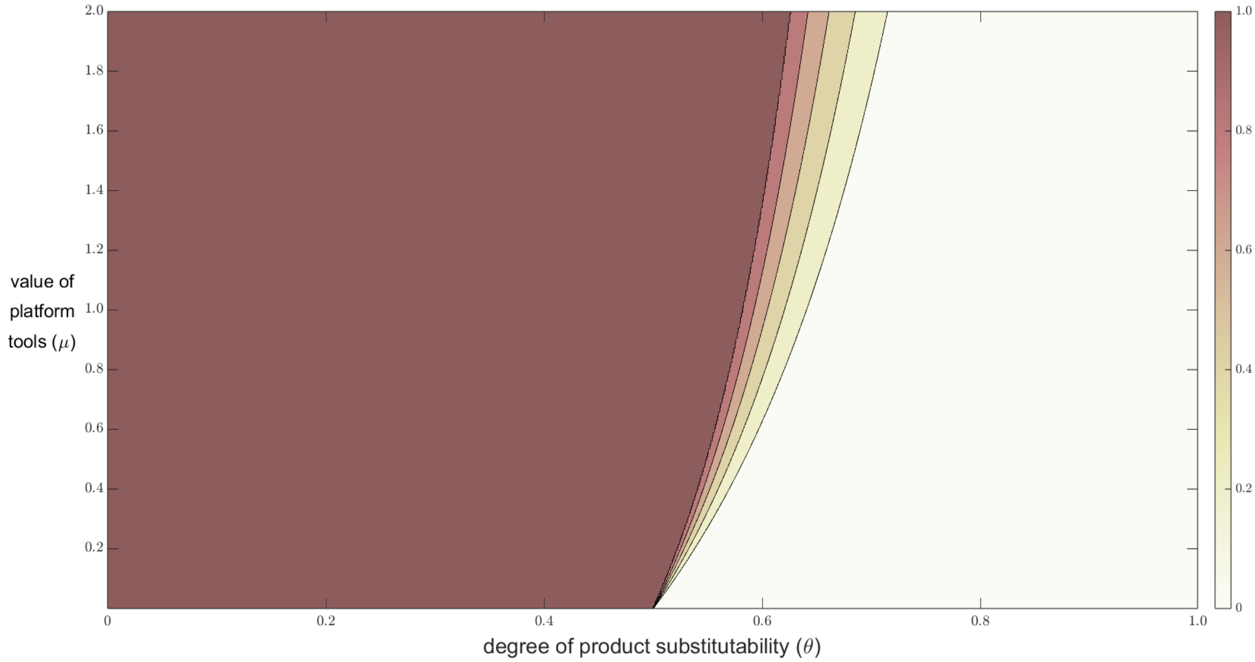

Figure 1: Heatmap of platform's optimal level of discoverability x^*
(darker color indicates higher x^*)

In Figure 1, we have mapped out the optimal x^* when θ is on the horizontal axis and μ is on the vertical axis. The figure shows combinations of the parameters (θ, μ) that support different levels of x^* from $x^* = 0$ (lightest colour) to $x^* = 1$ (darkest colour). The upward sloping relationship seen in the figure reflects that with higher θ, one would require a higher μ to leave the level of x^* unchanged.

One feature of our unit demand setting is that the lower prices resulting from seller competition do not lead to an increase in overall demand. In Online Appendix A.7 we use a less tractable elastic demand setting in which this effect is accounted for, and show that the main comparative static results are very similar

4 Extensions

In this section we explore several interesting extensions of the baseline model.

4.1 Platform brings in buyers

Often platforms attract buyers directly: these could be buyers obtained via the platform's own marketing efforts or they could be buyers that bought other products through the platform in the past (from other sellers or from the platform itself if it acted as a reseller at some point). It is therefore interesting to explore how optimal discoverability changes when the platform attracts buyers directly. To be clear, we will not model the platform competing with the sellers by offering its own first-party products.

Suppose the platform starts with a measure $\eta > 0$ of buyers.[13] We assume the platform's buyers are initially uninformed of the two sellers. By choosing x, the platform also determines the fraction x of the platform's buyers that discover the sellers listed on the platform.

The effect of these platform buyers is to increase each seller's captives by $\eta x (1 - \theta)$, since a fraction $1 - \theta$ of the platform's buyers that become informed of both sellers will buy from both sellers. (The remaining fraction θ of the platform's buyers purchase from the seller with the lower price, and so are not captive to either.) Thus, modifying (3), the condition for each seller to join given that the other does becomes

$$f \leq \left(b + (v - c) \frac{x (1 - 2\theta + \eta (1 - \theta))}{1 + x (1 - 2\theta + \eta (1 - \theta))} \right). \tag{7}$$

The platform's demand when it attracts both sellers is the same as before, plus an additional $\eta (\theta + 2 (1 - \theta)) x$ buyers that come directly from the platform. Thus, the platform's profit when both sellers join is now

$$\Pi (f, x) = (2 (1 + x (1 - \theta)) + \eta x (2 - \theta)) f. \tag{8}$$

Combining (7) and (8), and using the definition of μ, implies the platform's problem is to choose x to maximize

$$\Pi (x) = (2 (1 + x (1 - \theta)) + x\eta (2 - \theta)) \left(\mu + \frac{x ((1 - 2\theta) + \eta (1 - \theta))}{1 + x ((1 - 2\theta) + \eta (1 - \theta))} \right) (v - c). \tag{9}$$

The optimization over x is relegated to the Appendix: the optimal x^* is given by (16). We obtain the following comparative static results.

Proposition 2. *Suppose each seller starts with a measure one of captive buyers and the platform starts with η buyers of its own. The platform finds it optimal to induce both sellers to join. The optimal level of discoverability is decreasing in the level of substitutability between products (θ), increasing in the value offered by the platform's tools for sellers (μ), and increasing in the measure of platform buyers (η).*

Proposition 2 shows that the larger the installed base of buyers the platform starts with (relative to the measure of the sellers' initial captives), the higher the level of discoverability

[13] Given that the measure of each seller's captive buyers is normalized to one, this can also be interpreted as the ratio of platform buyers to each seller's captive buyers.

it will choose. This means the darkest area in Figure 1 (with highest x^*) expands to the right as η increases. Indeed, for large enough η, the platform chooses a positive level of discoverability for the full region of (θ, μ) shown in the figure. The rationale for this result is that informed platform buyers are a net benefit for the sellers, and they create more transactions for the platform. Furthermore, this increases the willingness of sellers to pay to participate (i.e. the fee that the platform can charge), thereby increasing the value of further expanding demand, which the platform does by increasing discoverability. Thus, we expect platforms that already have a lot of buyers coming to them directly, will be more likely to offer maximum discoverability. This also suggests that, over time, as sellers' initially captive buyers keep coming back to the platform to discover potentially new sellers, the platform will want to increase discoverability.

Up to now, we have treated the measure of platform buyers η as a constant. A natural extension is to allow η to increase in the level of discoverability x chosen by the platform: the idea is that more discoverability increases buyers' expected surplus of coming to the platform by increasing the chance such buyers will discover one or more sellers, and by increasing the chance sellers compete for buyers (so lowering their prices in a first-order stochastic dominance sense). It is straightforward to confirm that the platform's profit increases when it has more of its own buyers (i.e. $\partial\Pi(x)/\partial\eta > 0$). Given this, $d\Pi(x)/dx$, which can now be written as $\partial\Pi(x)/\partial x + (\partial\Pi(x)/\partial\eta)(d\eta/dx)$, is increasing in $d\eta/dx > 0$. Thus, when the measure of platform buyers increases with x, it adds an additional incentive for the platform to increase discoverability relative to the case when η was a constant. This implies the optimal x is now weakly higher than the one determined in Proposition 2. This makes sense. An additional benefit of offering discoverability is attracting buyers to come to the platform directly to discover sellers. This further attracts sellers to participate, allowing the platform to charge a higher fee. Both of these effects can already be seen in the expression (9) of $\Pi(x)$ above. Such cross-side network effects are a standard feature of two-sided platforms, and while this shows they are consistent with our framework, they are not required to show our main results.

4.2 Differential fees

Not surprisingly, the platform can do better setting more sophisticated fees than a single per-transaction fee. Without any constraints, it can achieve its first-best profits with a two-part tariff. This involves choosing maximum discoverability ($x^* = 1$) and charging a transaction fee ($f = v - c + b$) so that it extracts the entire margin of each seller's product when sellers both price at buyers' maximum willingness to pay $v + b$. This leaves sellers with zero profit on the platform, so the platform needs to offer them a lump-sum subsidy to match their outside option, which is equal to $v - c$. We formally show this result in Online Appendix A.8. Of course, the problem with such a subsidy is that it must be paid upfront, which is unrealistic since it would lead to moral hazard problems (e.g. sellers participate just to collect the subsidy but then don't do anything to serve buyers) and the platform may also face a budget constraint. In the Online Appendix we show that once the platform is limited in how much of a lump-sum subsidy it can offer, the comparative statics of the platform's optimal discoverability with respect to θ and μ remain the same as in Proposition 1, and moreover, the greater the subsidy the platform can offer, the more discoverability it will want to choose.

The same logic driving the optimal two-part tariff suggests that the platform can also do better by raising the fee it charges for transactions on which sellers compete (to push up seller prices) and lowering the fee it charges for transactions on which sellers do not compete (to keep sellers willing to participate). Indeed, in our model, charging $f_1 = v - c + b$ for transactions generated by buyers who are aware of both sellers but choose only one, and $f_2 = b + (v - c)\frac{1-2\theta}{2(1-\theta)} < f_1$ for all other transactions (such that the sellers are just willing to participate), would replicate the same outcome as with the optimal two-part tariff.

One problem with such a mechanism is that it requires the platform to distinguish between buyers for whom the sellers must compete more intensely from buyers for whom the sellers compete less (either because such buyers are only aware of one seller or because they view the products as independent rather than substitutes). The other problem is that f_2 above is negative when θ is sufficiently close to one. But even if the platform can implement such a mechanism, each seller has no way to verify which of its own initial captive buyers become aware of the rival seller, and so is subject to manipulation by the platform which could overstate the fraction of transactions for which it earns a higher fee.

Taking into account these practical limitations, we focus on a more realistic second-best mechanism which is only based on what the seller can also verify. We allow the platform to charge each seller different transaction fees for selling to the buyers they brought to the platform (their initially captive buyers) vs. for selling to buyers that discovered the seller through the platform (the other seller's initially captive buyers). Since each seller should know which buyers it brings onto the platform, it can in principle monitor the fees it pays are correct. This is indeed a practice that has been used. For example, Teachable, when experimenting with its Discover marketplace, charged instructors a lower fee for students that came via their own Teachable-powered sites and a higher fee for students that came via Teachable's discovery page (3-13% vs. 30%).[14] More broadly, large marketplaces (e.g.

[14]See https://teachable.com/blog/discover-by-teachable

Amazon.com, Etsy) increasingly derive revenues from per-click advertising fees, in addition to transaction fees. This means that a seller will pay a higher fee when transacting with buyers that discover it via ads on the marketplace than when selling to buyers who deliberately search for and come straight to the seller.

Intuitively, charging each seller a lower fee for transactions with buyers they brought to the platform vs. buyers that discovered them through the platform should make sellers more willing to participate and therefore allow the platform to increase the level of discoverability in order to increase the number of transactions enabled. As we will see, this intuition does not always hold.

To proceed, we allow the platform to charge each seller a transaction fee f_0 for transactions with the seller's initial captive buyers and a potentially different transaction fee f_1 for transactions with buyers that are not part of the seller's initial captive base.

Suppose both sellers join the platform (the payoffs from not joining are the same as in the baseline). The set of captive buyers for a seller is made up of three components:

- $1-x$ buyers on whom the seller incurs a marginal cost of $c+f_0-b$ and who only consider that seller
- $x(1-\theta)$ buyers on whom the seller incurs a marginal cost of $c+f_0-b$ and who consider both sellers
- $x(1-\theta)$ buyers on whom the seller incurs a marginal cost of $c+f_1-b$ and who consider both sellers.

Each seller's profit is then

$$(v+b-f_0-c)(1-x+x(1-\theta))+(v+b-f_1-c)\,x(1-\theta).$$

Indeed, the two sellers are symmetric, so each seller's expected profit from setting any other price in the support of its mixed strategy would have to be the same as what it can obtain simply by serving its captive buyers (which here incur different marginal costs).

The sellers' profits are increasing in x for given fees if and only if

$$\theta<\frac{v+b-c-f_1}{(v+b-c-f_0)+(v+b-c-f_1)}.$$

Compare this with the baseline, where sellers' profits were increasing in x if and only if $\theta<\frac{1}{2}$. The platform's ability to charge different fees makes it less likely for discoverability to be good for sellers' profits whenever $f_0<f_1$. The reason is that when $f_0<f_1$, each

seller makes a higher margin on its own initially captive buyers (for whom it prefers less discoverability) vs. on buyers that discovered it through the platform (for whom it prefers more discoverability), so overall each seller prefers less discoverability.

To determine the platform's revenue, consider the measure one of buyers that are initially captive to seller i. Out of these buyers, $1-x$ remain captive to seller i and buy from that seller only, so the platform makes $f_0(1-x)$ on them. Another fraction $x(1-\theta)$ are informed of both products and view them as independent, so they buy both and the platform makes $(f_0+f_1)x(1-\theta)$ on them. And the remaining fraction $x\theta$ are informed of both products and view them as substitutes, so they buy one product only. Given that the sellers are symmetric and therefore have the same price distributions in equilibrium, half of these buyers will buy from seller i and half will buy from seller j, so the platform makes $\frac{(f_0+f_1)x\theta}{2}$ on these buyers. Thus, in total, the platform's profit when both sellers join is

$$\begin{aligned} & 2f_0(1-x)+2(f_0+f_1)x(1-\theta)+(f_0+f_1)x\theta \\ = \; & f_0(2-x\theta)+f_1x(2-\theta). \end{aligned}$$

The platform's problem is to set x, f_0 and f_1 to maximize the above expression, subject to the following three constraints:

$$\begin{aligned} 0 &\leq f_0 \leq v+b-c \\ 0 &\leq f_1 \leq v+b-c \end{aligned}$$

$$(v+b-f_0-c)(1-x+x(1-\theta))+(v+b-f_1-c)x(1-\theta) \geq v-c.$$

The first two constraints rule out negative transaction fees[15] and ensure that buyers want to participate at the competitive price. The third constraint ensures that each seller wants to participate on the platform.

The optimization over (f_0, f_1, x) is relegated to the Appendix: the optimal x^* is given by (17)-(18). We obtain the following comparative static results.

Proposition 3. *Suppose the platform can charge each seller a fee f_0 for transactions with its initially captive buyers and f_1 for transactions with buyers it gains through discovery on the platform. Then the platform finds it optimal to induce both sellers to join and to set $f_1 > f_0$. The optimal level of discoverability is decreasing in the level of substitutability between products (θ) and increasing in the value offered by the platform's tools for sellers (μ).*

[15] Indeed, negative transaction fees are seldom used in practice because they create arbitrage-type problems (e.g. some sellers might join just to buy from themselves and thereby collect the subsidy).

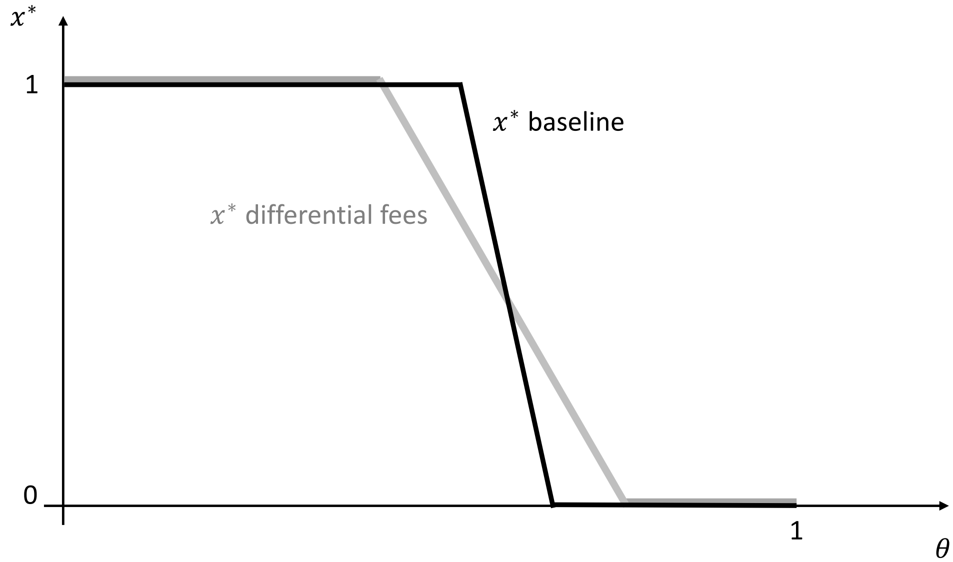

Figure 2: The platform's optimal level of discoverability x^*

In terms of fees, the key result is that the platform always finds it optimal to charge each seller a higher fee for transactions with buyers that are not part of the seller's initial captive base (f_1) than for transactions with the seller's initial captive buyers (f_0). The reason for this is that provided $0 < x < 1$, a larger share of a seller's transactions that come from the rival seller's buyers involve head-to-head competition (and therefore do not contribute to the seller's expected profit), relative to the seller's transactions that come from its own initially captive buyers. Indeed, the share of "discovery transactions" (i.e. transactions generated by the rival seller's buyers) that result in head-to-head competition is $\frac{x\theta}{x-\frac{x\theta}{2}}$, whereas the share of transactions with a seller's own buyers that result in head-to head competition is $\frac{x\theta}{1-\frac{x\theta}{2}} < \frac{x\theta}{x-\frac{x\theta}{2}}$. As a result, the platform prefers to set a higher fee for discovery transactions, because it is more likely to get passed through by the sellers.

It is important to emphasize that this differential fee strategy only works if $0 < x < 1$.[16] For this reason, the range over which partial discoverability is optimal (i.e. $0 < x^* < 1$) is now larger than in the baseline model where the platform could only charge a single fee. The platform prefers partial discoverability because it allows it to exploit this profitable differential fee strategy. This is illustrated in Figure 2, which is constructed with $\mu = 1$: the black line represents $x^*(\theta)$ in the baseline and the gray line represents $x^*(\theta)$ with differential fees.

This also leads to the following corollary, which compares the optimal level of discover-

[16]If $x = 0$ or $x = 1$, then the platform does not gain anything from charging differential fees. Indeed, if $x = 0$, then there is no discovery so f_1 is irrelevant, whereas if $x = 1$, then all buyers are equivalent, so only $f_0 + f_1$ matters.

ability here to the one from the baseline.[17]

Corollary 1. *Denote the optimal level of discoverability in the baseline by x_b^*, given by (6), and the optimal level of discoverability with differential fees by x_{df}^*, given by (17) when $\mu < 1$ and by (18) when $\mu > 1$. For every $\mu > 0$, there exists a unique $\theta_3 \in \left[\theta_1(\mu), \frac{\mu+1}{\mu+2}\right]$ such that $x_{df}^* \leq x_b^*$ if $\theta \leq \theta_3$ and $x_{df}^* \geq x_b^*$ if $\theta \geq \theta_3$, where $\theta_1(\mu)$ is defined in Proposition 1.*

Corollary 1 implies that being able to set different fees leads to less discoverability when θ is less than some threshold (denoted θ_3 in the Corollary) and leads to more discoverability when θ is more than that threshold. Moreover, the threshold always arises in the range where there is partial discoverability in the baseline, as can be seen from Figure 2.

4.3 More than two sellers

So far we have focused on the case with only two sellers. Suppose now there are $n \geq 2$ sellers: as in the two-seller case, each seller brings measure one of buyers who are informed of the particular seller they are captive to, but are not informed of any of the other sellers.[18]

Things are very similar to before, except now, a seller's expected profit from joining when it expects some number $m - 1 \leq n - 1$ of other sellers to join (where $2 \leq m \leq n$) is

$$(1 - x + mx(1-\theta))(v + b - f - c). \tag{10}$$

To understand this, note first that $1 - x$ of a seller's initial captives do not discover other sellers, so remain captive. The remaining fraction x discover all other $m - 1$ sellers, with a fraction $1 - \theta$ of these viewing all sellers' products as independent, so $x(1-\theta)$ also remain captive from each seller's perspective. Finally, each seller also sells to the $x(1-\theta)$ buyers it gets exposed to from each of the other $m - 1$ sellers' initial captives. So each of the m sellers who join ends up with

$$1 - x + x(1-\theta) + (m-1)x(1-\theta)$$

captives, thus leading to the result in (10).

Comparing (10) with the payoff $v - c$ from not joining, a seller will want to join the

[17] The proof involves a direct comparison of the two solutions for x^* and is provided in Online Appendix A.9 for completeness.

[18] Even if the measure of buyers who come with each seller decreased as more sellers are added, our results would remain unchanged. This reflects that the optimal level of discoverability does not depend on the measure of buyers each seller brings as long as it remains equal across sellers.

platform when it expects $m-1$ other sellers to do so if and only if

$$f \leq b + \frac{x\left(m\left(1-\theta\right)-1\right)}{x\left(m\left(1-\theta\right)-1\right)+1}\left(v-c\right). \tag{11}$$

As is clear from (10), there are positive network effects across sellers. The more sellers join, the higher the payoff from joining for each seller (and therefore the higher f the platform can charge). Note that we continue to assume favorable beliefs, in that sellers always coordinate on the highest number of available sellers joining that is an equilibrium, given the fee f charged by the platform.[19]

Suppose the platform attracts m sellers in total. Each of these sellers has $1-x$ captive buyers who are only informed of one product and buy that product only, so the platform demand generated by these buyers is $m\left(1-x\right)$. Meanwhile, a total measure of mx buyers discover all sellers on the platform. Out of these, a fraction $1-\theta$ view all products as independent so buy all of them, while the remaining fraction only buy one product. The platform demand generated by these informed buyers is $mx\left(\left(1-\theta\right)m+\theta\right)$. Total demand for the platform when m sellers join is thus

$$m\left(1+x\left(m-1\right)\left(1-\theta\right)\right). \tag{12}$$

Since (11) and (12) are both increasing in m, the platform obtains its maximum payoff by inducing all sellers to join ($m=n$) and setting f so (11) is binding when $m=n$. The resulting platform profit is

$$\Pi\left(x\right) = \left(b + \frac{x\left(n\left(1-\theta\right)-1\right)}{x\left(n\left(1-\theta\right)-1\right)+1}\left(v-c\right)\right)\left(n\left(1+x\left(n-1\right)\left(1-\theta\right)\right)\right). \tag{13}$$

The optimization over x is relegated to the Appendix: the optimal x^* is given by (19). We obtain the following comparative static results.

Proposition 4. *Suppose each of n sellers starts with a measure one of captive buyers. The platform finds it optimal to induce all sellers to join. The optimal level of discoverability is decreasing in the level of substitutability between products (θ), increasing in the value offered by the platform's tools for sellers (μ), and increasing in the number of available sellers (n).*

[19] Alternatively, we could assume sellers decide sequentially whether or not to join the platform. In Online Appendix A.10, we explore less favorable beliefs. These lower the platform's profit, as it has to set a lower fee to attract all sellers to join, but as was the case for the baseline setting with two sellers, such beliefs have no effect on the platform's optimal choice of x^*.

The new result is that additional sellers always increase the amount of discoverability the platform chooses. In part this reflects that our model over-emphasizes the positive network effect across sellers due to discoverability and de-emphasizes the negative substitution effect that can arise as more and more sellers are added. Indeed, the only thing that matters for a seller's expected profit is the profit from captives, which always increases when more sellers join the platform. Meanwhile, the number of sellers that compete for contested buyers turns out not to affect a given seller's equilibrium profit. With a more general demand function, each seller's profits from contested buyers would be decreasing in the number of participating sellers, so that adding more sellers can lower each seller's equilibrium profit. In Online Appendix A.11 we use a less tractable setting with elastic demand, and show that adding more sellers can reduce the optimal level of discoverability.

One way to interpret these different results is that our baseline demand specification captures that each additional seller serves a unique product category, with buyers sometimes only wanting to buy from one such product category (and viewing them as perfect substitutes), and other times wanting to buy from all of them. The result says the platform should increase the level of discoverability as it adds more product categories. In contrast, the alternative elastic-demand specification captures the idea of adding more sellers within a given product category. In that case, we find the platform should decrease the level of discoverability as it adds more sellers within a given product category.

Moreover, it is important to recall, in our model discoverability involves buyers seeing all listed sellers. There are many possible alternative formulations. For instance, discoverability could involve buyers seeing a fixed number of sellers, say $j \geq 2$ out of a total n participating sellers, where $j < n$. In this case, $1 - x$ of a seller's initially captive buyers do not get to see any other seller, and x of them get to observe $j - 1$ other randomly selected sellers. In Online Appendix A.12 we show that the optimal level of discoverability in this case is the same as it is for case above in which there are j sellers on the platform to start with. Thus, for instance, if buyers only look at most at two sellers, then the optimal level of discoverability is the same as in the baseline setting.

4.4 Heterogeneous sellers

So far we have assumed all sellers are identical, each starting with the same measure one of captive buyers. In this section we analyze two different cases where the sellers are not symmetric: in the first case we explore how asymmetry changes the optimal level of discoverability, and in the second case we illustrate the possibility that a platform may choose to only attract smaller sellers and leave larger sellers out by setting a high level of discoverability.

4.4.1 Two asymmetric sellers

Consider first the case with two sellers, where seller i has measure λ_i of initially captive buyers, and assume $\lambda_1 \geq \lambda_2$.

If seller i does not join the platform, then its profit is $\lambda_i (v - c)$. If only seller i joins the platform, its profit is $\lambda_i (v + b - f - c)$, while the profit of the non-joining seller is still $\lambda_j (v - c)$. If both sellers join the platform, then seller i and seller j will compete with different measures of captive buyers. The analysis in this case turns out to be more complicated, given that the seller with fewer captives will act more aggressively and its profit will be higher than what it can obtain by just charging the monopoly price on its captives. Despite this, as we show in the proof of Proposition 5 below, it is still the seller with more captives (seller 1) that turns out to constrain the fee the platform can set to induce the two sellers to participate in case $x > 0$. This is intuitive: that seller has a better outside option, and discoverability brings more of its buyers to the other seller, than vice-versa.

The captive buyers for seller 1 are now made up of seller 1's initial captives that did not discover seller 2 (measure $\lambda_1 (1 - x)$), seller 1's captives that discovered seller 2 but view the two sellers' products as independent (measure $\lambda_1 x (1 - \theta)$) and seller 2's initial captives that discovered seller 1 but view the two sellers' products as independent (measure $\lambda_2 x (1 - \theta)$). Thus, seller 1's profit is

$$(v + b - f - c) (\lambda_1 + \lambda_2) (\beta_1 (1 - x) + x (1 - \theta)),$$

where

$$\beta_1 = \frac{\lambda_1}{\lambda_1 + \lambda_2} \geq \frac{1}{2}$$

is seller 1's relative market share of initial captives. Comparing this with seller 1's profit if it doesn't join, the platform must set

$$f \leq b + (v - c) \left(1 - \frac{\beta_1}{\beta_1 (1 - x) + x (1 - \theta)} \right), \tag{14}$$

to ensure seller 1 participates, which as shown in the proof of Proposition 5 below, also ensures seller 2 participates.

The platform's demand when it attracts both sellers consists of the $(\lambda_i + \lambda_j)(1 - x)$ buyers who are informed of only one product (and so only buy that product), the $(\lambda_i + \lambda_j) x (1 - \theta)$ buyers who are informed of both products and view them as independent (they buy both), and the $(\lambda_i + \lambda_j) x\theta$ buyers who are informed of both products and view them as substitutes (they buy one product only). Thus, the platform's profit when both sellers join is

$$f\left(1+x\left(1-\theta\right)\right)\left(\lambda_1+\lambda_2\right).$$

Note that the platform can set $x=0$ and $f=b$ to obtain $b\left(\lambda_1+\lambda_2\right)$, which is strictly higher than $b\lambda_1$, the maximum profit it can achieve by attracting one seller only. Thus, it is optimal for the platform to attract both sellers.

The platform will therefore set f and x to maximize the last expression above subject to (14), which ensures both sellers participate. The optimization over f and x is relegated to the Appendix: the optimal x^* is given by (20). We obtain the following comparative static results.

Proposition 5. *Suppose seller i starts with a measure λ_i of captive buyers, where $\lambda_1 \geq \lambda_2$. The platform finds it optimal to induce both sellers to join. The optimal level of discoverability is decreasing in the level of substitutability between products (θ), increasing in the value offered by the platform's tools for sellers (μ), and decreasing in the extent of asymmetry between sellers (β_1).*

The new result is that the optimal level of discoverability is decreasing in β_1, the larger seller's market share of initial captives. Thus, the bigger the difference in initial market shares of captives, the less discoverability the platform will provide. The reason is that the binding participation constraint that the platform's fee and level of discoverability must respect is that of the larger seller. And the larger seller necessarily prefers less discoverability since it brings more buyers to the platform than it stands to gain from discoverability.

4.4.2 Why a large seller may not participate on the platform

As shown in Proposition 5, with two sellers, even if asymmetric, it is always profitable for the platform to attract both of them. With more than two sellers, if they are symmetric, the platform also wants to attract all of them (as we saw in Section 4.3). However, with more than two sellers, if they are heterogeneous, the platform may be better off setting its transaction fee and level of discoverability such that not all sellers participate.

In particular, the previous analysis with asymmetric sellers shows that it is the larger seller (in terms of their initial captives) that constrains the platform's transaction fee because it benefits less from joining the platform. This is consistent with real world observations: larger and more established brands are the ones least likely to participate on large marketplaces (e.g. Amazon.com), preferring to sell through their own channels instead.[20]

[20] See for example https://www.cnbc.com/2019/11/13/nike-wont-sell-directly-to-amazon-anymore.html

In what follows we confirm that in a setting with three sellers it may be optimal for the platform to set its fee and level of discoverability such that the larger seller does not participate in equilibrium. Suppose seller $i \in \{1,2,3\}$ has λ_i initial captives and assume $\lambda_1 > \lambda_2 = \lambda_3$. Denote

$$\beta = \frac{\lambda_1}{\lambda_1 + 2\lambda_2}.$$

First, it can never be optimal for the platform to induce only one seller to join because that implies no discovery, so the most the platform could obtain is $b\lambda_1$. The platform could do strictly better setting $x = 0$ and the same $f = b$, so all sellers are willing to join, yielding $b(\lambda_1 + 2\lambda_2)$ for the platform. Second, it can never be optimal to induce the large seller to join together with only one small seller. We prove this result as part of Proposition 6 below. The reason is essentially the same as above. The large seller is the least likely to wish to participate on the platform when other sellers are present, and given Bertrand competition for buyers who view the products as substitutes, having two small sellers join is actually better for the large seller than having just one small seller due to the possibility of discovery. So if the large seller participates, then the second small seller is even more willing to participate, and the platform certainly benefits from having three rather than two sellers via an increased number of transactions.

Taking these two results into account, the platform's optimal strategy is either to induce all three sellers to join, or to only induce the two small sellers to join. If the platform induces all three sellers to join, the binding constraint on the platform's optimal fee is once again the participation of the large seller (we show this in the proof of Proposition 6 below), so the platform's profits in this case are

$$\begin{aligned} &\max_x \left\{ f(\lambda_1 + 2\lambda_2)(1 + x(1-\theta)) \right\} \\ \text{subject to } \lambda_1 (v - c) \quad &\leq \quad (v + b - f - c)(\lambda_1(1-x) + (\lambda_1 + 2\lambda_2)x(1-\theta)), \end{aligned}$$

which is equal to

$$\max_x \left\{ (\lambda_1 + 2\lambda_2)(v-c)\left(\mu + 1 - \frac{1}{1 - x + \frac{x(1-\theta)}{\beta}}\right)(1 + x(1-\theta)) \right\}.$$

Meanwhile, if the platform only induces the two small sellers to join, then the analysis is the same as in the baseline model, so the platform's profits in this case are

$$\max_x \left\{ 2\lambda_2 (v-c)\left(\mu + 1 - \frac{1}{1 + x(1-2\theta)}\right)(1 + x(1-\theta)) \right\}.$$

Consider the tradeoff between these two options. The total number of buyers is larger when attracting all three sellers ($\lambda_1 + 2\lambda_2$ instead of $2\lambda_2$), but the transaction fee can be higher when attracting just the two small sellers:

$$(v-c)\left(\mu+1-\frac{1}{1+x\left(1-2\theta\right)}\right) > (v-c)\left(\mu+1-\frac{\lambda_1}{1-x+\frac{x(1-\theta)}{\beta}}\right)$$

which is true if and only if

$$\beta > \frac{1}{2}.$$

Thus, the large seller has to be at least as large as the two small sellers combined in order for the maximum transaction fee that can be charged with two small sellers to be higher than that charged with all three sellers. In this case, the size disparity between the large seller and the two small sellers is so big that for the same level of discoverability, the platform must charge a lower transaction fee if it wants to attract all three sellers than when it wants to attract only the two small sellers.

By contrast, if the large seller is close in size to each of the small sellers ($\frac{1}{2} < \beta \leq 1$), then there is no tradeoff and the platform always prefers to induce all three sellers to join, consistent with the results from Section 4.3 with multiple sellers: profits are increasing in the number of (equal) sellers that join.

The following proposition confirms this by focusing on the specific case when tools have no value ($\mu = 0$), so the only valuable service that the platform can provide is discovery.

Proposition 6. *Suppose there are three sellers, one large of size λ_1 and two identical smaller sellers, each of size $\lambda_2 < \lambda_1$. Suppose also $\mu = 0$. When $\theta \geq \frac{1}{2}$, the platform strictly prefers to induce all three sellers to join if $\beta < 1-\theta$ and is indifferent between two or three sellers joining (with zero resulting profits) when $\beta \geq 1-\theta$. When $\theta < \frac{1}{2}$, the platform prefers to induce all three sellers to join if $\beta \leq \frac{1}{1+2\theta}$ and prefers to induce only the two small sellers to join if $\beta > \frac{1}{1+2\theta}$. If it is optimal to induce only the two small sellers to join, the optimal level of discoverability is as in the baseline. If it is optimal to induce all three sellers to join, the optimal level of discoverability is higher than in the baseline, strictly so if $\frac{1}{2} \leq \theta < 1-\beta$.*

The proposition shows that when $\mu = 0$, for any θ, there exists a threshold such that the platform prefers to induce all three sellers to join when β is below that threshold and prefers (strictly only if $\theta < \frac{1}{2}$) to induce only the two small sellers to join when β is above that threshold. This is an artifact of the assumption that $\mu = 0$, so the platform has no valuable tools to offer aside from discovery. Indeed, this implies that when the large seller becomes sufficiently large relative to the two small sellers (i.e. β becomes large), the platform prefers to drop the large seller because attracting it means choosing almost no discovery and therefore vanishingly small profits in the absence of valuable tools.

In general however, with $\mu > 0$ so the platform offers valuable tools, if the large seller becomes sufficiently big relative to the small sellers, then the platform once again strictly prefers inducing all three sellers to join (which it can always do by setting x equal or close to zero), for the simple reason that the large seller is too big to leave out and it can be served profitably with tools. For the platform to prefer inducing only the two small sellers to join, λ_1 has to be in some intermediate range relative to λ_2 (given θ). This is confirmed in Figure 3, which shows the platform's optimal choice of sellers (either all three sellers, or the two small sellers only) as a function of θ and β when $\mu = 0.1$.

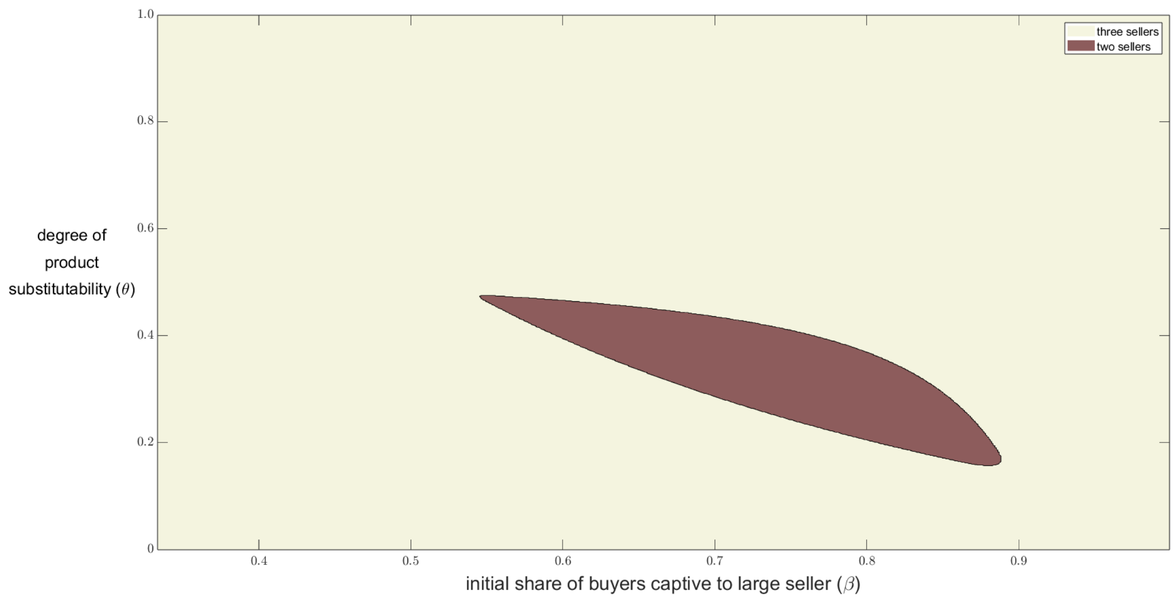

Figure 3: Parameter region where two sellers join and where three sellers join

As μ increases, the region in Figure 3 where the platform prefers to only induce the two sellers to join shrinks, and we note that for any $\mu \geq 0.2$, there is no θ and β for which the platform ever prefers only selling to the two small sellers.

4.5 Competing platforms

So far we have assumed there is a single monopoly platform. In this subsection we consider two extensions to explore the implications of platform competition. The first provides a direct extension of our baseline setting to the case with two identical competing platforms and n sellers. The second explores a setting in which one platform has the advantage of its own set of buyers (e.g. Amazon) and the other has the advantage of providing superior seller tools (e.g. Shopify). We show the possibility of sustaining an equilibrium where the platforms endogenously differentiate themselves by offering different levels of discoverability, and thereby attracting different sellers.

4.5.1 Symmetric platform competition

Suppose there are two identical platforms, platforms 1 and 2, and n symmetric sellers, each of which has measure one of captive buyers. We assume that when platform i offers a positive level of discoverability, i.e. $x_i > 0$, it incurs some small fixed cost $\varepsilon > 0$. Consistent with this assumption, we model the platforms' decisions as being made sequentially, with platform 1 choosing x_1 first and platform 2 choosing x_2 second after observing x_1. After observing each other's choices of x_i, they simultaneously choose fees f_1 and f_2 in period 2.

Assuming $\varepsilon > 0$ ensures the platform that offers a positive level of discoverability in equilibrium can earn positive profits from doing so. It is also justified by noting that in practice, discoverability requires investments in search and matching features. The assumption that platform 1 chooses x_1 before platform 2 chooses x_2 is made without loss of generality. Moreover, even if the two platforms choose their x_1 and x_2 simultaneously, the same equilibria characterized below would still hold.[21]

From (10), the payoff to a seller that joins platform i alongside $m-1$ other sellers (where $2 \leq m \leq n$) is

$$(1 - x + m x_i (1 - \theta)) (v + b - f_i - c) .$$

This payoff is increasing in m. Meanwhile, the payoff to a seller that joins platform i alone is

$$v + b - f_i - c.$$

There are therefore two possibilities. If $\theta < 1 - \frac{1}{n}$, then the maximum payoff that each seller can obtain (given platforms' fees and levels of discovery) is when they all join the same platform, which means this is the only possible outcome in equilibrium. So in this case, platform 1 attracts all sellers in the fee-setting period if and only if $x_1 \geq x_2$. Working backwards to the discovery-setting period, we must have $x_1 = 1$ (otherwise platform 2 could profitably deviate to $x_2 = 1$ and attract all sellers in the second period) and $x_2 = 0$ (otherwise platform 2 would make negative profits given the fixed cost ε). This leads to $f_1 = (v + b - c) \left(1 - \frac{1}{n(1-\theta)}\right)$ and $f_2 = 0$. Platform 1's profit is $(v + b - c) \left(1 - \frac{1}{n(1-\theta)}\right) n (1 + (n-1)(1-\theta)) - \varepsilon$, which is positive given $\theta < 1 - \frac{1}{n}$ pro-

[21] In this case, other equilibria would also arise: the symmetric equilibrium, where the roles of the two platforms are switched, and a mixed strategy equilibrium where both platforms invest in discoverability with some positive probability less than one.

vided ε is sufficiently small.

Now suppose $\theta \geq 1 - \frac{1}{n}$. In this case, the highest payoff that platform i can ever offer each seller in the fee-setting period is $v - c + b$, by setting $x_i = 0$ in the discovery-setting period. Thus, there is never any advantage for either platform to set $x_i > 0$, which means the only possible equilibrium is $x_1 = x_2 = 0$ and $f_1 = f_2 = 0$.

Comparing this equilibrium outcome to the outcome with a monopoly platform and n-sellers (see expression 19 in the proof of Proposition 4), it is apparent that for $\theta \leq 1 - \frac{1}{n}$ and for $\theta \geq \frac{(\mu+1)(n-1)}{(\mu+1)(n-1)+1}$, the levels of discoverability across monopoly and competing platforms cases coincide. For θ in between these two cutoffs, there is strictly more discoverability with a monopoly platform than with competing platforms. Thus, we have proven the following proposition.

Proposition 7. *Assume ε is sufficiently small. There are two cases. If $\theta < 1 - \frac{1}{n}$, then in equilibrium platform 1 offers maximum discoverability, platform 2 offers no discoverability, and all sellers join platform 1. If $\theta \geq 1 - \frac{1}{n}$, then in equilibrium neither platform offers any discoverability and each seller joins either platform. Platform competition decreases discoverability.*

The reason platform competition deceases discoverability is that when platforms compete, they focus on maximizing the payoff to sellers in order to attract them, and sellers generally prefer less discoverability than the platforms.

4.5.2 Competition between differentiated platforms

Suppose there are two platforms. Platform 1 has $\eta > 0$ of its own captive buyers modelled in as in Section 4.1, but we normalize the value of its seller tools to $b = 0$. One can think of platform 1 as Amazon. Platform 2 has no captive buyers but offers $b > 0$. One can think of it as Shopify.

We assume there are three sellers. One seller is large, and has measure λ initially captive buyers. Two sellers are "small" and come with no captive buyers. The two platforms choose their respective values of x_1 and x_2 simultaneously in period 1.[22] We obtain the following result (the proof is in the Appendix).

Proposition 8. *Assume $\lambda \geq \frac{(1-\theta)\eta}{\theta}$ and $b \leq \frac{\theta}{2(1-\theta)}(v - c)$. In the unique equilibrium, platform 1 attracts the two small sellers, setting $f_1^* = v - c$ and $x_1^* = 1$, and platform 2 attracts the large seller, setting $f_2^* = b$ and $x_2^* = 0$.*

[22]The equilibrium characterization that follows is unchanged if the platforms choose x_1 and x_2 sequentially, in either order, and is also robust to platforms facing some small fixed cost ε of choosing $x_i > 0$.

This result shows the possibility for the co-existence of two competing platforms, one that attracts the larger seller by not offering discoverability, and one that attracts the two small sellers by offering maximum discoverability. The assumption that the small sellers bring no buyers of their own was made to keep things tractable. The logic underlying the result in Proposition 8 extends to the case when the small sellers have positive measures of captive buyers, though the calculations would be more complicated. Here, the small sellers benefit from discoverability by gaining access to platform 1's captive buyers—this captures the value proposition of Amazon for many of its sellers. By contrast, the large seller stands to lose more than it gains from discoverability when joining the same platform as the small sellers, so it is attracted by platform 2 which offers more valuable tools.

The possibility that platforms endogenously differentiate themselves by attracting different sellers is reminiscent of Karle et al. (2020), where platforms may compete to attract different sellers, thereby allowing sellers to avoid head-to-head competition. The assumptions in Proposition 8 are sufficient for the existence and uniqueness of the equilibrium, but are not necessary. The assumptions require the large seller be large enough and the value of the tools that platform 2 provides not be too high. Both assumptions can always be satisfied by taking θ sufficiently close to one.

5 Managerial implications

Discoverability is an essential feature of marketplaces and multi-sided platforms. It is a key driver of the strength and defensibility of their network effects. Buyers come to the platform either because they know the platform or because they follow a seller they are familiar with, and then they discover other sellers on the platform, which leads to more transactions. Thus, more discoverability means stronger and more defensible network effects. This is why established platforms (e.g. Amazon, eBay, Etsy, DoorDash, Udemy) prioritize buyers discovering as many new sellers as possible.

However, it is important to realize that discoverability creates a tradeoff by making it harder to attract sellers that have their own installed base of buyers. Thus, our analysis is particularly relevant for (a) established platforms trying to attract branded/established sellers, (b) emerging platforms seeking to solve the chicken-and-egg problem by attracting sellers that have their own installed base of buyers, (c) B2B providers of tools (e.g. Shopify, Substack, Teachable) that need to decide how much (if any) discoverability to enable in order to generate network effects.

For such platforms, our analysis delivers a few useful implications.

First, discoverability only creates a problem for sellers of competing products, but is welcomed by sellers of complementary or independent products. Thus, platforms should maximize the discoverability of complementary or independent products/sellers and only tread carefully with respect to competing products/sellers. More specifically, platforms should enable more discoverability across distinct product categories and less discoverability within the same product category (where sellers tend to be close competitors). Finally, if it is possible to use customer data and AI to learn about the extent to which buyers view sets of products as substitutable or not, the platform should create more discoverability for the sets of products that they view as more independent or complementary.

Second, the more established a platform becomes in terms of how many buyers are loyal to it, relative to buyers loyal to individual sellers, the more discoverability it can afford to enable, so as to maximize transactions. This is what most of today's large established platforms have done over time: once their respective sellers are dependent on them because they have the buyers these sellers want to reach, the platforms have taken steps to commoditize the sellers. Still, once sellers realize what is happening, they may want to jump ship to platforms that don't commoditize them (i.e. with less or no discoverability). That's why Shopify has been thriving despite the apparent dominance of Amazon in the U.S. and some other markets, and why Olo and others can thrive despite the presence of several large online food delivery platforms (e.g. DoorDash, Grubhub, Uber Eats). Committing not to commoditize sellers via public statements (e.g. Shopify's "arm the rebels",[23] Teachable's "escape the algorithm"[24]) can then be part of a strategy to attract sellers and differentiate from large, established marketplaces.

Third, if feasible, a good way to mitigate the fundamental tradeoff created by discoverability is to charge sellers differential transaction fees for buyers they brought to the platform vs. buyers that discovered them through the platform. This is intuitively fairer and when implemented well, can help the platform incentivize sellers to bring more buyers to the platform, while still enabling a high level of discoverability—so the platform can have its cake and eat it too. The problem is that to implement this differential pricing, platforms need to reliably distinguish buyers that were brought to the platform by a given seller from buyers that discovered the seller via the platform, which may be difficult in some cases. One practical way to address this is to give sellers specific registration links to share with their buyers and/or create a separate website or app where buyers can find the platform directly (e.g. Shopify's Shop.app). Another solution is to charge sellers an additional per-click fee

[23]https://blakeir.medium.com/arming-the-rebels-of-the-future-d61b3fe30515

[24]https://teachable.com/

when buyers are searching in a seller's product category as opposed to when they search the seller's name or go straight to its listings.

Fourth, discoverability is more attractive to sellers with small installed bases of loyal buyers and less attractive to sellers with large bases of loyal buyers. This is why large established platforms like Amazon.com have trouble attracting larger brands which can afford to avoid commoditization at the hands of a platform that enables too much discoverability. This means that in general platforms face a choice between two broad options: they can serve sellers that have the ability to attract their own buyers with tools but limited discoverability (implying weaker network effects), or they can mostly focus on smaller sellers (with limited ability to attract their own buyers) while enabling high levels of discoverability. Which also means discoverability can be a source of endogenous differentiation among platforms: if an established platform enables high levels of discoverability (e.g. Amazon or DoorDash), then a new entrant can focus on tools and credibly commit to keep discoverability low (e.g. Shopify or Olo) in order to attract sellers that have or can attract their own buyers.

6 Conclusion

We have provided a framework for analyzing the extent to which platforms want to allow buyers who are brought in by participating sellers to discover rival sellers. While we have covered many different extensions of the simple baseline setting in the paper, there remain many other avenues to explore in future work.

Further analysis of competing platforms seems warranted, although this remains challenging. For instance, it would be interesting to explore other types of heterogeneity between sellers, to understand how different seller characteristics drive their preferences over platforms that offer different levels of discoverability. In our analysis of competing platforms, we assumed a seller would only go to one platform or the other, bringing all its buyers onto the chosen platform. Another possibility would be to allow the seller to determine the portion of its initially captive buyers it brings onto each platform, or possibly to both.

Extending our analysis to allow the platform to offer first-party products would be another interesting avenue to pursue. On the one hand, the presence of first-party products should make third-party sellers more reluctant to participate, which may require the platform to dial back its discoverability. On the other hand, the platform now derives an additional benefit from discoverability: its own products being discovered by buyers that were brought by other sellers.

Finally, considering a dynamic setting where the sellers' initial captives become loyal to the platform after some time would possibly provide a rationale for platforms to increase the extent of discoverability they offer over time. This dynamic inconsistency problem creates a role for credible commitments by the platform.

7 Appendix

We provide the remaining details for the proofs of each proposition.

7.1 Proof of Propositions 1 and 2

We directly prove Proposition 2, which is more general. The proof of Proposition 1 then follows by setting $\eta = 0$ (i.e. the platform starts with no buyers of its own).

Factoring out the constant term $v - c$, the derivative of (9) with respect to x is

$$(2(1-\theta)+\eta(2-\theta))(\mu+1) - \frac{(2+\eta)\theta}{(1+x((1-2\theta)+\eta(1-\theta)))^2}. \tag{15}$$

If $\theta \leq \frac{1+\eta}{2+\eta}$, then (15) is increasing in x and is non-negative when evaluated at $x = 0$, so we must have $x^* = 1$. If $\theta > \frac{1+\eta}{2+\eta}$, then (15) is decreasing in x, so the second-order condition (SOC) holds. Setting (15) equal to zero and solving for x implies the unconstrained solution

$$x(\theta) = \frac{1 - \sqrt{\frac{(2+\eta)\theta}{(2(1-\theta)+\eta(2-\theta))(\mu+1)}}}{2\theta - 1 - \eta(1-\theta)}.$$

Given $x(\theta)$ is decreasing in θ for $\theta > \frac{1+\eta}{2+\eta}$, and given $x\left(\frac{1+\eta}{2+\eta}\right) > 1$ and $x(\theta) < 0$ for θ sufficiently high, the constrained solution is given by

$$x(\theta) = \begin{cases} 1 & if \quad 0 < \theta \leq \theta_1(\eta,\mu) \\ \frac{1-\sqrt{\frac{(2+\eta)\theta}{(2(1-\theta)+\eta(2-\theta))(\mu+1)}}}{2\theta-1-\eta(1-\theta)} & if \quad \theta_1(\eta,\mu) \leq \theta \leq \theta_2(\eta,\mu) \\ 0 & if \quad \theta \geq \theta_2(\eta,\mu) \end{cases} \tag{16}$$

and $\theta_1(\eta,\mu) \in \left(\frac{1+\eta}{2+\eta}, \theta_2(\eta,\mu)\right)$ is the unique solution in θ to $x(\theta) = 1$, and where $\theta_2(\eta,\mu) = \frac{2(1+\eta)(\mu+1)}{(2+\eta)(\mu+2)} > \frac{1+\eta}{2+\eta}$ is the unique solution to $x(\theta) = 0$.

Setting $\eta = 0$, we obtain the results in Proposition 1. Note that $\theta_2(0,\mu) = \frac{\mu+1}{\mu+2} < 1$, but with $\eta > 0$, we can have $\theta_2(\eta,\mu) > 1$.

Finally, the comparative static results hold since the cutoffs $\theta_1(\eta,\mu)$ and $\theta_2(\eta,\mu)$, as well as the interior solution for x^*, are all increasing in η and μ, and in the case of the interior solution for x^*, decreasing in θ.

7.2 Proof of Proposition 3

The sellers' participation constraint can be rewritten as

$$f_1 x (1-\theta) + f_0 (1 - x\theta) \leq (v-c) x (1-2\theta) + b (1 + x (1-2\theta)).$$

Suppose the sellers' participation constraint is not binding at the optimum. Then we must have $f_0 = f_1 = v + b - c$, otherwise the platform could profitably increase either f_0 or f_1. But then the sellers' participation constraint is equivalent to $v - c \leq 0$, which is not possible.

So the sellers' participation constraint must be binding at the optimum, i.e. we must have

$$f_1 x (1-\theta) + f_0 (1 - x\theta) = (v-c) x (1-2\theta) + b (1 + x (1-2\theta)).$$

We can use this to express f_1 as a function of f_0. The platform's profits can then be written as

$$-f_0 \frac{\theta (1-x)}{1-\theta} + \frac{(2-\theta)}{1-\theta} ((v-c+b) x (1-2\theta) + b),$$

which the platform maximizes over (f_0, x) subject to

$$0 \leq f_0 \leq v + b - c$$

and

$$0 \leq (v-c) \frac{1-2\theta}{1-\theta} + b \frac{1 + x (1-2\theta)}{x (1-\theta)} - f_0 \frac{(1-x\theta)}{x (1-\theta)} \leq v + b - c.$$

Since the last expression of platform profits is decreasing in f_0, we must either have

$$f_0 = 0$$

or

$$(v-c) \frac{1-2\theta}{1-\theta} + b \frac{1 + x (1-2\theta)}{x (1-\theta)} - f_0 \frac{(1-x\theta)}{x (1-\theta)} = v + b - c.$$

Considering each case, and the corresponding maximizing choice of x, involves lengthy and tedious algebra, which we have relegated to Online Appendix A.13. There we confirm that $f_1 > f_0$, and establish that the platform's optimal choice of x^* is:

$$x^* = \begin{cases} 1 & \text{if} \quad 0 < \theta \leq \frac{1}{2} \\ \frac{\mu}{(\mu+1)\theta} & \text{if} \quad \frac{1}{2} \leq \theta \leq \frac{2}{3+\mu} \\ \frac{1-\sqrt{\frac{\theta}{2(\mu+1)(1-\theta)}}}{\theta} & \text{if} \quad \frac{2}{3+\mu} \leq \theta \leq \frac{2(\mu+1)}{2(\mu+1)+1} \\ 0 & \text{if} \quad \theta \geq \frac{2(\mu+1)}{2(\mu+1)+1} \end{cases} \tag{17}$$

when $\mu \leq 1$, and by

$$x^* = \begin{cases} 1 & \text{if} \quad 0 < \theta \leq \theta_0(\mu) \\ \frac{1-\sqrt{\frac{\theta}{2(\mu+1)(1-\theta)}}}{\theta} & \text{if} \quad \theta_0(\mu) \leq \theta \leq \frac{2(\mu+1)}{2(\mu+1)+1} \\ 0 & \text{if} \quad \theta \geq \frac{2(\mu+1)}{2(\mu+1)+1} \end{cases} \tag{18}$$

when $\mu \geq 1$, where $\theta_0(\mu)$ is the unique solution to

$$\frac{\theta}{(1-\theta)^3} = 2(\mu+1).$$

Note inspecting the above cases, x^* is either constant or increasing in μ, and is decreasing in θ.

7.3 Proof of n-seller case

Following the same steps as in the two-seller case, the platform's optimal level of x is given by

$$x^* = \begin{cases} 1 & if \quad 0 < \theta \leq \theta_1(\mu, n) \\ \frac{1-\sqrt{\frac{\theta}{(\mu+1)(1-\theta)(n-1)}}}{1-n(1-\theta)} & if \quad \theta_1(\mu, n) \leq \theta \leq \theta_2(\mu, n) \\ 0 & if \quad \theta \geq \theta_2(\mu, n) \end{cases}, \tag{19}$$

where $\theta_1(\mu, n) \in \left(1 - \frac{1}{n}, \theta_2(\mu, n)\right)$ is the unique solution in θ to

$$\frac{\theta}{(1-\theta)^3} = n^2(n-1)(\mu+1)$$

and $\theta_2(\mu, n) = \frac{(\mu+1)(n-1)}{(\mu+1)(n-1)+1}$. To show the result, we can define the function $X(\theta) = n^2(n-1)(1-\theta)^3(1+\mu)$, which is strictly decreasing in θ with $X\left(1-\frac{1}{n}\right) > 1 - \frac{1}{n}$. This implies $1 - \frac{1}{n} < \theta_1$ and $X(\theta_2) < \theta_2$, implying $\theta_2 > \theta_1$. As a result, for $\theta \leq \theta_1$ we have $x^* = 1$, and for $\theta > \theta_2$ we have $x^* = 0$. Finally, the comparative static results hold since the cutoffs $\theta_1(\mu, n)$ and $\theta_2(\mu, n)$, as well as the interior solution for x^*, are all increasing in n and μ, and in the case of the interior solution for x^*, decreasing in θ.

7.4 Proof of Proposition 5

Let the measure of captives that seller i obtains be denoted λ_i'. To handle this case we use the result in Proposition 1 of Myatt and Ronayne (2023) to determine each seller's expected profit.[25] Their result covers the case of two sellers i and j with $\lambda_i' > \lambda_j'$ captives and the same marginal costs c. Seller i is the less aggressive seller as it has more captives, meaning $p_i^+ > p_j^+$ in their notation. Then seller j's expected profit is

$$\left(\lambda_j' + \phi\right)\left(p_i^+ - c\right) = \frac{\lambda_j' + \phi}{\lambda_i' + \phi}\lambda_i'\left(v - c\right) > \lambda_j'\left(v - c\right),$$

while seller i's expected profit is $\lambda_i'\left(v - c\right)$, where ϕ is the measure of buyers informed of both sellers and view them as substitutes.

Following the same logic for the measure of captives of seller 1 in the main text, the captive buyers for seller i in general are

$$\lambda_i' = \lambda_i\left(1 - x\right) + \lambda_i x\left(1 - \theta\right) + \lambda_j x\left(1 - \theta\right).$$

Given $\lambda_1 > \lambda_2$, we have $\lambda_1' > \lambda_2'$. Moreover, $\phi = \left(\lambda_i + \lambda_j\right) x\theta$.

Thus, seller 1's profit is

$$\left(v + b - f - c\right)\left(\lambda_1 + \lambda_2\right)\left(\beta_1\left(1 - x\right) + x\left(1 - \theta\right)\right)$$

and seller 2's profit is

$$\left(v + b - f - c\right)\left(\lambda_1 + \lambda_2\right)\frac{\left(1 - \beta_1\right)\left(1 - x\right) + x}{\beta_1\left(1 - x\right) + x}\left(\beta_1\left(1 - x\right) + x\left(1 - \theta\right)\right),$$

where $\beta_1 \in \left[\frac{1}{2}, 1\right]$ is defined in the main text.

Seller 1 participates iff (14) holds and seller 2 participates iff

$$f \le b + \left(v - c\right)\left(1 - \frac{\left(1 - \beta_1\right)\left(\beta_1\left(1 - x\right) + x\right)}{\left(\left(1 - \beta_1\right)\left(1 - x\right) + x\right)\left(\beta_1\left(1 - x\right) + x\left(1 - \theta\right)\right)}\right).$$

Since $\beta_1 \ge \frac{1}{2}$, we have

$$\frac{\beta_1}{\beta_1\left(1 - x\right) + x\left(1 - \theta\right)} \ge \frac{\left(1 - \beta_1\right)\left(\beta_1\left(1 - x\right) + x\right)}{\left(\left(1 - \beta_1\right)\left(1 - x\right) + x\right)\left(\beta_1\left(1 - x\right) + x\left(1 - \theta\right)\right)},$$

so the binding constraint is (14) of seller 1. Clearly f will be set at the maximum value

[25] For completeness, we've restated the relevant part of Proposition 1 of Myatt and Ronayne in the Online Appendix A.5, which is much more general than the result stated here.

allowed by the constraint, so the platform maximizes

$$\left(\lambda_1+\lambda_2\right)(v-c)\left(\mu+1-\frac{1}{1-x+\frac{x(1-\theta)}{\beta_1}}\right)(1+x(1-\theta)).$$

over x.

If $\theta \leq 1-\beta_1$, then $1-x+\frac{x(1-\theta)}{\beta_1}$ is increasing in x, so the profit expression above is increasing in x, which means $x^*=1$. Now suppose $\theta>1-\beta_1$. The derivative of the profit expression above in x is

$$\left(\lambda_1+\lambda_2\right)(v-c)\left((\mu+1)(1-\theta)-\frac{2-\theta-\frac{1-\theta}{\beta_1}}{\left(1-x+\frac{x(1-\theta)}{\beta_1}\right)^2}\right).$$

Since $2-\theta-\frac{1-\theta}{\beta_1} \geq 0$ and we have assumed $\theta>1-\beta_1$, the last expression above is decreasing in x, so the SOC holds. From this, we directly conclude that the optimal level of discovery is given by

$$x^*=\begin{cases} 1 & if \quad 0<\theta \leq \theta_1\left(\mu, \beta_1\right) \\ \frac{1-\sqrt{\frac{2-\theta-\frac{1-\theta}{\beta_1}}{(\mu+1)(1-\theta)}}}{1-\frac{1-\theta}{\beta_1}} & if \quad \theta_1\left(\mu, \beta_1\right) \leq \theta \leq \theta_2\left(\mu, \beta_1\right) \\ 0 & if \quad \theta \geq \theta_2\left(\mu, \beta_1\right)\end{cases}, \tag{20}$$

where

$$\theta_2\left(\mu, \beta_1\right)=\frac{\mu+\frac{1}{\beta_1}-1}{\mu+\frac{1}{\beta_1}} \in\left[1-\beta_1, 1\right]$$

and $\theta_1\left(\mu, \beta_1\right) \in\left(1-\beta_1, \theta_2\left(\mu, \beta_1\right)\right)$ is the unique solution to

$$\frac{2-\frac{1}{\beta_1}+\left(\frac{1}{\beta_1}-1\right) \theta}{(1-\theta)^3}=\frac{\mu+1}{\beta_1^2}.$$

Finally, the comparative static results hold since the cutoffs $\theta_1\left(\mu, \beta_1\right)$ and $\theta_2\left(\mu, \beta_1\right)$, as well as the interior solution for x^*, are all increasing in μ and decreasing in β_1, and in the case of the interior solution for x^*, decreasing in θ.

7.5 Proof of Proposition 6

As argued in the main text, it can never be optimal for the platform to induce only one seller to join. Furthermore, it can never be optimal for the platform to induce the large seller to join together with only one small seller. Indeed, if this was the case, the large seller must prefer joining together with a small seller than its outside option, i.e. we would have

$$(v+b-f-c)\left((1-x)\lambda_1+(\lambda_2+\lambda_1)x(1-\theta)\right)\geq(v-c)\lambda_1$$

Meanwhile, the condition for the second small seller to prefer not joining when the other two sellers have joined is

$$(v+b-f-c)\left((1-x)\lambda_2+(\lambda_1+2\lambda_2)x(1-\theta)\right)<(v-c)\lambda_2.$$

It can be easily verified that these two conditions are incompatible, so there cannot be an equilibrium with one large seller and one small seller joining for any (f,x). Nor would the platform want to force the outcome in which only one large seller and one small seller join. Indeed, from the analysis above, the maximum transaction fee it could charge would be

$$f=b+(v-c)\left(1-\frac{1}{(1-x)+x(1-\theta)\frac{\lambda_2+\lambda_1}{\lambda_1}}\right).$$

At this fee, we know that the second small seller would also be willing to join. The platform's profits with one large seller and one small seller are

$$f(\lambda_1+\lambda_2)(1+x(1-\theta)),$$

whereas with all three sellers participating, the platform would make

$$f(\lambda_1+2\lambda_2)(1+x(1-\theta)),$$

which is strictly larger.

Thus, there are only two possibilities for the platform's optimal strategy: either all three sellers join the platform or only the two small sellers join.

In the case where only the two small sellers join, the platform sets x as in the baseline, except here we have assumed $b=\mu=0$, so

$$x_2^*=\begin{cases}1 & if \quad 0<\theta\leq\frac{1}{2}\\ 0 & if \quad \theta\geq\frac{1}{2}\end{cases},$$

The platform's optimal fee and resulting profits for this case are

$$\begin{aligned} f_2^* &= \begin{cases} (v-c)\,\frac{1-2\theta}{2(1-\theta)} & if \quad 0<\theta\leq\frac{1}{2} \\ 0 & if \quad \theta\geq\frac{1}{2} \end{cases} \\ \Pi_2^* &= \begin{cases} \lambda_2\,(v-c)\,\frac{(1-2\theta)(2-\theta)}{1-\theta} & if \quad 0<\theta\leq\frac{1}{2} \\ 0 & if \quad \theta\geq\frac{1}{2} \end{cases} . \end{aligned}$$

In the case where all three sellers join the platform, the large seller's profit is

$$(v-c-f)\,(\lambda_1\,(1-x)+(\lambda_1+2\lambda_2)\,x\,(1-\theta)),$$

while the two small sellers each make a profit equal to

$$(v-c-f)\,(\lambda_2\,(1-x)+(\lambda_1+2\lambda_2)\,x\,(1-\theta)).$$

For the large seller to participate we must have

$$f\leq(v-c)\left(1-\frac{\lambda_1}{\lambda_1\,(1-x)+(\lambda_1+2\lambda_2)\,x\,(1-\theta)}\right)$$

For the small sellers to participate we must have

$$f\leq(v-c)\left(1-\frac{\lambda_2}{\lambda_2\,(1-x)+(\lambda_1+2\lambda_2)\,x\,(1-\theta)}\right).$$

Since $\lambda_1>\lambda_2$, the binding constraint must be that of the large seller, so for f to be optimal, it must be that

$$f=(v-c)\left(1-\frac{1}{1-x+\frac{x(1-\theta)}{\beta}}\right).$$

Platform profits are then

$$\begin{aligned} & f\,(\lambda_1+2\lambda_2)\,(1+x\,(1-\theta)) \\ = \; & (\lambda_1+2\lambda_2)\,(v-c)\left(1-\frac{1}{1-x+\frac{x(1-\theta)}{\beta}}\right)(1+x\,(1-\theta)) \end{aligned}$$

In this case, the optimal level of discoverability is

$$x_3^*=\begin{cases} 1 & \text{if} \quad 0<\theta\leq 1-\beta \\ 0 & \text{if} \quad \theta\geq 1-\beta \end{cases} .$$

And the platform's profits are

$$\Pi_3^* = \begin{cases} (\lambda_1 + 2\lambda_2)(v-c)\frac{(1-\beta-\theta)(2-\theta)}{1-\theta} & \text{if } \ 0 < \theta \leq 1-\beta \\ 0 & \text{if } \quad \theta \geq 1-\beta \end{cases}.$$

Thus, when $\theta \geq \max\left\{\frac{1}{2}, 1-\beta\right\}$, we have $\Pi_3^* = \Pi_2^* = 0$, and when $\theta \leq \min\left\{\frac{1}{2}, 1-\beta\right\}$, we have $\Pi_3^* \geq \Pi_2^*$ iff $\lambda_2 \geq \theta\lambda_1$. If $1-\beta \leq \theta < \frac{1}{2}$ (which can only happen when $\beta > \frac{1}{2}$), then $\Pi_2^* > \Pi_3^*$. And if $\frac{1}{2} \leq \theta < 1-\beta$ (which can only happen when $\beta > \frac{1}{2}$), then $\Pi_3^* > \Pi_2^*$. From this we can conclude:

- If $\theta \geq \frac{1}{2}$, then $\Pi_3^* > \Pi_2^*$ for all $\beta < 1-\theta$ and $\Pi_3^* = \Pi_2^* = 0$ for all $\beta \geq 1-\theta$.
- If $\theta < \frac{1}{2}$, then $\Pi_2^* > \Pi_3^*$ iff $\beta > \frac{1}{1+2\theta}$.

7.6 Proof of Proposition 8

We work backwards. Suppose first that $x_1 = 1$ and $x_2 = 0$. We want to show $f_1 = v - c$ and $f_2 = b$ constitute an equilibrium in the period 2 subgame. Note given $x_2 = 0$ and given platform 2 has no captive buyers, it cannot offer any positive surplus to the small seller(s). To make any profit it must attract the large seller. It offers the large seller a surplus of $\lambda(v + b - c - f_2)$. The best it can possibly do therefore in period 2 is to charge $f_2 = b$ to obtain

$$\Pi_2 = \lambda b.$$

In the proposed period 2 equilibrium, platform 1 attracts the two small sellers with the fee $f_1 = v - c$. This leaves the two small sellers indifferent about joining, so platform 1 cannot charge more than this. Platform 1 obtains a profit of

$$\begin{aligned} \Pi_1 &= \eta(\theta + 2(1-\theta))(v-c) \\ &= \eta(2-\theta)(v-c). \end{aligned}$$

There is just one possible deviation to check in period 2, which is the deviation where platform 1 tries to attract the large seller, either by itself or in conjunction with the two small sellers.

If the large seller joins platform 1 alone it gets $(\lambda + \eta)(v - c - f_1)$, so this requires

$$(\lambda + \eta)(v - c - f_1) \geq \lambda(v - c),$$

i.e.

$$f_1 \leq \frac{\eta (v - c)}{\lambda + \eta}.$$

This means that if platform 1 restricts entry to the large seller, the maximum profit it can make is

$$(\lambda + \eta) f_1 = \eta (v - c),$$

which is less than the profit $\eta (2 - \theta)(v - c)$ it makes in the proposed equilibrium.

Suppose now platform 1 attracts all three sellers. The two small sellers join as soon as $f_1 < v - c$ because their outside option is zero. The large seller's payoff from joining platform 1 alongside the two small sellers is

$$(\lambda + \eta)(1 - \theta)(v - c - f_1)$$

since $(\lambda + \eta)(1 - \theta)$ of the time it faces captive buyers (and it is the less aggressive seller, so its expected profit is just determined by its captives). So to attract the large seller alongside the two small sellers, platform 1 would need to set f_1 so that

$$(\lambda + \eta)(1 - \theta)(v - c - f_1) \geq \lambda (v - c),$$

i.e.

$$f_1 \leq \frac{((1 - \theta)(\eta + \lambda) - \lambda)(v - c)}{(1 - \theta)(\eta + \lambda)}.$$

Given the assumption $\lambda > \frac{(1-\theta)\eta}{\theta}$, the RHS above is negative, so there can be no profitable deviation for platform 1 in stage 2 in which it attracts all three sellers.

We have thus proven that there are no profitable deviations in stage 2. It remains to check neither platform can benefit from deviating in their choices of x.

Suppose platform 1 deviates to $x_1 < 1$. If platform 1 ends up attracting just the two small sellers, this cannot be a profitable deviation because it would obtain at most $\eta x_1 (2 - \theta)(v - c)$, which is less than $\eta (2 - \theta)(v - c)$, the profit it obtained in the proposed equilibrium with $x_1 = 1$. Similarly, if it ends up attracting just the large seller, its profit would be $(\eta x_1 + \lambda) f_1$, where we must have

$$(\lambda + x_1 \eta)(v - c - f_1) \geq \lambda (v - c)$$

for the large seller to prefer joining over its outside option. This means

$$(\eta x_1 + \lambda) f_1 \leq \eta x_1 (v - c) < \eta (2 - \theta)(v - c),$$

so the deviation cannot be profitable for platform 1 in this case. Suppose now that with the deviation to $x_1 < 1$, platform 1 attracts all three sellers. The large seller's payoff from coming to platform 1 together with the two small sellers is then

$$((1-x_1)\lambda + x_1(\lambda+\eta)(1-\theta))(v-c-f_1).$$

So to attract the large seller, platform 1 would need to set f_1 so that

$$((1-x_1)\lambda + x_1(\lambda+\eta)(1-\theta))(v-c-f_1) > \lambda(v-c),$$

which can be rewritten

$$f_1 \leq \left(1 - \frac{\lambda}{(1-x_1)\lambda + x_1(\lambda+\eta)(1-\theta)}\right)(v-c).$$

And it is easily verified that the RHS is negative under the condition $\lambda > \frac{(1-\theta)\eta}{\theta}$. So no deviation to $x_1 < 1$ can be profitable for platform 1.

Now consider platform 2 deviating to $x_2 > 0$. The only way in which it could obtain higher profits than in the proposed equilibrium would be to attract all three sellers (attracting just the two small sellers doesn't work because it has no captive buyers and attracting just the large seller would do at most just as well as the proposed equilibrium). If it attracts all three sellers, then the large seller's payoff from coming to platform 2 together with the two small sellers is $\lambda(1-x_2\theta)(v+b-c-f_1)$. So to attract the large seller, platform 2 would need to set f_1 so that $\lambda(1-x_2\theta)(v+b-c-f_1) > \lambda(v-c)$, i.e.,

$$f_1 \leq b - \frac{x_2\theta}{1-x_2\theta}(v-c).$$

And platform 2's corresponding profits would be $\lambda(1+2x_2(1-\theta))f_1$. Thus, the maximum profits that platform 2 would be able to extract with this deviation is

$$\lambda(1+2x_2(1-\theta))\left(b - \frac{x_2\theta}{1-x_2\theta}(v-c)\right).$$

The derivative with respect to x_2 is increasing in b, so evaluating this derivative at $b = \frac{\theta(v-c)}{2(1-\theta)}$, the upper bound on b in our assumption above, the derivative in x_2 is negative for any $0 < x_2 \leq 1$ and equal to zero when $x_2 = 0$. Thus, platform 2's profit is maximized at $x_2 = 0$, which yields λb, i.e. its profit in the proposed equilibrium.

References

Anderson, S. and Ozlem, B-D. (2022) "Online trade platforms: Hosting, selling, or both?" *International Journal of Industrial Organization*, 84, No. 102861.

Baojun, J. and Zou, T. (2020) "Consumer Search and Filtering on Online Retail Platforms" Journal of Marketing Research, Volume 57, Issue 5.

Barach, M., Joseph M., Golden J., and Horton, J. (2020) "Steering in Online Markets: The Role of Platform Incentives and Credibility" Management Science. Vol. 66, No. 9.

Baye, M. R. and Morgan, J. (2001) "Information gatekeepers on the internet and the competitiveness of homogeneous product markets," American Economic Review, Vol. 91(3), 454-474.

Casner, B. (2020) "Seller Curation in Platforms" *International Journal of Industrial Organization*, Vol 72. No. 102659.

Casner, B. and Teh, T-H. (2023) "Content-hosting platforms: discovery, membership, or both?" Working paper.

Choi, J-P. and Jeon, D-S. (2023) "Platform Design Biases in Ad-Funded Two-Sided Markets" *RAND Journal of Economics*, forthcoming.

de Corniére, A. (2016) "Search Advertising" *American Economic Journal: Microeconomics*, Vol 8, No. 3, 156-188.

Dukes, A. and Liu, L. (2016). "Online shopping intermediaries: The strategic design of search environments" *Management Science*, Vol. 62 No. 4. pp. 1064–1077.

Hagiu, A. and Altman, E. (2017) "Finding the Platform in Your Product," *Harvard Business Review*, 95(4), 94-100.

Hagiu, A. and Jullien, B. (2011) "Why do intermediaries divert search?" *RAND Journal of Economics*, Volume 42, No. 2, pp. 337-362.

Hagiu, A., Jullien, B., and Wright, J. (2020) "Creating platforms by hosting rivals" *Management Science*, Vol 66. No. 7, 3234–3248.

Hagiu, A. and Wright, J. (2021) "Don't Let Platforms Commoditize Your Business," *Harvard Business Review*. May–June 2021.

Hagiu, A., Teh, T-H., and Wright, J. "Should platforms be allowed to sell on their own marketplaces?" *RAND Journal of Economics*, Vol. 53. No. 2, pp. 297-327.

Hagiu, A. and Wright, J. (2023) "Marketplace leakage" *Management Science*, forthcoming.

Karle, H., Peitz, M., and Reisinger, M. (2020) "Segmentation versus agglomeration: Competition between platforms with competitive sellers," Journal of Political Economy, 128(6), 2329-2374.

Madsen, E. and Vellodi, N. (2022) “Insider Imitation” Working paper.

Mayya, R. and Li, Z. (2022) “Growing Platforms by Adding Complementors Without Consent: Evidence from On-Demand Food Delivery Platforms” Working paper.

Myatt, D. P. and D. Ronayne (2023) “Asymmetric Models of Sales” Working paper.

Ronayne, D. (2021) “Price comparison websites,” International Economic Review, Vol. 62(3), 1081-1110.

Ronayne, D., and Taylor, G. (2022) “Competing sales channels with captive consumers,” The Economic Journal, Vol. 132(642), 741-766.

Shi, Z., Srinivasan, K., and Zhang, K. (2023) “Design of Platform Reputation Systems: Optimal Information Disclosure” *Marketing Science*, forthcoming.

Teh, T-H. (2022) “Platform Governance” *American Economic Journal: Microeconomics*, Vol 14, No. 3, pp. 213-54.

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

Zhong, Z. (2023) “Platform Search Design: The Roles of Precision and Price” *Marketing Science*, Vol. 42, No. 2

Zhou, B. and Zou, T. (2023) “Competing for Recommendations: The Strategic Impact of Personalized Product Recommendations in Online Marketplaces” *Marketing Science*, Vol. 42, No. 2.