

Creating platforms by hosting rivals*

Andrei Hagiu[†], Bruno Jullien[‡] and Julian Wright[§]

November 2018

Abstract

We explore conditions under which a multiproduct firm can profitably turn itself into a platform by “hosting rivals,” i.e. by inviting rivals to sell products or services on top of its core product. Hosting eliminates the additional shopping costs to consumers of buying a specialist rival’s competing version of the multiproduct firm’s non-core product. On the one hand, this makes it easier for the rival to compete on the non-core product. On the other hand, hosting turns the rival from a pure competitor into a complementor: the value added by its product now helps raise consumer demand for the multi-product firm’s core product. As a result, hosting can be both unilaterally profitable for the multi-product firm and jointly profitable for both firms.

JEL classification: D4, L1, L5

Keywords: multi-sided platforms, shopping costs, bundling, competition, complementarity.

1 Introduction

Recently a lot of attention has been given to multi-sided platforms such as those operated by Airbnb, Alibaba, eBay, Expedia, Facebook and Tencent, to name a few. In part, this reflects that many of the most valuable companies in the world today generate a lot of their revenue from platform businesses, focusing on facilitating interactions or transactions between different parties (e.g. buyers and sellers) rather than selling products or services that they own or produce themselves.

However, in many cases, existing product (or service) companies have the potential to become (multi-sided) platforms too. The most straightforward way for a product company to do so is by

*We thank Patrick Rey and Alex White for very helpful comments, as well as participants at 44th EARIE in Maastricht, the Online Platform Competition Conference at the University of Florida Gainesville, 16th ZEW Conference on ICT in Mannheim, and the 2018 platform research symposium at Boston University.

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

[‡]Toulouse School of Economics, CNRS, Toulouse. E-mail: bruno.jullien@tse-fr.eu.

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

inviting third parties to sell their products or services on top of the original product the company already sells. Two well-known and successful examples are Apple’s iPhone and Salesforce’s customer relationship management (CRM) software. After launching the iPhone in 2007 as a stand-alone product with no support for third-party apps, Apple quickly realized it would benefit from the creativity of third-party developers. As a result, in 2008 the company turned the iPhone into a platform by opening it up to third-party apps. Salesforce was founded in 1999 as a seller of CRM software products to small-to-medium-size businesses. In 2005, the company created a platform (Force.com) and an app marketplace (AppExchange) around its offering, which allows third-party software developers to build and sell other software to Salesforce’s CRM customers. Today there are over one million Force.com registered developers and over 2,500 apps offered on AppExchange.

When the third-party products are complementary (or unrelated) to the original product, the benefits of inviting them to sell on top of the original product are obvious. However, some of the third-party products may actually be (partial) substitutes to the original product, in which case the benefit of “hosting” the third-parties is not so obvious. This motivates our paper, which explores the conditions under which hosting a third-party that produces a rival product or service can be profitable.

In our view, there are many potential opportunities for existing companies to turn themselves into platforms by hosting rivals. Many of these opportunities are hypothetical for the time being, since the firms involved have yet to explore them. There are, nonetheless, a good number of cases where existing firms have successfully completed or at least embarked on the transition. In the two examples mentioned above, Apple allows some apps which compete with functionality already existing in the iPhone: e.g. Google Maps competes with Apple Maps, Google Chrome competes with the Apple’s Safari browser which is pre-installed. Salesforce’s AppExchange allows customers to purchase some third-party apps that directly compete with functionality included in Salesforce’s CRM product: e.g. Survey Monkey, GetFeedback and QuestionPro compete for consumer surveys, or CongaGrid and GridBuddy for data management.

Similarly, consider Intuit, the seller of QuickBooks, which is the leading software for accounting, financial management and tax compliance for small businesses in the United States. Over the past five years, Intuit has progressively turned QuickBooks into a multi-sided platform. Specifically, the company opened up application-programming interfaces, created a developer program, and launched an app store, all of which allow third-party developers to build and sell software products to QuickBooks’ customer base. Some of these products compete with features already included in QuickBooks: for instance, third-party payroll management apps such as TimeTracker and TimeRewards are direct substitutes for Intuit’s own TSheets. Intuit has also created QuickBooks Capital, a marketplace where both Intuit and a number of selected third-party lenders can offer loans to QuickBooks customers. Intuit negotiates attractive rates with the third-party lenders and makes it easier for QuickBooks customers to apply for loans directly from QuickBooks by directly providing the relevant information to lenders.¹

In many countries Cable and satellite TV providers have allowed Netflix to sell to their subscribers

¹See <https://quickbooks.intuit.com/capital/>

through their own platforms, even though Netflix competes with the cable companies’ video-on-demand services.² In the financial sector, a European company named “Open Banking” is facilitating the movement of banking to a platform model. Open Banking provides a software infrastructure that allows banks to offer third party deposit products to their own customers through their existing accounts. As an example, Deutsche Bank offers its German account holders the chance to access fixed deposits of selected rival banks through its “ZinsMarkt”.

Gyms provide a more “physical” example. Recently, some “big box gyms” have begun renting out space in their facilities to specialty studios, where the latter can offer classes to the gym’s members. For instance, the New York Sports Club (NYSC) hosts cycling classes offered by Cyc Fitness, a boutique cycling studio, within several of the NYSC’s gym locations in New York City.³ Country clubs work in a similar way, sometimes hosting third-parties that provide classes or specialized services to the clubs’ members (e.g. swimming or tennis coaching), where these were previously (or sometimes still are) provided by the clubs.

We provide a simple model which captures some of the key tradeoffs that arise when a firm decides whether to turn itself into a platform by hosting a (partial) competitor. In the model there is a multiproduct firm M that provides two types of products A and B , and a specialized firm S that just offers a superior version of B . There are two types of consumers, some who just want product A and some who want both products. Consumers incur a shopping cost of going to each firm, and have the option to go to both (i.e. multi-stop shop). In this model, if M “hosts” S , it means that consumers can go to M and buy any subset of product A , M ’s version of product B and S ’s version of product B , while incurring the shopping cost only once.

Modelling a platform as arising endogenously to save customers’ shopping costs captures a fundamental role played by real-world platforms, which is to provide a common infrastructure that allows multiple products or services to be produced and/or sold to customers. Our focus is then on the strategic interaction between the potential platform owner and the specialized firm, which determines whether or not the multiproduct firm will indeed become a platform.

By eliminating the additional shopping cost consumers incur when they wish to buy A from M and B from S , hosting eliminates M ’s ability to price discriminate across the two types of consumers. Essentially, when S is hosted, M unbundles the A and B products, so hosting turns competition for the market into competition *within* the market. This means M can no longer make a profit by selling B , given that S offers a superior version of B and the firms now compete on a level playing field in B . On the other hand, hosting allows M to potentially gain by raising its price on product A because shopping costs are now taken care of by the surplus offered by S ’s superior version of B . In this sense, hosting allows M to gain by turning a substitute into a complement. This logic extends to the joint profit perspective, which becomes relevant when the two firms can make lump-sum transfers to each other under hosting. Taking into account both firms’ profits expands the region of parameters where hosting dominates, but the tradeoff and its underlying logic remain.

²For instance, Sky’s Sky Q service in the UK and Italy provides users access to Internet video and music services (YouTube, Netflix and Spotify) that compete with Sky Q’s own offerings.

³<https://www.newyorksportsclubs.com/cyc>

If we also allow for the possibility that M can monitor S 's transactions under hosting and charge a variable (per transaction) fee, then hosting always dominates non-hosting from a joint profit perspective in the absence of fixed hosting costs. The reason is that the variable fee provides a new instrument for M to extract surplus from consumers who buy S 's superior version of B . Although it is constrained by outside competition, this instrument is sufficient to make up for the loss of its ability to price discriminate using the price of its own version of B under hosting.

Non-hosting can dominate from a joint profit perspective even if M can charge a variable fee and even in the absence of fixed hosting costs once we allow for the possibility that some consumers are not aware of the specialist's existence. If the number of these consumers is sufficiently high, they can result in competition being softened in the absence of hosting as M focuses more on exploiting uninformed consumers than just competing for informed consumers. On the other hand, if uninformed consumers learn of S 's existence and prices whenever they visit M , the firms' profits under hosting do not change compared to the full-information setting. Thus, firms may prefer not to host in order to prevent too many consumers becoming informed about their rivals.

Finally, we consider what happens when there are multiple competing specialists and the value added by each specialist under hosting is uncertain. This allows us to determine the optimal number of specialists that M wants to host. In case shopping costs are not too high and M cannot use variable fees, we find M will want to limit the number of specialists it hosts given that whenever it hosts more than one specialist that turns out to offer this additional value, the specialists will compete it away.

2 Related literature

Our paper relates to several strands of literature.

Some existing papers have analyzed the tradeoff between the platform business model and more traditional alternatives: marketplace vs. reseller in Hagiu and Wright (2015a), platform vs. vertically integrated firm in Hagiu and Wright (2015b and 2018), agency vs. wholesale pricing in Abhishek et al. (2016) and Johnson (2017). In these papers, the main difference between the platform and the traditional business model is the allocation of control over the key factors that are relevant for customers (e.g. prices, marketing decisions, product delivery, etc.). A distinction relative to the current paper is that this literature does not allow the same product or service to be offered by the firm in competition with its agents (suppliers or professionals). Thus, this strand of literature does not address the issue of a "traditional" firm hosting rivals to become a platform, the central question of the current paper.

Somewhat closer is the literature discussing a firm's decision to open itself up to third-party developers. The key issues that this literature has focused on are how much technology to share with platform participants (Boudreau, 2010, Parker and Van Alstyne, 2018, Niculescu et al., 2018), whether to make a piece of software open source (August et. al, 2013 and 2018), and whether a firm should become a "platform" and collaborate with producers of complementary products (Mantovani and Ruiz-Aliseda, 2016). In these papers, the main benefit of opening up to third-parties is to encour-

age innovation and contributions by outside developers or complementors, who are not rivals of the platform. The current paper is different in that it shows “opening up” can make sense even if the third-parties are rivals.

There is a small literature studying whether a platform that caters to third-party providers should offer its own products/services. For example, Hagi and Spulber (2013) study a platform’s incentive to introduce first-party content alongside third-party content. They show that doing so can be beneficial in mitigating the chicken-and-egg problem coordination problem in user participation. Zhu and Liu (2018) study this question in the context of Amazon, showing empirically that Amazon is more likely to compete with its marketplace sellers in product categories that are more successful (in terms of sales). Relative to these papers, the current paper does the reverse: it studies whether a product firm should introduce third-party sellers on its own platform. Moreover, neither of these papers models platforms as reducing shopping costs.

To some extent, the platform as modeled in our paper can be viewed as a vertically integrated firm that uses the upstream input (product A in our model) to offer downstream products (the product A and the various versions of product B). The vertically integrated firm can consider selling access to its upstream facility to rival downstream firms (S in our model). The literature on vertical foreclosure has studied incentives to provide such access when the upstream facility is essential for downstream firms to sell in the downstream market and when the upstream firm charges tariffs to the downstream firms for access (see Rey and Tirole, 2007, for a summary). Our setting is different in several respects. First, hosting does not reduce S ’s cost in our model, but rather allows consumers to save on shopping costs. Second, the platform is not essential, so the hosted firm can still sell outside the platform. Third, the firm providing the platform is a multiproduct provider with market power in both goods. Because of the shopping cost, hosting transforms the competing specialist firms into complementors to the monopolized good (product A), whereas in a standard vertical setting, access to the input increases competition. For this reason, hosting may be profitable even without financial compensation from the hosted firm or any wholesale contract.

The empirical study of Facebook’s integration of Instagram by Li and Agarwal (2017) has a similar flavour of turning a rival into a complementor, but this is done through outright acquisition. By contrast, when a platform hosts a rival in the current paper, the rival maintains pricing autonomy (or more generally other forms of control)—this autonomy is fundamental to being a platform as opposed to a vertically integrated firm. Furthermore, the key driving force in Li and Agarwal (2017) is the consumption complementarity between Facebook and Instagram, whereas in the current paper it is the shopping cost saving. In a similar vein, Eisenmann, Parker and Van Alstyne (2011) interpret vertical integration decisions by platforms as bundling, which increases convenience for users. Although the bundling theme is similar to our paper, they focus on bundling with complementors, not with rivals.

Finally, our paper belongs to a burgeoning economics literature which explores the implications of competition with multiproduct firms when, due to shopping costs, consumers have a demand for one-stop shopping. In this literature, like our paper, different products sold at the same firm become de facto complementary to one another due to consumers wanting to save on shopping costs. Important

early contributions to this literature are Lal and Matutes (1989, 1994), and Chen and Rey (2012), who focused on the implications for loss-leader pricing. Recent works such as those by Zhou (2014) and Rhodes and Zhou (2018) have focused on the implications of one-stop shopping for pricing, bundling and product range decisions in search contexts. We are the first to focus on the implications of one-stop shopping for hosting a rival’s product.

3 Model

We start with a simple benchmark model. There are two types of products, A and B . Suppose there is a multiproduct firm M that offers both product A and its version of B , denoted B_M , and a specialized firm S that just offers its version of product B , denoted B_S . This means product A is monopolized by M while product B can be supplied by M or S . We normalize both firms’ costs to zero. To keep things concrete, we will illustrate the model setup with the gym example described earlier. Thus, A can be thought of as the gym’s core offering that is included in the membership, and B as a specialized class that can be offered by the gym M or by a specialist firm S (i.e. Cyc).

The total measure of consumers is normalized to one. Among them, there are two types. A fraction $\lambda_A > 0$ of consumers just want to purchase one unit of A and are not interested in B (i.e. they value both versions of B at zero). In the gym example, these are consumers who just want to use the gym’s core facilities and are not interested in specialized classes. We call these monoprodukt consumers “ A -type” consumers. The remaining fraction $\lambda_B > 0$ (which equals $1 - \lambda_A$) of consumers want to purchase one unit of A and one unit of B . We call these multiproduct consumers “ B -type” consumers. All consumers value product A at $u_A > 0$.⁴ The B -types value B_M at $u_B > 0$ and B_S at $u_S = u_B + \Delta$, where $\Delta \geq 0$. Thus, B -types view B_S as superior to B_M .⁵

All consumers incur a shopping cost $\sigma > 0$ when going to each firm, regardless of how many products they buy from it. Thus, if consumers go to both M and S (i.e. multi-stop shop), they will incur σ twice. Consumers can always purchase an outside option which gives them a payoff normalized to zero. Throughout the paper we assume that

$$\sigma < \min \{u_A, u_B\},$$

i.e. the shopping cost is low enough that M could potentially sell either product alone. We also make the additional assumption that

$$\Delta \leq \sigma,$$

i.e. the shopping cost exceeds the added value of S ’s product B_S . This implies that without hosting, M can have an advantage in selling product B , provided it makes it attractive for consumers to want

⁴In Section A of the Online Appendix we relax this assumption to allow positive or negative correlation between the values different types of consumers place on products A and B , showing the benchmark tradeoff between hosting and non-hosting is similar.

⁵Our results do not depend crucially on the assumption that B -type consumers are all the same. In Section B of the Online Appendix, we consider the variation of our baseline model in which B -type consumers have heterogeneous tastes over products B_M and B_S , and obtain a similar tradeoff between hosting and non-hosting.

to buy A . In Section 4.4, we briefly discuss what happens in the less interesting case in which $\Delta > \sigma$.

4 Benchmark results

In this section we analyze the above model, first in the case without hosting (Section 4.1), and then in the case when M hosts S , so that consumers can buy A and B_S at M without incurring the shopping cost twice (Section 4.2). We then determine M 's incentives to host S , first without any transfers or fees (Section 4.3), next when a lump-sum transfer payment can be made between the firms as part of the hosting contract (Section 4.4) and finally, when M can monitor transactions using S and so can use both a lump-sum transfer and variable fees in the hosting contract (Section 4.5).

4.1 Without hosting

In this section, we characterize the equilibrium that arises without hosting, i.e. consumers must go to S if they want to buy B_S . In the gym example, this means that Cyc Fitness only sells classes in its own studios, whereas the NYSC offers its own cycling classes in its gyms.

Consider first A -type consumers who are only interested in A . If M charges a price of p_A , these consumers will buy A provided $p_A \leq u_A - \sigma$. Now consider B -types. If M charges a price of p_B and S charges a price p_S , they have four relevant options:

- buy A and B_M , obtaining utility $u_A + u_B - p_A - p_B - \sigma$
- buy A and B_S , obtaining utility $u_A + u_B + \Delta - p_A - p_S - 2\sigma$
- buy B_M only, obtaining utility $u_B - p_B - \sigma$
- buy B_S only, obtaining utility $u_B + \Delta - p_S - \sigma$

Because the shopping cost σ outweighs S 's added value in product B (i.e. Δ), there is a unique equilibrium outcome in which M makes all the sales.⁶ Formal proofs for this result and others are provided in the appendix.

Proposition 1 *In the baseline model without hosting there is a unique equilibrium outcome in which prices are $p_A^* = u_A - \sigma$, $p_B^* = \sigma - \Delta$, $p_S^* = 0$. All A -type consumers purchase A , and all B -type consumers buy A and B_M from M . Profits are $\pi_M^* = u_A - \sigma + \lambda_B(\sigma - \Delta)$ and $\pi_S^* = 0$.*

Some comments are in order. In equilibrium, B -types choose to buy both products from M because (i) avoiding the additional shopping cost σ of multi-stop shopping is worth more to B -types than getting the higher utility from S 's better version of B , and (ii) getting the additional utility from A is worth more to B -types than getting the higher utility from S 's better version of B . Furthermore,

⁶Strictly speaking, there are other equilibria, in which $p_S^* < 0$. We rule out such equilibria because they involve S setting a price that it would prefer to change if some consumers actually purchased from it (i.e. off the equilibrium path).

M 's equilibrium prices for A and B_M are such that the net surplus B -types derive from buying A and B_M exactly matches the surplus they get from the two next best alternatives: buying A from M and B_S from S , or buying only B_S from S . Note that S 's presence constrains the amount that M can extract from selling its two products to B -types to $u_A - \Delta$.

Given that B -types buy both products from M , M collects p_A from A -types and $p_A + p_B$ from B -types, which means it can set its best price for A -types (i.e., $p_A^* = u_A - \sigma$) separately from its best (competitive) price for B -types (i.e., $p_A^* + p_B^* = u_A - \Delta > u_A - \sigma$, so $p_B^* = \sigma - \Delta$). This achieves the same outcome as if M could use third-degree price discrimination, which is possible because B -types always buy both A and B_M from M due to the high shopping cost σ .

Finally, the equilibrium in the proposition still holds even if $\lambda_A = 0$, so that there are no A -types. However, in that case equilibrium prices are not uniquely defined. Nevertheless, all equilibria result in the same profits. Specifically, M could either choose (i) $p_A = u_A - \sigma$ and $p_B = \sigma - \Delta$ as in Proposition 1 or (ii) $u_A - \sigma < p_A \leq u_A$ and $p_B = u_A - \Delta - p_A$, in which case consumers only want to buy A if they also buy B_M , and they compare buying A and B_M with just buying B_S from S . In this context, adding some A -types creates a new role for M 's price p_A and eliminates the range of equilibria in which $u_A - \sigma < p_A \leq u_A$.

4.2 Hosting

Now suppose S is hosted by M , meaning B -types can buy B_S from S through M without incurring the additional shopping cost σ . In the gym example, this means that the NYSC now hosts cycling classes offered by Cyc Fitness, so a NYSC member interested in Cyc's classes does not have to go to a separate Cyc studio. We still allow S to sell directly, at price p_S (Cyc Fitness did not stop offering classes in its studio after being hosted by the NYSC).⁷ Meanwhile, let \hat{p}_S denote the price S charges when it sells B_S through M .

We assume there is a fixed cost of hosting, denoted $F > 0$. For instance, hosting a specialized cycling class in a gym may require re-arranging and customizing the space with the relevant equipment and branding, as well as updating software systems for scheduling and reservations to include the specialized class. Similarly, for a bank there could be significant system costs (software, compliance, training) of allowing rival providers to sell their term deposits to its customers. Any negotiating and legal costs associated with writing a hosting contract would also be included in F , as should anticipated costs of integrating systems and employees more generally. In practice, both M and S may incur such costs. Since throughout most of the paper we will focus on the solution in which a lump-sum transfer can be made between the two firms (i.e. through a fixed fee), it will make no difference which firm actually incurs the fixed costs of hosting. Thus, for convenience, we will assume F is always incurred by M .

In equilibrium, S will only sell through M , so consumers will never multi-stop shop. The reason is that selling directly has the disadvantage of having B -type consumers incur an additional shopping cost σ or foregoing the additional utility $u_A - p_A$ of being able to purchase A on M . Thus, selling

⁷The hosting equilibrium we derive in this section remains valid even if S no longer sells B_S directly.

directly is less profitable for S than selling through M .

By removing the additional shopping cost for consumers interested in both A and B , hosting puts both firms on an even playing field when competing to make sales of B . Since $\Delta > 0$, in equilibrium S always wins this competition, and sells B_S to all B -types at $\hat{p}_S = \Delta$, while $p_B = 0$. On the other hand, this leaves M free to sell A to both types of consumers, without worrying about how this affects consumers' willingness to buy from it versus S . Thus, one can think of hosting as leading M to unbundle A and B_M . In this case, A -types buy A provided $p_A \leq u_A - \sigma$, and B -types will buy A provided that $p_A \leq u_A$ and that they want to go to M in the first place, which they do since they obtain a surplus of

$$u_A + u_B + \Delta - p_A - \hat{p}_S - \sigma = u_A + u_B - p_A - \sigma \geq u_B - \sigma > 0.$$

Thus, M has two options. It can either set $p_A = u_A - \sigma < u_A$ and sell A to all consumers, obtaining $\pi_M = u_A - \sigma$, or set $p_A = u_A$ and sell A only to B -types, obtaining $\pi_M = \lambda_B u_A$. Using that $\lambda_A = 1 - \lambda_B$, we obtain the following result.⁸

Proposition 2 *In the baseline model with hosting, there are two cases to consider:*

- If $\lambda_A \leq \frac{\sigma}{u_A}$, the hosting equilibrium with the highest profit for M involves prices $p_A^* = u_A$, $p_B^* = 0$, and $\hat{p}_S^* = p_S^* = \Delta$. The A -types do not purchase, while the B -types all buy A and B_S through M . Profits are $\pi_M^* = \lambda_B u_A - F$ and $\pi_S^* = \lambda_B \Delta$.
- If $\lambda_A > \frac{\sigma}{u_A}$, the equilibrium prices are $p_A^* = u_A - \sigma$, $p_B^* = 0$, and $\hat{p}_S^* = p_S^* = \Delta$. The A -types purchase A , and the B -types all buy A and B_S through M . Profits are $\pi_M^* = u_A - \sigma - F$ and $\pi_S^* = \lambda_B \Delta$.

The margin Δ that S obtains on B -types reflects that under hosting, with no shopping cost disadvantage, S has a competitive advantage of Δ in selling B , which it can fully extract. By contrast, recall that without hosting, S was at a shopping cost disadvantage and had to compete against the bundle of A and B_M , which prevented it from making any profit.

Under hosting, if M could engage in third-degree price discrimination, it would want to charge $u_A - \sigma$ to A -types and u_A to B -types (indeed, the B -types' shopping costs are now covered by the surplus offered by B_S). However, given that S now competes and wins the market for B on the platform created by M , such price discrimination is no longer possible. This drives a tradeoff between hosting and non-hosting, which we will explore in the next section.⁹

⁸Strictly speaking, here too there are other equilibria. These involve $p_B^* < 0$ given that M does not sell B_M in equilibrium, and equilibria in which $p_S^* < \hat{p}_S^*$ given that S does not sell B_S directly in equilibrium. Among all the possible equilibria we focus on the best equilibrium for M (which also turns out to be the equilibrium that maximizes joint profits). This avoids equilibria in which firms set prices (specifically, p_B and p_S) such that they would refuse to sell if some consumers actually asked to purchase from them (i.e. off the equilibrium path).

⁹If instead M only sold a bundle of A and B_M , it would choose between setting $p_{AB} = u_A - \sigma$ and selling the bundle to all consumers, or setting $p_{AB} = u_A$ and selling the bundle to B -type consumers only. In either case it obtains no more profit from bundling than in the equilibrium described in the proposition above.

An implicit assumption in our analysis above is that M does not remove B_M when hosting S . In any proposed equilibrium in which M does not compete by trying to sell B_M (so that S has a monopoly over product B), clearly M can do better by offering B_M . Thus, the only way M would not offer B_M is if it could commit *ex-ante* to not offer it. This would always increase joint profits in our baseline model. However, this may not be possible in practice if it requires M to write a contract specifying that it will not compete with S on M 's platform. Indeed, this type of contract would likely raise antitrust concerns because it could be viewed as a form of collusion. A commitment to remove B_M may therefore require a technological commitment, which may not always be feasible. Even if such a commitment is feasible, it may not be jointly profitable once firms take into account the realistic possibility that other firms would then want to enter to sell B . This is indeed the situation when there are multiple specialists competing, which we consider in Section 5.2.

4.3 Unilateral incentive to host

We initially consider whether M is better off with hosting or without hosting, while ignoring the possibility of any transfer payments between the firms. This allows us to provide some initial intuition about the tradeoffs associated with hosting. We obtain the following result.

Proposition 3 *In the baseline model without any transfer payments between the firms, hosting is preferred by M if and only if $\lambda_A \leq \frac{\sigma}{u_A}$ and $\Delta > \frac{\lambda_A(u_A - \sigma) + F}{1 - \lambda_A}$, and hosting is always preferred by S .*

It is easiest to interpret this result when there are almost no A -types, i.e. when $\lambda_A \rightarrow 0$. Then, hosting allows M to increase its profit by Δ , so hosting is preferred provided Δ exceeds the fixed cost F of hosting. The gain of Δ that M obtains from hosting comes from a gain of σ on the A product and a loss of $\sigma - \Delta$ on the B product (which is smaller). Indeed, hosting allows M to charge u_A for A instead of $u_A - \sigma$, because shopping costs are now taken care of by S through the surplus obtained from B_S . This is the sense in which hosting S (and thereby eliminating the shopping cost necessary to access B_S) allows M to gain by turning a substitute into a complement. On the other hand, under hosting M no longer extracts $\sigma - \Delta$ from its sale of B_M as sales of B are now made by S . Thus, turning competition for the market into competition within the market means that M gives up on its profit in the B market. Put differently, hosting unbundles the products and levels the playing field in product B competition, which means M can no longer make a profit on B . In contrast, S can now extract the profit $\lambda_B \Delta$, selling B_S to B -types, which is why it strictly prefers hosting.

Now consider what happens when there are some (but not too many) A -types, so $\lambda_A \leq \frac{\sigma}{u_A}$. In this case, if M charges u_A for A instead of $u_A - \sigma$, it loses the A -types, who no longer purchase. This means the additional surplus extracted by M from product A under hosting may no longer dominate the negative effect of hosting on M 's profit in the B market. Put differently, the presence of A -types constrains M 's ability to extract more from product A under hosting since A -types do not care about the extra surplus generated from hosting S 's superior version of product B . Thus, M can be worse off under hosting even without taking into account the fixed costs of hosting. This happens when the loss of B -type sales under hosting (which recall is equal to $\sigma - \Delta$) is large, i.e. when Δ is small.

Finally, consider the case when there are many A -types, so $\lambda_A > \frac{\sigma}{u_A}$. Then M will not want to increase p_A at all as a result of hosting since it does not want to give up on selling to the A -types, and so there is no gain on product A from hosting to offset the loss on product B . In this case, hosting always lowers M 's profit, reflecting M 's inability to price discriminate across the two types of consumers by charging more to B -types through p_B when it hosts.

The logic here bears some similarity with that of bundling. Shopping costs give to the multi-product firm M a competitive edge over single-product firms whenever products A and B are bought together by the consumer. At the same time, however, non-hosting puts competitive pressure on the monopolized product A , just like bundling does (see Whinston, 1990). Hosting (like unbundling) relaxes this competitive pressure.

4.4 Joint incentives to host without monitoring

So far we have ignored any transfer payments that could be made between the firms. Suppose M cannot monitor sales by S and charge for them, which could be because the monitoring technology is too costly to implement or because S does not want to share customer transaction data with M . Suppose, however, that firms can make lump-sum transfers. Then hosting will arise whenever the two firms can be made jointly better off with hosting, after taking into account the fixed costs of hosting. Throughout the rest of the paper we will focus on the impact of hosting on the firms' joint profit. For our baseline setting, we obtain the following result.

Proposition 4 *In the baseline model with no variable fees in the hosting contract, hosting is jointly preferred if and only if (i) $\lambda_A \leq \frac{\sigma}{u_A}$ and $\Delta > \frac{\sigma}{2} + \frac{F - (\sigma - \lambda_A u_A)}{2(1 - \lambda_A)}$, or (ii) $\lambda_A > \frac{\sigma}{u_A}$ and $\Delta > \frac{\sigma}{2} + \frac{F}{2(1 - \lambda_A)}$.*

Comparing Proposition 4 with Proposition 3, it is clear that the tradeoff between hosting and non-hosting in terms of joint profits is similar to the one when we focused only on M 's profit. The only difference is that now under hosting we must add S 's profit $\lambda_B \Delta$, which expands the region of parameter values for which hosting dominates. Specifically, if $\Delta > \frac{\sigma}{2}$, then hosting raises joint profits, so in this case it always dominates non-hosting for F sufficiently small.¹⁰ This extends to the case we ruled out by assumption, namely $\Delta > \sigma$, i.e. S 's efficiency advantage is so high that it more than offsets the shopping cost advantage of M . In that case, aside from the fixed cost F , there is no downside to hosting given that M does not sell B either way (with or without hosting). In particular, it is easily shown that, when $\Delta > \sigma$, hosting is jointly preferred if $\lambda_B \sigma + \max\{\sigma - \lambda_A u_A, 0\} > F$.

The tradeoff between hosting and non-hosting in Proposition 4 shifts towards hosting when Δ increases and towards non-hosting when λ_A , F or u_A increase. When σ increases, the tradeoff shifts towards hosting if $\lambda_A \leq \frac{\sigma}{u_A}$ and towards non-hosting if $\lambda_A > \frac{\sigma}{u_A}$. To understand why the tradeoff is non-monotonic in the level of the shopping cost σ , note that without hosting, because M is able to price discriminate, the amount it collects is lowered by the shopping cost for the A -types only (i.e. by $\lambda_A \sigma$). On the other hand, under hosting, if M wants to sell to A -types, it has to lower its price

¹⁰The impact of hosting on consumer surplus and welfare is explored in Online Appendix C.

by σ on all consumers. This means when there are relatively many A -types, so they are served by M under hosting, a higher σ shifts the tradeoff in favor of non-hosting. On the other hand, when there are relatively few A -types, so M gives up selling to them under hosting, a higher σ shifts the tradeoff in favor of hosting.

4.5 Joint incentives to host with monitoring

Suppose now that M is able to monitor S 's sales through M and charge for them, which we assumed was not possible previously. This means M can also set a per-transaction fee (or variable fee) τ in addition to a lump-sum fee in the hosting contract, so that S pays τ to M for each unit it sells on M .

The timing remains as before: after the contract is specified (including variable and lump-sum fees), the two firms set their prices simultaneously, taking the variable fee τ specified in the contract as given. The pricing game given τ turns out to have multiple equilibria: to keep the analysis streamlined, we always select the equilibrium that maximizes joint profits of M and S for every given τ . In the appendix we prove the following result.

Proposition 5 *In the baseline model, when M can specify a variable fee τ to charge S and a lump-sum transfer in the hosting contract, the optimal variable fee is always $\tau^* = \sigma$ and*

- *If $\lambda_A \leq \frac{\sigma}{u_A}$, then the highest joint profit that can be achieved in the hosting equilibrium is $\lambda_B(u_A + \sigma + \Delta) - F$, so hosting is jointly preferred if and only if $\Delta > \frac{F - (\sigma - \lambda_A u_A)}{2(1 - \lambda_A)}$.*
- *If $\lambda_A > \frac{\sigma}{u_A}$, then the highest joint profit that can be achieved in the hosting equilibrium is $u_A - \sigma + \lambda_B(\sigma + \Delta) - F$, so hosting is jointly preferred if and only if $\Delta > \frac{F}{2(1 - \lambda_A)}$.*

Comparing Proposition 5 with Proposition 4, it is straightforward to see that the possibility of using a variable fee unambiguously shifts the tradeoff towards hosting. Hosting S and charging it a variable fee $\tau = \sigma$ allows M to preserve the same competitive edge as without hosting while relaxing the competitive pressure on good A .¹¹ Thus, joint profits under hosting increase by exactly $\lambda_B \sigma$ as a result of using variable fees. As a result, we now need $F > 0$ for the tradeoff to be non-trivial: if $F = 0$, then hosting is always jointly preferred when M can charge a variable fee. The reason is that the variable fee τ is another instrument that M can use to price discriminate between A -types and B -types whenever M chooses to keep selling to A -types: τ allows M to extract the increase in surplus offered to B -types by S 's superior product. Thus, τ makes up for the loss of the ability to price discriminate using the price p_B for B_M . Meanwhile, M can use p_B to control any double marginalization problem that would otherwise arise with variable fees. Nevertheless, with $F > 0$, the key tradeoff determining whether hosting will be chosen is still increasing in Δ and decreasing in λ_A , as in the case without variable fees.

¹¹Note, however, that outside competition constrains the variable fee to be no more than σ , which prevents M from achieving the vertically integrated monopoly solution.

The tradeoff between hosting and non-hosting in Proposition 5 shifts towards hosting when Δ increases or when σ increases if $\lambda_A \leq \frac{\sigma}{u_A}$, and towards non-hosting when λ_A , F or u_A increase. The only qualitative difference compared to the case without variable fees is that in the parameter range $\lambda_A > \frac{\sigma}{u_A}$, the tradeoff no longer changes when σ increases. The reason is that now, under hosting, M can use the variable fee to extract an additional margin of σ from B -types, which is the same margin S could extract under non-hosting.

5 Extensions

In this section we consider two important extensions of the baseline model. The first extension shows that once some consumers have imperfect information, hosting specialists may not always dominate, even if variable fees can be used and even if there are no fixed costs associated with hosting. The second extension allows for more than one specialist and addresses the issue of how many specialists M should host by assuming there is uncertainty over the specialists' added value.

5.1 Hosting as information

So far we have assumed all consumers know about the existence of the specialist firm, but this is of course not always the case. In fact, one of the key benefits of being hosted on a platform for a specialist is to increase awareness of its services among the platform's customers. For example, it is likely that before hosting, some members of NYSC will not have heard of Cyc Fitness. However, after hosting, by seeing Cyc located within NYSC, members will become aware of its offerings, and, if they have an interest in cycling classes, will also find out its prices. To capture this situation within our model, suppose that without hosting, a fraction $\eta \in (0, 1)$ of B -type consumers are not informed about the specialist's existence. Once the specialist is hosted, consumers learn of its existence and prices whenever they visit M . In particular, we assume that when consumers find out about S 's existence upon visiting M , they also learn about its direct channel (outside M) and price charged there.

We start by considering the case without hosting. The co-existence of informed and uninformed consumers implies that for some parameter range there is no equilibrium in pure strategies. In that case, we determine the (unique) equilibrium in mixed strategies. The following proposition summarizes the outcome without hosting.

Proposition 6 *When a fraction η of B -type consumers are uninformed of S 's existence, the non-hosting equilibrium is determined as follows:*

- For $\eta \leq \frac{\sigma - \Delta}{u_B}$, the unique equilibrium is in pure strategies and is identical to the benchmark case (see Proposition 1).
- For $\eta > \frac{\sigma - \Delta}{u_B}$, the unique equilibrium is in mixed strategies and involves $p_A = u_A - \sigma$, p_B drawn from the CDF

$$G_B(p_B) = \frac{p_B - \eta u_B}{p_B - \sigma + \Delta}$$

with support $p_B \in [\eta u_B, u_B]$ and a mass point at $p_B = u_B$,

$$\Pr(p_B = u_B) = \frac{\eta u_B + \Delta - \sigma}{u_B + \Delta - \sigma} < 1,$$

p_S drawn from the CDF

$$G_S(p_S) = 1 - \frac{\eta}{1 - \eta} \frac{u_B + \Delta - \sigma - p_S}{p_S + \sigma - \Delta}$$

with support $p_S \in [\eta u_B + \Delta - \sigma, u_B + \Delta - \sigma]$. Expected equilibrium profits are

$$\begin{aligned}\pi_M &= u_A - \sigma + \lambda_B \eta u_B \\ \pi_S &= \lambda_B (1 - \eta) (\eta u_B + \Delta - \sigma).\end{aligned}$$

Proposition 6 implies that if the fraction of uninformed consumers is small, the existing benchmark equilibrium without hosting (Proposition 1) still applies. In this equilibrium M sells B_M to informed B -type consumers since avoiding the additional shopping cost of multi-stop shopping is worth more to informed B -types than the higher utility offered by S 's better version of B . Given the fraction η of uninformed consumers is small, M does not find it profitable to deviate by increasing the price of B_M so as to exclusively target uninformed consumers and give up selling B_M to informed consumers. On the other hand, if the fraction of uninformed consumers is sufficiently high, starting from this pure-strategy equilibrium, M would prefer to increase p_B all the way to u_B in order to extract the entire surplus of uninformed consumers, thus giving up on selling to the informed B -types. However, this cannot be an equilibrium since S would best respond by setting a higher price, which would then entice M to once again try selling to informed B -types. Instead, there is a unique mixed strategy equilibrium where both M and S randomize over the price of their respective versions of B , while M still sells A to all consumers by setting $p_A = u_A - \sigma$. In the parameter range on which this mixed strategy equilibrium prevails, both firms obtain higher expected profits than in the benchmark case: S obtains positive profits instead of zero, while M 's profits are higher by $\lambda_B (\eta u_B - (\sigma - \Delta)) > 0$. This is because the presence of uninformed consumers relaxes price competition.

Next consider the case with hosting. In principle, the fact that some consumers do not know about the specialist's existence before visiting M means that this scenario is somewhat different from the benchmark hosting case, where all consumers were informed of S 's existence and presence on M even before going to M . However, it turns out that this difference does not affect the analysis (since all consumers are induced to shop at M in any equilibrium, they all end up informed of S 's existence). Thus, the same hosting equilibrium as in the benchmark case prevails—both with and without variable fees. For convenience, the next proposition summarizes the joint profits under hosting for each case.

Proposition 7 *When a fraction η of B -type consumers are uninformed of S 's existence, the hosting equilibrium and firm profits are the same as in the benchmark case:*

- When M cannot monitor S 's sales and charge variable fees, if $\lambda_A \leq \frac{\sigma}{u_A}$, then joint profits

in equilibrium are $\lambda_B(u_A + \Delta) - F$, whereas if $\lambda_A > \frac{\sigma}{u_A}$, then joint equilibrium profits are $u_A - \sigma + \lambda_B \Delta - F$.

- When M can monitor S 's sales and charge variable fees, if $\lambda_A \leq \frac{\sigma}{u_A}$, then joint profits are $\lambda_B(u_A + \sigma + \Delta) - F$, whereas if $\lambda_A > \frac{\sigma}{u_A}$, then joint profits are $u_A - \sigma + \lambda_B(\sigma + \Delta) - F$.

We can now compare the outcomes under hosting and non-hosting using the previous two propositions. The first proposition below focuses on the case when variable fees cannot be used.

Proposition 8 *Suppose a fraction η of B -type consumers are uninformed of S 's existence under non-hosting, but they become informed about S under hosting if they visit M . When variable fees cannot be used:*

- If $\eta \leq \frac{\sigma - \Delta}{u_B}$, the conditions for hosting to be jointly preferred are identical to those in Proposition 4.
- If $\eta > \frac{\sigma - \Delta}{u_B}$ and $\lambda_A \leq \frac{\sigma}{u_A}$, hosting is jointly preferred if and only if

$$\Delta > \frac{F - (\sigma - \lambda_A u_A)}{\eta(1 - \lambda_A)} + \frac{(2 - \eta)(\eta u_B - \sigma)}{\eta} + \frac{\sigma}{\eta}.$$

- If $\eta > \frac{\sigma - \Delta}{u_B}$ and $\lambda_A > \frac{\sigma}{u_A}$, hosting is jointly preferred if and only if

$$\Delta > \frac{F}{\eta(1 - \lambda_A)} + \frac{(2 - \eta)(\eta u_B - \sigma)}{\eta} + \frac{\sigma}{\eta}.$$

The first case in Proposition 8 is identical to the benchmark analysis since, as explained above, the equilibrium analysis with hosting is unchanged and the equilibrium analysis without hosting is also unchanged when there are not many uninformed consumers. In the remaining two cases in Proposition 8, when the fraction of uninformed consumers is sufficiently high, the tradeoff unambiguously shifts towards non-hosting. To see this note that joint profit is the same under hosting, but the joint profits without hosting are higher in Proposition 6 when $\eta > \frac{\sigma - \Delta}{u_B}$ than they are in Proposition 1. When there are enough of them, the presence of uninformed consumers softens the competition for B -types without hosting, reflecting that M will sometimes exploit the uninformed B -types by setting a high price, and that S will best respond by also sometimes setting a high price p_S . In contrast, by promoting the specialist, hosting removes the friction that prevented S from reaching all consumers, thereby intensifying competition.

The effect of the factors $(u_A, \lambda_A, F, \sigma$ and $\Delta)$ on the tradeoff between hosting and non-hosting in Proposition 8 is qualitatively the same as in Proposition 4, with two exceptions when $\eta > \frac{\sigma - \Delta}{u_B}$. First, when $\lambda_A > \frac{\sigma}{u_A}$, the tradeoff now shifts towards hosting when σ increases, reflecting that the shopping cost σ limits the surplus that can be extracted from uninformed B -types when competition is relaxed.

Second, an increase in u_B now affects the tradeoff, shifting it towards non-hosting. This reflects that with higher u_B , the benefit of relaxing competition by focusing on exploiting uninformed B -types under non-hosting is higher since there is more surplus that can be extracted from such consumers.

Now consider the case when variable fees can be used by M under hosting. Then the joint profit comparison between hosting and non-hosting is given by the following proposition.

Proposition 9 *Suppose a fraction η of B -type consumers are uninformed of S 's existence under non-hosting, but they become informed about S under hosting if they visit M . When variable fees can be used:*

- If $\eta \leq \frac{\sigma - \Delta}{u_B}$, the conditions for hosting to be jointly preferred are identical to those in Proposition 5.
- If $\eta > \frac{\sigma - \Delta}{u_B}$ and $\lambda_A \leq \frac{\sigma}{u_A}$, hosting is jointly preferred if and only if

$$\Delta > \frac{F - (\sigma - \lambda_A u_A)}{\eta(1 - \lambda_A)} + \frac{(2 - \eta)(\eta u_B - \sigma)}{\eta}.$$

- If $\eta > \frac{\sigma - \Delta}{u_B}$ and $\lambda_A > \frac{\sigma}{u_A}$, hosting is jointly preferred if and only if

$$\Delta > \frac{F}{\eta(1 - \lambda_A)} + \frac{(2 - \eta)(\eta u_B - \sigma)}{\eta}.$$

The first case in Proposition 9 is identical to the benchmark analysis (see Proposition 5) since the equilibrium analysis with hosting is unchanged and the equilibrium analysis without hosting is also unchanged when there are not many uninformed consumers. In the second and third cases in Proposition 9, the tradeoff is once again unambiguously shifted in favor of non-hosting relative to the benchmark in Proposition 5. In particular, in this case non-hosting can be jointly preferred even though M can use variable fees and even if there is no fixed cost associated with hosting (i.e. if $F = 0$), provided η is high enough.

In terms of comparative static results, again we find that the effect of the factors (u_A , λ_A , F , σ and Δ) on the tradeoff between hosting and non-hosting in Proposition 9 is qualitatively the same as in the corresponding case with full information (Proposition 5), with the exception that when $\eta > \frac{\sigma - \Delta}{u_B}$ and $\lambda_A > \frac{\sigma}{u_A}$, the tradeoff now shifts towards hosting when σ increases. Moreover, again we find that when $\eta > \frac{\sigma - \Delta}{u_B}$, an increase in u_B shifts the tradeoff towards non-hosting. The rationale is the same as above.

5.2 Multiple specialists and uncertainty

Thus far, we have focused on the case of a single specialist S . Obviously, in reality there may be several competing specialists and M has a choice of how many of them to host. For instance, while Salesforce hosts several third-party apps for consumer surveys (e.g. Survey Monkey, GetFeedback,

QuestionPro) and data management (e.g. CongaGrid, GridBuddy) that compete with functionality included in Salesforce's CRM product, there are other commonly used survey apps (e.g. SurveyGizmo, KeySurvey, Praiseworthy) and data management apps (e.g. SAS) that are not hosted on AppExchange. Nor does Salesforce's platform host any competitor to its core CRM functionality, such as Pipeliner CRM.¹² Similarly, on its QuickBooks platform, Intuit hosts some third-party payroll management apps (e.g. TimeTracker and TimeRewards) that compete with Intuit's own Tsheets, but other prominent alternatives to Tsheets (e.g. Hubstaff, Toggl, Freckle) are not hosted.¹³

To capture this, suppose now there are $n \geq 2$ specialists who offer $u_B + \Delta$ when visited outside M . When any of these specialists is hosted, we also assume there is uncertainty regarding the value offered to M 's customers: for each hosted specialist, the value is $u_B + \Delta$ with probability θ and u_B with probability $1 - \theta$, where the realizations for different specialists are drawn independently and are the same for all consumers. This captures the general idea that there may be uncertainty over how the specialists will perform when hosted. Specifically, firms may be uncertain whether a given specialist's value added is specific to its location or carries over when it is hosted on M . This uncertainty is assumed to be resolved after the contract has been signed and the specialists have been hosted. Thus, if M decides to host one or multiple specialists, we assume M must commit to the fixed transfer and any variable fee τ before the uncertainty is resolved, while the firms set their prices (to consumers) after the uncertainty is resolved. Finally, we assume the fixed cost of hosting is F regardless of how many specialists are hosted. Note the case in which $\theta = 1$ captures the simple extension of the benchmark model to allow for multiple specialists (i.e. without uncertainty).

In the absence of hosting, allowing for multiple competing specialists does not change anything since in the benchmark case the specialist was already at a disadvantage when competing against M for sales of B (given $\Delta < \sigma$). Thus, all specialists price at zero and make zero profits, whereas M sets $p_A = u_A - \sigma$ and $p_B = \sigma - \Delta$, obtaining profits $u_A - \sigma + \lambda_B (\sigma - \Delta)$. This is the same outcome as in the benchmark case. With hosting, two differences arise: competition outside the platform drives prices to zero, and two hosted specialists offering the same value obtain zero profit. M must then decide how many specialists to host, and compare the resulting joint profits with the no-hosting outcome in deciding whether to host them.

The next two propositions characterize the hosting outcome and provide the conditions for hosting to be jointly preferred to non-hosting, first for the case when M cannot use variable fees under hosting (Proposition 10) and then for the case when M can use variable fees (Proposition 11).

Proposition 10 *Suppose there are $n \geq 2$ specialists who offer $u_B + \Delta$ when visited outside M , but when hosted offer $u_B + \Delta$ with probability θ and u_B with probability $1 - \theta$. When variable fees cannot be used:*

- *If $\sigma \leq \lambda_A u_A$, then M prefers to host $k^* = \arg \max_{k \in \{1, 2, \dots, n\}} \{k(1 - \theta)^{k-1}\}$ specialists, and*

¹²See <https://www.forbes.com/sites/christinecrandell/2014/12/06/salesforce-opens-the-door-to-competitors/#b54699f18c34>

¹³See <https://blog.hubstaff.com/tsheets-alternatives/>

hosting is jointly preferred if and only if

$$\Delta > \frac{\sigma(1 - \lambda_A) + F}{(1 - \lambda_A)(1 + \theta k^* (1 - \theta)^{k^*-1})}.$$

- If $\lambda_A u_A < \sigma < \lambda_A u_A + \lambda_B \Delta$, then M prefers to host

$$k^* = \arg \max_{k \in \{1, 2, \dots, n\}} \left\{ (1 - \theta)^k (\lambda_A u_A - \sigma) + k \theta (1 - \theta)^{k-1} (\lambda_A u_A - \sigma + \lambda_B \Delta) \right\}$$

specialists and hosting is jointly preferred if and only if

$$\Delta > \frac{\sigma(1 - \lambda_A) - \left(1 - (1 - \theta)^{k^*} - k^* \theta (1 - \theta)^{k^*-1}\right) (\sigma - \lambda_A u_A) + F}{(1 - \lambda_A)(1 + k^* \theta (1 - \theta)^{k^*-1})}.$$

- If $\sigma \geq \lambda_A u_A + \lambda_B \Delta$, then M prefers to host all n specialists and hosting is jointly preferred if and only if

$$\Delta > \frac{\lambda_A (u_A - \sigma) + F}{(1 - \lambda_A)(1 - (1 - \theta)^n)}.$$

By ignoring integer constraints and focusing on the case in which M hosts specialists, we can gain more insight into the optimal number of specialists to host. This is given by

$$k^* = \begin{cases} -\frac{1}{\ln(1-\theta)} & \text{if } \sigma \leq \lambda_A u_A \\ -\frac{1}{\ln(1-\theta)} + \frac{(1-\theta)(\sigma - \lambda_A u_A)}{\theta(\lambda_A u_A + \lambda_B \Delta - \sigma)} & \text{if } \lambda_A u_A < \sigma < \lambda_A u_A + \lambda_B \Delta \\ n & \text{if } \sigma \geq \lambda_A u_A + \lambda_B \Delta \end{cases},$$

assuming k^* belongs to $[1, n]$ in each case. The optimal number of specialists to host is (weakly) increasing in σ , and is decreasing in θ , Δ and u_A .

To understand the result, note first that if $\sigma \geq \lambda_A u_A + \lambda_B \Delta$, then shopping costs are so high that, regardless of how many hosted specialists turn out to be of high value, M does best to sell A only to B -type consumers, who have their shopping cost covered by buying the B product. In this case, all specialists make zero profits, while M 's profits are $\lambda_B u_A$, except in the case when no hosted specialist turns out to be of high value, when M 's profits are $\lambda_B (u_A - \Delta)$. Thus, joint profits are higher when at least one of the specialists hosted on M turns out to offer the added value Δ , because the hosted specialists free the platform from outside competitive pressure on good A that arises when the value of good B is larger outside than inside. Since the probability of at least one specialist turning out to offer high value is increasing in k , M will host all available specialists.

On the other hand, if $\sigma \leq \lambda_A u_A$, then shopping costs are so low that, regardless of how many hosted specialists turn out to be of high value, M does best to sell A to all consumers by setting the price $p_A = u_A - \sigma$. In this case, joint profits are the same (equal to $u_A - \sigma$) across the different realizations, unless exactly one hosted specialist turns out to offer high value—then the two firms can

extract the additional value Δ from B -types since it will not be competed away. Thus, M chooses k to maximize the probability $k\theta(1-\theta)^{k-1}$ of exactly one specialist offering high value. The optimal k is therefore decreasing in the probability θ that any individual specialist turns out to be of high value.

For intermediate levels of the shopping costs (i.e. $\lambda_A u_A < \sigma < \lambda_A u_A + \lambda_B \Delta$), M does best selling A only to B -types when two or more specialists turn out to be of high value. Otherwise, M does best selling A to all consumers. Thus, the only change relative to the case with low shopping costs is that now, when two or more specialists turn out to be of high value, joint profits are higher, equal to $\lambda_B u_A$ instead of $u_A - \sigma$. Since the probability of two or more specialists turning out to be of high value is increasing in k ,¹⁴ the optimal number of specialists to host is higher than in the case with low shopping costs, but still possibly below n .

Consider now the case when M can charge variable fees under hosting.

Proposition 11 *Suppose there are $n \geq 2$ specialists who offer $u_B + \Delta$ when visited outside M , but when hosted offer $u_B + \Delta$ with probability θ and u_B with probability $1 - \theta$. When variable fees can be used under hosting, M prefers hosting all n specialists and hosting is jointly preferred over non-hosting if and only if*

$$\Delta > \frac{F}{(1 - \lambda_A)(1 - (1 - \theta)^n)}.$$

Like in the benchmark case, hosting is always preferred if $F = 0$. This is not surprising: when variable fees can be used, whether only one or multiple hosted specialists have an efficiency advantage, M can always extract the efficiency gain σ due to hosting by setting $\tau = \sigma$.

The key difference relative to the benchmark model is that now, regardless of how many hosted specialists turn out to be of high quality, they have to compete with the outside specialists of high quality that price at cost. This explains why, going from the case with just one specialist (Proposition 5) to the case with two or more competing specialists (Proposition 11), the tradeoff shifts towards non-hosting. This holds even if $\theta = 1$, i.e. there is no uncertainty over the added value of the specialists.

The presence of the competing outside specialists constrains the price of hosted specialists. Specifically, the maximum price \hat{p}_S that can be charged by hosted specialist(s) of high quality is $\hat{p}_S = \sigma$ if $p_A = u_A - \sigma$ (so M sells A to both types), or $\hat{p}_S = 0$ if $p_A = u_A$ (so M sells A to B -types only). Clearly, in the latter case, M cannot extract any variable fee from the hosted specialists because they would make a loss, whereas in the former case, M can extract σ from them. Thus, it can never be profitable for M to only sell A to B -types by setting $p_A > u_A - \sigma$, so the only possible equilibrium now must involve M selling A to both types.

This has two implications. First, M 's profits are $u_A - \sigma + \lambda_B \sigma$ when at least one hosted specialist turns out to be of high quality and $u_A - \sigma + \lambda_B (\sigma - \Delta)$ when none of the hosted specialists turns out to be of high quality. Thus, M will host as many specialists as possible in order to maximize the

¹⁴Indeed, $1 - (1 - \theta)^k - k\theta(1 - \theta)^{k-1}$ is increasing in k .

chance that at least one specialist will have an efficiency advantage. Second, the hosting vs. non-hosting tradeoff is very similar to the case $\lambda_A > \frac{\sigma}{u_A}$ in the benchmark model, i.e. the case in which M sold to A -types in equilibrium. In particular, σ has no effect on the tradeoff and the effects of the parameters Δ , λ_A and F are the same as in the benchmark model. Furthermore, taking into account the fixed cost of hosting, the tradeoff shifts towards hosting when the number of competing specialists increases, i.e. when n increases.

6 Managerial implications

There are several factors that determine whether a multiproduct firm can gain by hosting a rival specialist to sell over its common infrastructure, thereby creating a platform.

Regardless of whether firms can make use of fixed transfers or variable fees based on monitoring transactions on the platform, our results imply that the multiproduct firm M should host a rival when the fraction of consumers who value both the core-product and the specialist product is high (i.e. there are many B -type consumers), when the added value of the specialist's version of the product is high (i.e. Δ is high), when the utility that consumers get from the core product is low (i.e. u_A is low), and when the fixed cost of hosting is low (i.e. F is low). If there are sufficiently many consumers who value both products, then hosting should be preferred when consumers benefit more from one-stop shopping (i.e. when σ is high). This helps explain for example why big box gyms (like the New York Sports Club) are increasingly willing to host specialty fitness studios (like the Cyc) under revenue sharing contracts: many users are interested in both standard gym amenities (e.g. weight equipment) and specialized classes (e.g. cycling), and there is a clear benefit from having both colocated. However, if there are a lot of consumers who are interested in the core product A only and M cannot monitor transactions and charge a variable fee, the tradeoff between hosting and non-hosting actually shifts towards non-hosting when shopping costs are high.¹⁵

Unsurprisingly, we find that the ability to monitor the rival's transactions on the platform, and so charge it a variable fee, makes hosting more profitable. With ongoing improvements in monitoring technologies, we therefore expect to see such hosting become more prevalent over time, i.e. a greater number of firms turning their products into platforms.

An important factor that makes hosting a less desirable choice is if many consumers are uninformed about the specialist firm in the absence of hosting. In this case, hosting can intensify competition between the firms by making consumers informed about the specialist rival's existence. Of course, in reality, consumers may already have partial information about the specialist rival before hosting, but it can still be that hosting makes them realize just how good the specialist is, which otherwise they would not know. More generally, to the extent hosting helps provide new and positive information about the specialist, that should make hosting less appealing for M . In the case of Salesforce, for example, their CRM customers might be skeptical of a newcomer like Pipeliner CRM (and they may

¹⁵To understand this last result, note that since M cannot price discriminate under hosting, it will have to lower its price by the shopping cost on all consumers if it wants to keep selling to A -types, which it does when there are a lot of them.

worry it would be hard to switch over to Pipeliner), but if Salesforce hosted Pipeliner, they would realize how easy it is to switch and use Pipeliner.

Finally, our results showed that M 's decision to host rivals also depends on the availability of multiple competing specialists. The existence of two or more competing specialists makes it less likely that M will host: indeed, hosting is no longer sufficient to avoid head-to-head competition because competition will continue to exist outside of the platform. In this case, the number of specialists to host depends on whether transactions on the platform can be monitored. If transactions can be monitored, then firms will want all available specialists to be hosted. In our setting, this maximizes the chances that one or more specialists will turn out to work well on the platform, and M can use a variable fee to extract the resulting surplus. More generally, beyond our specific model, hosting many firms will also intensify on-platform competition, so it should be desirable provided the platform has sufficient instruments to extract the resulting surplus. This seems to be the case for Apple's iPhone App Store: if Apple hosts rival app developers for a product category, then it is open to all such qualified apps. Clearly, Apple has instruments to extract surplus from these rivals, such as taking a cut on all app purchases (including in-app purchases), as well as via taking a share of their advertising revenues. On the other hand, if the platform does not have enough pricing instruments or pricing power to extract surplus from the hosted firms, and shopping costs are sufficiently low, firms should limit how many specialists are hosted in order to avoid having multiple high-quality specialists on the platform that compete away the additional surplus they offer. For instance, on its QuickBooks Capital lending marketplace for its QuickBooks small business customers, Intuit only allows a limited number of selected lenders in order to avoid aggressive price competition (i.e. competition on loan rates).

7 Future directions

Our paper is the first to provide a formal analysis of how a multiproduct firm can create a platform by hosting rivals. Naturally, there are other factors relevant to the decision whether or not to become a platform in this way that we did not capture. Thus, there are many interesting extensions that future research can explore.

A key benefit of hosting that our model does not capture is that by inviting multiple differentiated specialists (that appeal to different consumers), the platform can attract more consumers, and these additional consumers make it more attractive for specialists to be hosted by the platform, and so on. The resulting network effects can reinforce the benefits of hosting. However, since the benefits obtained by network effects are fairly well understood, we chose to abstract from them in the current paper, for simplicity.

On the other hand, even when firms can use transaction fees, if hosted specialists are unwilling or unable to charge different prices on M relative to their prices outside M (perhaps because they have a sufficiently large base of customers who they sell to directly), then M can no longer charge very high variable fees under hosting. This in turn may make hosting jointly less profitable relative to non-hosting.

There are also several longer-term risks associated with hosting which we have not considered in our formal modelling, but which could be considered in future work. Hosting may allow the multiproduct firm to learn from the rival specialist, after which it can offer its own better version, thus making the specialist regret the hosting partnership. One could argue Amazon has done this to some extent, by starting to sell certain product in its own name after seeing them become popular thanks to the sales efforts of third-party sellers on its marketplace. Thus, specialists need to protect themselves against this risk if the advantage they offer can be easily copied. Conversely, by being hosted, the specialist might be able to learn how to provide the multiproduct firm’s core product (e.g. by obtaining access to its customers), which can allow it to supplant the multiproduct firm itself. Finally, hosting may also subject each party to a hold-up risk to the extent they each need to incur some non-recoverable fixed costs of setting up, and so would make the firms vulnerable to ex-post exploitation via contract renegotiation.

References

- [1] Abhishek, V., K. Jerath and Z.J. Zhang (2016) “Agency Selling or Reselling? Channel Structures in Electronic Retailing,” *Management Science*, 62(8), 2259–2280.
- [2] August, T., H. Shin and T.I. Tunca (2013) “Licensing and competition for services in open source software,” *Information Systems Research*, 24(4), 1068–1086.
- [3] August, T., H. Shin and T. I. Tunca (2017) “Generating Value Through Open Source: Software Service Market Regulation and Licensing Policy,” *Information Systems Research*, 29(1), 186–205.
- [4] Boudreau, K. (2010) “Open platform strategies and innovation: Granting access vs. devolving control,” *Management Science*, 56(10), 1849–1872.
- [5] Chen, Z. and P. Rey (2012) “Loss Leading as an Exploitative Practice,” *American Economic Review*, 102(7), 3462–3482.
- [6] Eisenmann, T., G. Parker and M. Van Alstyne (2011) “Platform envelopment,” *Strategic Management Journal*, 32, 1270–1285.
- [7] Hagiu, A. and D. Spulber (2013) “First-Party Content and Coordination in Two-Sided Markets,” *Management Science*, 59(4), 933–949.
- [8] Hagiu, A. and J. Wright (2015a) “Marketplace or Reseller?” *Management Science*, 61(1), 184–203.
- [9] Hagiu and Wright (2015b) “Multi-sided Platforms,” *International Journal of Industrial Organization*, 43, 162–174.
- [10] Hagiu and Wright (2018) “Controlling versus Enabling,” *Management Science*, forthcoming.

- [11] Johnson, J. (2017) “The Agency Model and MFN Clauses” *Review of Economic Studies*, 84(3), 1151–1185.
- [12] Lal, R. and C. Matutes (1989) “Price Competition in Multimarket Duopolies,” *RAND Journal of Economics*, 20(4), 516–537.
- [13] Lal, R. and C. Matutes (1994) “Retail Pricing and Advertising Strategies,” *Journal of Business*, 67(3), 345–370.
- [14] Li, Z. and A. Agarwal (2017) “Platform Integration and Demand Spillovers in Complementary Markets: Evidence from Facebook’s Integration of Instagram,” *Management Science*, 63(10), 3438–3458.
- [15] Mantovani, A. and F. Ruiz-Aliseda (2016) “Equilibrium innovation ecosystems: the dark side of collaborating with complementors,” *Management Science*, 62(2), 534–549.
- [16] Niculescu, M. F., D.J. Wu and L. Xu (2018). “Strategic intellectual property sharing: Competition on an open technology platform under network effects,” *Information Systems Research*, forthcoming.
- [17] Parker, G. and M. Van Alstyne (2018) “Innovation, openness, and platform control,” *Management Science*, 64(7), 3015–3032.
- [18] Rey, P. and J. Tirole (2007) “A primer on foreclosure,” In Armstrong, M. and R.H. Porter (eds.), *Handbook of Industrial Organization III*, North-Holland: Elsevier.
- [19] Rhodes, A. and J. Zhou (2018) “Consumer Search and Retail Market Structure”, *Management Science*, forthcoming.
- [20] Whinston, M. (1990) “Tying, Foreclosure, and Exclusion,” *American Economic Review*, 80(4), 837–59.
- [21] Zhou, J. (2014) “Multiproduct Search and the Joint Search Effect,” *American Economic Review*, 104(9), 2918–2939
- [22] Zhu, F. and Q. Liu (2018) “Competing with Complementors: An Empirical Look at Amazon.com,” *Strategic Management Journal*, forthcoming.

8 Appendix

8.1 Proof of Proposition 1

First, we show why the prices in Proposition 1 characterize an equilibrium. Note the price p_A^* leaves A -type consumers indifferent between buying and not buying. In equilibrium the surplus of B -types is $v_B^* = u_A + u_B - p_A^* - p_B^* - \sigma = u_B + \Delta - \sigma > 0$ since $\sigma < u_B$, which just makes B -type consumers

indifferent between buying A and B_M , buying B_S alone, or buying A and B_S . If B -type consumers instead just buy B_M their surplus is $u_B - p_B^* - \sigma = u_B + \Delta - 2\sigma$, which is lower than v_B^* since $\sigma > 0$.

Obviously, S cannot do better lowering its price (and making a loss) or raising its price (since it still will not sell to any consumers). Since B -types just care about the total price $p_A + p_B$ charged for A and B_M , M always does better setting the maximum price possible to sell to the A -types and adjusting p_B so as to compete with S . If M raises p_A it will lose A -type consumers, and also lose B -type consumers unless it lowers p_B by a corresponding amount, which would imply no gain in profit from the B -types. Similarly, lowering p_A will cause M to make less from the A -type consumers, and also to make less from the B -type consumers unless it raises p_B by a corresponding amount, which implies no gain in profit from the B -types. The same logic applies for a deviation in p_B , which requires M make an offsetting adjustment in p_A in order to keep consumers, which either causes A -type consumers to drop out (if p_A is higher) or for M to make less profit from A -types (if p_A is lower). Note that requiring consumers to buy the bundle of A and B_M also wouldn't help since B -types already buy the bundle and A -types would not want to buy the bundle at the equilibrium prices, and furthermore, M cannot induce either type to pay more than they are currently paying by offering the two products as a bundle. Thus, neither firm has a profitable deviation.

We now rule out other possible equilibria. Obviously $p_A^* \leq u_A$, otherwise M would obtain no profit given that it has an inferior version of B . We can then rule out any equilibrium with $p_A > u_A - \sigma$. Indeed, in this case A -type consumers do not buy anything and the B -type consumers would not get a positive surplus from just buying A from M . Thus, in equilibrium, these consumers either buy A and B_M from M , obtaining a surplus of $u_A + u_B - p_A - p_B - \sigma$, or just B_S from S , obtaining a surplus of $u_B + \Delta - p_S - \sigma$. Given $u_A > \Delta$, in the proposed equilibrium we must have $p_S = 0$, and $p_A + p_B = u_A - \Delta$, giving M a profit of $\lambda_B(u_A - \Delta)$. But by deviating to p_A^* and p_B^* given in the proposition, M can obtain π_M^* , which is strictly higher since it also sells to the A -types.

The remaining possibility is an equilibrium in which $p_A \leq u_A - \sigma$ so that both types of consumers would always want to buy A . There cannot be an equilibrium involving the B -types buying B_S , since even if $p_S = 0$, M can always do better selling to B -types by setting the positive price $p_B = \sigma - \Delta$ to extract additional revenue by inducing these consumers to buy B_M , while keeping the price for A unchanged. Finally, note that in equilibrium we cannot have $p_A < u_A - \sigma$ since M always does better setting the maximum price possible to sell to the A -types (i.e. $p_A = u_A - \sigma$) and adjusting p_B so as to compete with S , given B -types only care about the total amount they pay for A and B_M .

8.2 Proof of Propositions 2 – 5

We assume M can charge S a variable fee τ per transaction when S is hosted (i.e. the setting of Section 4.5). Then the result without any transfer between the firms (or with a fixed transfer only) is obtained at the end by setting $\tau = 0$.

Consider the case with hosting. We start by solving for the equilibrium in the second stage for a given τ . If $\tau > \sigma$, then S prefers to sell directly instead of through M , and then the outcome is the same as under non-hosting. Thus, we focus on $\tau \leq \sigma$. The simpler case with $\tau = 0$, which establishes

Proposition 2, will be discussed at the end. In equilibrium, S must sell B_S to all B types through M because this is how it can offer the highest utility for the B product. There are two cases depending on p_A . (Note throughout this proof and subsequent proofs characterizing the hosting equilibrium we ignore the fixed cost of hosting F , since it is irrelevant for the analysis, but we do consider it when determining the tradeoff with non-hosting.)

Suppose first that $p_A = u_A - \sigma$, so M sells A to both A types and B types (there is no incentive to set p_A any lower to sell to all consumers). In this case, for B type consumers to prefer buying A from M and B_S from S on M rather than buying A and B_M from M or buying A from M and B_S from S outside M , we must have:

$$u_B + \Delta - \hat{p}_S \geq \max \{u_B - p_B, u_B + \Delta - p_S - \sigma\}.$$

Clearly, this must hold with equality in equilibrium, otherwise S could increase \hat{p}_S . Thus, we must have

$$\hat{p}_S = \min \{p_B + \Delta, p_S + \sigma\}.$$

If $p_S + \sigma < p_B + \Delta$, then we would have $\hat{p}_S = p_S + \sigma$ and S could profitably increase \hat{p}_S and p_S by the same amount. Thus, we must have $p_S + \sigma \geq p_B + \Delta$ and therefore

$$\hat{p}_S = p_B + \Delta.$$

Combined with $\tau \leq \sigma$, this implies that

$$\hat{p}_S - \tau = p_B + \Delta - \tau \geq p_B + \Delta - \sigma,$$

so S does not want to deviate by setting p_S slightly below $p_B + \Delta - \sigma$ and a sufficiently high \hat{p}_S , such that B types prefer to buy B_S from S outside M .

Furthermore, τ must not be above \hat{p}_S (so S makes non-negative profits) and M must not want to deviate by slightly decreasing p_B and selling B_M instead of getting τ from S . This means we must have

$$p_B \leq \tau \leq p_B + \Delta.$$

Finally, M must not want to increase p_A and only serve B types. The best such deviation for M is to set p'_A such that

$$u_A - p'_A + u_B + \Delta - \hat{p}_S - \sigma = u_B + \Delta - \sigma - p_S,$$

provided $p'_A \leq u_A$. So the best deviation is

$$p'_A = u_A + \min \{0, p_S - p_B - \Delta\}.$$

Note that $p'_A \geq u_A - \sigma$ because $p_S + \sigma \geq p_B + \Delta$. Deviation profits are therefore $\lambda_B (u_A + \min \{0, p_S - p_B - \Delta\} + \tau)$,

whereas M 's equilibrium profits are $u_A - \sigma + \lambda_B \tau$. For the deviation not to be profitable, we need

$$\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} + \min \{0, p_S - p_B - \Delta\} \leq 0.$$

S 's profits are $\lambda_B (p_B + \Delta - \tau)$. Since neither profit depends on p_S , we can always choose the lowest possible p_S , i.e. $p_S = p_B + \Delta - \sigma$, so this equilibrium always exists because $\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} < \sigma$. And since joint profits are increasing in p_B , we can focus on the equilibrium with the highest joint profits, which involves $p_B = \tau$. This implies $\hat{p}_S = \tau + \Delta$ and $p_S = \tau + \Delta - \sigma$, so M 's profits are $u_A - \sigma + \lambda_B \tau$, while S 's profits are $\lambda_B \Delta$.

Next, suppose $u_A - \sigma < p_A \leq u_A$, so M only sells to B types. For B types to prefer buying A from M and B_S from S on M rather than buying A and B_M from M or just buying B_S from S outside M , we must have:

$$u_A - p_A + u_B + \Delta - \hat{p}_S - \sigma \geq \max \{u_A - p_A + u_B - p_B - \sigma, u_B + \Delta - p_S - \sigma\}.$$

As in the previous case, this must hold with equality in equilibrium, so

$$\hat{p}_S = \min \{p_B + \Delta, u_A - p_A + p_S\}.$$

For the same reason as in the previous case, we must have $u_A - p_A + p_S \geq p_B + \Delta$ and therefore

$$\hat{p}_S = p_B + \Delta.$$

Also, given $\tau \leq \sigma$, as in the previous case, S does not want to deviate by setting p_S slightly below $p_B + \Delta - \sigma$ and a sufficiently high \hat{p}_S , such that B types prefer to buy B_S from S outside M . Furthermore, as in the previous case, we must have

$$p_B \leq \tau \leq p_B + \Delta.$$

Equilibrium profits are $\lambda_B (p_A + \tau)$ for M and $\lambda_B (p_B + \Delta - \tau)$ for S . Clearly, M wants to increase p_A as much as possible, so given $u_A - p_A + p_S \geq p_B + \Delta$, it must be that $p_A = u_A + \min \{0, p_S - p_B - \Delta\}$.

Finally, M must not want to decrease p_A to $u_A - \sigma$ and sell A to all consumers. This deviation would result in profits $u_A - \sigma + \lambda_B \tau$, whereas M 's equilibrium profits are $\lambda_B (u_A + \min \{0, p_S - p_B - \Delta\} + \tau)$. For this deviation not to be profitable we must have

$$\min \{0, p_S - p_B - \Delta\} + \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} \geq 0.$$

There are two possibilities for this equilibrium:

- If $p_S \geq p_B + \Delta$, then the equilibrium exists if and only if $\sigma \geq \lambda_A u_A$ and equilibrium profits are $\lambda_B (u_A + \tau)$ for M and $\lambda_B (p_B + \Delta - \tau)$ for S . We focus on the equilibrium with the highest joint profits, so $p_B = \tau$, $\hat{p}_S = \tau + \Delta$, $p_S \geq \tau + \Delta$ and $p_A = u_A$. Equilibrium profits are $\lambda_B (u_A + \tau)$

for M and $\lambda_B \Delta$ for S .

- If $p_S \leq p_B + \Delta$, then equilibrium profits are $\lambda_B (u_A + \tau + p_S - p_B - \Delta)$ for M and $\lambda_B (p_B + \Delta - \tau)$ for S and this equilibrium exists if and only if $\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} \geq p_B + \Delta - p_S \geq 0$. Joint profits are $\lambda_B (u_A + p_S) \leq \lambda_B (u_A + p_B + \Delta) \leq \lambda_B (u_A + \tau + \Delta)$. Thus, joint equilibrium profits in this case are dominated by joint profits in the previous case.

We conclude that the equilibrium in which M only sells to B types exists if and only $\sigma \geq \lambda_A u_A$ and the maximum joint profits that can be attained in this equilibrium are $\lambda_B (u_A + \tau + \Delta)$.

Since joint profits are increasing in τ for both the equilibrium in which M sells to both types of consumers and the equilibrium in which M just sells to B types, M will set the highest possible τ compatible with hosting, which is $\tau = \sigma$. In the equilibrium with M selling to both A types and B types, joint profits are then $u_A - \sigma + \lambda_B (\sigma + \Delta)$, whereas in the equilibrium with M selling to B types only (provided it exists, i.e. $\sigma \geq \lambda_A u_A$), joint profits are $\lambda_B (u_A + \sigma + \Delta)$. It is easily verified that the latter equilibrium has higher joint profits whenever it exists, thus we can conclude that the outcome that maximizes joint profits under hosting involves $\tau = \sigma$ and:

- If $\sigma \geq \lambda_A u_A$, then the highest equilibrium joint profit under hosting is $\lambda_B (u_A + \sigma + \Delta)$. Comparing with the equilibrium joint profits without hosting $u_A - \sigma + \lambda_B (\sigma - \Delta)$, and taking into account the fixed cost of hosting F , hosting is preferred if and only if $\Delta > \frac{F - (\sigma - \lambda_A u_A)}{2(1 - \lambda_A)}$.
- If $\sigma < \lambda_A u_A$, then the highest equilibrium joint profit under hosting is $u_A - \sigma + \lambda_B (\sigma + \Delta)$. In this case, hosting is preferred if and only if $\Delta > \frac{F}{2(1 - \lambda_A)}$.

Finally, to obtain the results for the case in which variable fees are not feasible (Propositions 2, 3 and 4), we can simply use the above analysis for the case $\tau = 0$. The equilibrium in which M sells to both types of consumers always exists and involves $p_A = u_A - \sigma$, $p_B = 0$, $\hat{p}_S = \Delta$ and $p_S = \Delta - \sigma$, so M 's profits are $u_A - \sigma$, while S 's profits are $\lambda_B \Delta$. The equilibrium in which M only sells to B types exists if and only if $\sigma \geq \lambda_A u_A$ and involves $p_A = u_A$, $p_B = 0$, $\hat{p}_S = \Delta$ and $p_S \geq \Delta$, so M 's profits are $\lambda_B u_A$, while S 's profits are $\lambda_B \Delta$. Thus:

- If $\sigma \geq \lambda_A u_A$, then the highest equilibrium joint profit under hosting is $\lambda_B (u_A + \Delta)$. In this case, hosting is preferred if and only if $\Delta > \frac{\lambda_A (u_A - \sigma) + F}{2(1 - \lambda_A)}$.
- If $\sigma < \lambda_A u_A$, then the highest equilibrium joint profit under hosting is $u_A - \sigma + \lambda_B \Delta$. In this case, hosting is preferred if and only if $\Delta > \frac{\sigma}{2} + \frac{F}{2(1 - \lambda_A)}$.

8.3 Proof of Proposition 6

We first show there cannot be any pure-strategy equilibrium in which M sells B_M to uninformed B -types only. We then characterize the conditions under which there is a pure-strategy equilibrium in which M sells B_M to both informed and uninformed B -types. Finally, we characterize the mixed-strategy equilibrium, which turns out to exist if and only if the pure-strategy equilibrium does not exist.

To show there is no pure-strategy equilibrium in which M sells B_M to uninformed B -type consumers only, note this can only be an equilibrium if M extracts the entire surplus from uninformed B -types, which means we must have $p_A + p_B = u_A + u_B - \sigma$. Furthermore, M can extract the entire surplus from A -types by setting $p_A = u_A - \sigma$ and $p_B = u_B$. Given these prices, S 's best response is to price just below $p_S = u_B + \Delta - \sigma$, which extracts almost the entire surplus from informed B -types. However, M could then deviate by slightly lowering p_B and attracting all informed B -type consumers as well, which results in a discrete increase in its profits. Thus, this cannot be part of a pure-strategy equilibrium.

Consider instead a possible pure-strategy equilibrium in which M sells to both informed and uninformed B -types. There are two cases to consider within this scenario: (a) M sets $p_A = u_A$ and only sells to B -types, and (b) M sets $p_A = u_A - \sigma$ and sells to both A -types and B -types. Case (a) is easily ruled out. Given that M sells to both informed and uninformed B -types, for this to be an equilibrium, we must have $p_B = -\Delta$ and $p_S = 0$, so M 's profits are $\lambda_B (u_A - \Delta)$. But then M could deviate to $p_A = u_A - \sigma$ and $p_B = \sigma - \Delta$, which yields strictly higher profit $u_A - \sigma + \lambda_B (\sigma - \Delta)$.

Consider case (b) in which $p_A = u_A - \sigma$. We must have $p_S = 0$ and $p_B = \sigma - \Delta$. Thus, profits are $\pi_M = u_A - \sigma + \lambda_B (\sigma - \Delta)$ and $\pi_S = 0$. Clearly, S cannot profitably deviate. M has three possible deviations. The first one is to set $p_B = u_B$ in order to only attract the uninformed B -types, in which case it must keep $p_A = u_A - \sigma$. This deviation yields profits $u_A - \sigma + \lambda_B \eta u_B$. The second possible deviation is to set $p_A > u_A - \sigma$ and $p_A + p_B = u_A - \Delta$, which ensures that M still sells to all informed and uninformed B -types but no longer sells to A -types. This yields profits $\lambda_B (u_A - \Delta)$, which can never be a profitable deviation. The third possible deviation is to set $p_A > u_A - \sigma$ and $p_A + p_B = u_A + u_B - \sigma$, which ensures that M only sells to uninformed B -types. This yields profits $\lambda_B \eta (u_A + u_B - \sigma)$, which are lower than the profits obtained through the first deviation. Thus, the pure-strategy equilibrium under case (b) exists if and only if $u_A - \sigma + \lambda_B (\sigma - \Delta) \geq u_A - \sigma + \lambda_B \eta u_B$, which is equivalent to $\eta < \frac{\sigma - \Delta}{u_B}$. Thus, the only possible pure-strategy equilibrium has $p_A = u_A - \sigma$, $p_B = \sigma - \Delta$ and $p_S = 0$, and it exists if and only if $\eta \leq \frac{\sigma - \Delta}{u_B}$.

We next determine the mixed-strategy equilibria. Denote by G_B the CDF of M 's price distribution for p_B and by G_S the CDF of S 's price distribution for p_S . Again, there are two possibilities: (a) $p_A = u_A$, so M sells only to B -types, and (b) $p_A = u_A - \sigma$ so M sells to both A -types and some B -types.

As before we can rule out case (a) arising in equilibrium. To see this, suppose $p_A = u_A$, so we must have $p_B \leq u_B - \sigma$. In this case, for any p_B that M plays with positive probability, its profit is

$$\lambda_B (u_A + p_B) (\eta + (1 - \eta) \Pr(p_B < p_S - \Delta)) = \lambda_B (u_A + p_B) (\eta + (1 - \eta) (1 - G(p_B + \Delta))).$$

However, by setting $p_A = u_A - \sigma$ and $\tilde{p}_B = p_B + \sigma$, the profit achieved by M becomes

$$\begin{aligned} & u_A - \sigma + \lambda_B \tilde{p}_B (\eta + (1 - \eta) (1 - G(\tilde{p}_B - \sigma + \Delta))) \\ &= u_A - \sigma + \lambda_B (p_B + \sigma) (\eta + (1 - \eta) (1 - G(p_B + \Delta))), \end{aligned}$$

which is strictly higher. Thus, setting $p_A = u_A$ cannot be part of an equilibrium.

Turning next to case (b), given $p_A = u_A - \sigma$, we must have $p_B \leq u_B$. In particular, M can guarantee profits $u_A - \sigma + \lambda_B \eta u_B$ by setting $p_B = u_B$. This implies that M will never set p_B below ηu_B (even if it attracted both informed and uninformed consumers at this price, it would not do better than ηu_B in total profits from selling B_M). Thus, the support of $G_B(\cdot)$ is $[\eta u_B, u_B]$. This implies that the support of $G_S(\cdot)$ is $[\eta u_B + \Delta - \sigma, u_B + \Delta - \sigma]$.

We determine $G_S(\cdot)$ by imposing that any price p_B in the support of G_B yields the same profit as setting $p_B = u_B$. This is equivalent to:

$$u_A - \sigma + p_B \lambda_B (\eta + (1 - \eta) (1 - G_S(p_B + \Delta - \sigma))) = u_A - \sigma + \lambda_B \eta u_B.$$

Rearranging and with the change of variables $p_S \equiv p_B + \Delta - \sigma$, this is equivalent to

$$G_S(p_S) = 1 - \frac{\eta}{1 - \eta} \frac{u_B + \Delta - \sigma - p_S}{p_S + \sigma - \Delta}.$$

Note $G_S(p_S)$ is increasing in p_S , and $G_S(u_B + \Delta - \sigma) = 1$ and $G_S(\eta u_B + \Delta - \sigma) = 0$, so $G_S(\cdot)$ has no mass points.

Similarly, we determine $G_B(\cdot)$ by imposing that any price p_S in the support of G_S yields the same profit as setting $p_S = \eta u_B + \Delta - \sigma$, thereby capturing all the informed customers. This is equivalent to

$$\lambda_B (1 - \eta) (1 - G_B(p_S + \sigma - \Delta)) p_S = \lambda_B (1 - \eta) (\eta u_B + \Delta - \sigma).$$

Rearranging and with the change of variables $p_B \equiv p_S - \Delta + \sigma$, this is equivalent to

$$G_B(p_B) = \frac{p_B - \eta u_B}{p_B - \sigma + \Delta}.$$

Note that $G_B(\eta u_B) = 0$. Furthermore, $G_B(p_B)$ is increasing in p_B if and only if $\eta > \frac{\sigma - \Delta}{u_B}$. If $\eta \leq \frac{\sigma - \Delta}{u_B}$, then $G_B(\cdot)$ is weakly decreasing and therefore there is no mixed-strategy equilibrium. Note also that $\eta > \frac{\sigma - \Delta}{u_B}$ implies $G_B(u_B) < 1$, so there is a mass point at $p_B = u_B$,

$$\Pr(p_B = u_B) = \frac{\eta u_B + \Delta - \sigma}{u_B + \Delta - \sigma} < 1.$$

In this mixed-strategy equilibrium, expected profits are

$$\begin{aligned} \pi_M &= u_A - \sigma + \lambda_B \eta u_B \\ \pi_S &= \lambda_B (1 - \eta) (\eta u_B + \Delta - \sigma). \end{aligned}$$

Finally, we need to check that M cannot profitably deviate by setting $p_A = u_A$ and thereby giving up selling to A -types. If it did, then M would have to set $p_B \leq u_B - \sigma$. When M sets p_B , it sells to

informed B -types only if $p_S > p_B + \Delta$. Thus, M 's deviation profits as a function of p_B are

$$\begin{aligned} & (u_A + p_B) \lambda_B (\eta + (1 - \eta) (1 - G_S(p_B + \Delta))) \\ = & (u_A + p_B) \lambda_B \eta \frac{u_B}{p_B + \sigma}, \end{aligned}$$

which is decreasing in p_B . Thus, M 's best deviation is to set $p_B = \eta u_B - \sigma$. The deviation profits are then $\lambda_B (u_A + \eta u_B - \sigma)$, which is clearly lower than the equilibrium profits $u_A - \sigma + \lambda_B \eta u_B$. Thus, the deviation is not profitable and the mixed-strategy equilibrium we have determined exists if and only if $\eta > \frac{\sigma - \Delta}{u_B}$.

8.4 Proof of Proposition 7

We prove that the equilibrium outcome and profits are the same as in the benchmark setting (Propositions 2 and 5). Indeed, if M sets $p_A \leq u_A - \sigma$ and so sells to all consumers, then all consumers will be informed when deciding which version of B to buy and from where. Thus, everything is the same as in the analysis of hosting in the benchmark setting. The other case is when M sets $p_A > u_A - \sigma$. The $1 - \eta$ informed consumers have the same options as before. The η uninformed consumers only consider whether to go to M and buy both A and B_M . However, in order for M to be able to somehow extract more from these consumers, it would need to attract them. But when these consumers arrive at M , they will find out about S , and they will have the same options of choosing to purchase B_S instead of B_M if M tries to extract more from them (e.g., if it sets a higher τ in its contract and a higher p_B). This reflects that this was the relevant constraint on M 's pricing before, and ensured that $\tau \leq \sigma$. There is no new price deviation for M that would allow it to exploit the uninformed consumers. (Note the possibility of M setting τ to make it impossible for S to compete on M would just cause everything to revert to the non-hosting outcome. If joint profits are higher in this case, M would not want to host in the first place, so it doesn't make sense to consider such a τ .) As a result, the equilibrium is the same as in the benchmark model.

8.5 Proof of Propositions 10 - 11

We start with the more complicated case in which M can charge a variable fee τ when it hosts specialists, i.e. Proposition 11. We then derive the proof of Proposition 10 from the equilibrium of the subgame in which $\tau = 0$ under hosting.

Regardless of how many specialists are hosted, Bertrand competition outside the platform implies that in equilibrium we can always restrict attention to the case where all specialists (hosted and not hosted) price at zero outside the platform.

Suppose M hosts $k \in [1, n]$ specialists and has set a variable fee $\tau \geq 0$ in its contract with the hosted specialists. Given that all specialists are identical, there are only three distinct cases to consider: (i) all the hosted specialists turn out to offer u_B only, (ii) exactly one of the hosted specialists turns out to offer $u_B + \Delta$ and all other specialists offer u_B only, (iii) two or more of the hosted specialists turn out to offer $u_B + \Delta$.

Consider case (i) first. If all hosted specialists turn out to offer only u_B , then all specialists make zero profits and the outcome is as follows:

- if $\tau \geq \sigma - \Delta$, then M keeps selling to both types and makes the sales of B_M by setting $p_A = u_A - \sigma$ and $p_B = \sigma - \Delta$, obtaining profits $u_A - \sigma + \lambda_B (\sigma - \Delta)$, which is the same as without hosting.
- if $\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} - \Delta \leq \tau < \sigma - \Delta$, then M keeps selling to both types and makes the sale of B_M by setting $p_A = u_A - \sigma$ and $p_B = \tau$, obtaining profits $u_A - \sigma + \lambda_B \tau$.
- if $\tau < \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} - \Delta$, then M sells to B types only by setting $p_A = u_A - \tau - \Delta$ and $p_B = \tau$, obtaining profits $\lambda_B (u_A - \Delta)$.

Note that joint profits are (weakly) increasing in τ . They are maximized for $\tau \geq \sigma - \Delta$, when they are equal to $u_A - \sigma + \lambda_B (\sigma - \Delta)$.

Consider case (ii), when exactly one of the hosted specialists turns out to offer $u_B + \Delta$. With some abuse of terminology, we will refer to the specialist offering $u_B + \Delta$ as S . If $\tau > \sigma$, then S cannot make non-negative profits while avoiding that consumers prefer going to an outside specialist. The same goes for the other hosted specialists. In this case, M sets $p_A = u_A - \sigma$ and $p_B = \sigma - \Delta$ and the outcome is the same as without hosting.

Suppose then $\tau \leq \sigma$. In equilibrium, S must sell B_S to all B types because it can offer the highest utility for the B product. All other hosted specialists price at τ on M . There are two cases depending on p_A .

Suppose first that $p_A = u_A - \sigma$, so all consumers buy the A product (there is no need to set p_A any lower to attract all consumers). In this case, for S to make the sales of B_S , we must have

$$u_B + \Delta - \hat{p}_S \geq \max \{u_B - p_B, u_B + \Delta - \sigma\}.$$

Clearly, this must hold with equality in equilibrium, otherwise S could increase \hat{p}_S . Thus, we must have

$$\hat{p}_S = \min \{p_B + \Delta, \sigma\}.$$

Furthermore, τ must not be above \hat{p}_S (so S makes non-negative profits) and M must not want to deviate by setting p_B slightly below $\hat{p}_S - \Delta$ and selling B_M instead of getting τ from S . This means we must have

$$\min \{p_B, \sigma - \Delta\} \leq \tau \leq \min \{p_B + \Delta, \sigma\}.$$

Finally, M must not want to increase p_A and only serve B types. The best such deviation for M is to set p'_A such that

$$u_A - p'_A + u_B + \Delta - \hat{p}_S - \sigma = u_B + \Delta - \sigma,$$

provided the solution in p'_A is below u_A . So the best deviation is

$$p'_A = u_A - \hat{p}_S = u_A - \min \{p_B + \Delta, \sigma\}.$$

If $p_B \geq \sigma - \Delta$, then $p'_A = u_A - \sigma$, so this deviation does not do any better. If $p_B \leq \sigma - \Delta$, then $p'_A = u_A - \Delta - p_B$, so M 's deviation profits are $\lambda_B (u_A - \Delta - p_B + \tau)$, whereas M 's equilibrium profits are $u_A - \sigma + \lambda_B \tau$. For the deviation not to be profitable, we then need

$$\lambda_A u_A + \lambda_B (\Delta + p_B) \geq \sigma.$$

Thus, if $p_B \geq \sigma - \Delta$, then $(p_A = u_A - \sigma, \hat{p}_S = \sigma, p_S = 0)$ is an equilibrium given τ if and only if

$$\sigma - \Delta \leq \tau \leq \sigma.$$

Equilibrium profits are $u_A - \sigma + \lambda_B \tau$ for M and $\lambda_B (\sigma - \tau)$ for S . Since these profits do not depend on p_B , we can just focus on $p_B = \sigma - \Delta$.

If $p_B \leq \sigma - \Delta$, then $(p_A = u_A - \sigma, \hat{p}_S = \Delta + p_B, p_S = 0)$ is an equilibrium given τ if and only if

$$p_B \leq \tau \leq p_B + \Delta \text{ and } \lambda_A u_A + \lambda_B (\Delta + p_B) \geq \sigma.$$

Since joint profits are increasing in p_B , we focus on the highest possible p_B , which is $p_B = \min \{\tau, \sigma - \Delta\}$. If $\sigma - \Delta \leq \tau$, then $p_B = \sigma - \Delta$, so this is the same equilibrium as in the case $p_B \geq \sigma - \Delta$. If $\sigma - \Delta > \tau$, then $p_B = \tau$ and $(p_A = u_A - \sigma, p_B = \tau, \hat{p}_S = \Delta + \tau, p_S = 0)$ is an equilibrium given τ if and only if

$$\tau \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} - \Delta.$$

In this equilibrium, profits are $u_A - \sigma + \lambda_B \tau$ for M , $\lambda_B \Delta$ for S , and zero for all other specialists.

To conclude this case, the equilibrium with M selling to A -types exists if and only if $\tau \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} - \Delta$ and profits are $u_A - \sigma + \lambda_B \tau$ for M and $\lambda_B \min \{\Delta, \sigma - \tau\}$ for S . Thus, joint profits are $u_A - \sigma + \lambda_B \min \{\Delta + \tau, \sigma\}$.

Suppose now $u_A - \sigma < p_A \leq u_A$, so M does not sell to A types. In this case, for S to make the sales of B_S , we must have

$$u_A - p_A + u_B + \Delta - \hat{p}_S \geq \max \{u_A - p_A + u_B - p_B, u_B + \Delta\}.$$

Clearly, this must hold with equality in equilibrium, otherwise S could increase \hat{p}_S . Thus, we must have

$$\hat{p}_S = \min \{\Delta + p_B, u_A - p_A\}.$$

Furthermore, τ must not be above \hat{p}_S (in order that S does not make a loss) and M must not want to deviate by setting p_B slightly below $\hat{p}_S - \Delta$ and selling B_M instead of getting τ from S . This means we must have

$$\min \{p_B, u_A - p_A - \Delta\} \leq \tau \leq \min \{\Delta + p_B, u_A - p_A\}$$

Finally, M must not want to decrease p_A to $u_A - \sigma$ and sell A to all consumers. This deviation would result in profits $u_A - \sigma + \lambda_B \tau$, whereas M 's equilibrium profits are $\lambda_B (p_A + \tau)$. For this deviation

not to be profitable we must have

$$u_A - p_A \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}.$$

There are two possibilities for this equilibrium:

- If $p_B \leq u_A - p_A - \Delta$, then this equilibrium exists if and only if

$$p_B \leq \tau \leq p_B + \Delta \text{ and } u_A - p_A \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}.$$

Equilibrium profits are then $\lambda_B (p_A + \tau)$ for M and $\lambda_B (p_B + \Delta - \tau)$ for S . Clearly, M would want to increase p_A as much as possible, so it must be that $p_A = u_A - p_B - \Delta$. Thus, for any p_B in the interval $[\tau - \Delta, \tau]$, equilibrium profits are $\lambda_B (u_A - p_B + \tau - \Delta)$ for M and $\lambda_B (p_B + \Delta - \tau)$ for S , and this is an equilibrium if and only if $p_B + \Delta \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$. Thus, there exists a p_B in the interval $[\tau - \Delta, \tau]$ such that this is an equilibrium if and only if $\tau \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$. Note that joint profits are $\lambda_B u_A$ and so do not depend on p_B or τ .

- If $p_B \geq u_A - p_A - \Delta$, then this equilibrium exists if and only if

$$u_A - p_A - \Delta \leq \tau \leq u_A - p_A \text{ and } u_A - p_A \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}.$$

Equilibrium profits are then $\lambda_B (p_A + \tau)$ for M and $\lambda_B (u_A - p_A - \tau)$ for S . Clearly, M would want to increase p_A as much as possible, so it must be that $p_A = u_A - \tau$ and profits are $\lambda_B u_A$ for M and 0 for S . This is an equilibrium if and only if $\tau \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$.

Combining these two possibilities, we conclude that the equilibrium in which M does not sell to A types exists if and only if $\tau \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$ and joint profits in this equilibrium are always $\lambda_B u_A$.

Thus, summarizing case (ii), we find:

- If $\tau \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} - \Delta$, then only the equilibrium in which M does not sell to A types exists and joint profits are $\lambda_B u_A$ (there are multiple equilibria depending on how profits are shared between M and S).
- If $\tau \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, then only the equilibrium in which M sells A types exists and joint profits are $u_A - \sigma + \lambda_B \min \{\Delta + \tau, \sigma\}$.
- If $\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} - \Delta \leq \tau \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, then both types of equilibria exist, so joint profits are either $\lambda_B u_A$ or $u_A - \sigma + \lambda_B \min \{\Delta + \tau, \sigma\}$.

Clearly, joint profits are weakly increasing in τ , so they are maximized for $\tau = \sigma$, which leads to profits $u_A - \lambda_A \sigma$ for M and zero for S .

Finally, consider case (iii), in which two or more of the hosted specialists turn out to offer $u_B + \Delta$. In this case, Bertrand competition pins down the prices of all the hosted specialists on M so that $\hat{p}_S = \tau$, while Bertrand competition outside M continues to pin down the specialists' outside prices at

$p_S = 0$. As in the previous case, if $\tau > \sigma$, then the hosted specialists cannot compete with the outside specialists and make non-negative profits, so the outcome is the same as under non-hosting. Suppose then $\tau \leq \sigma$. Then B -type consumers prefer buying B_S and A on M instead of only buying B_S outside if and only if $p_A + \tau \leq u_A$.

Consider first the equilibrium when M sells to A types. Then $p_A = u_A - \sigma$, so clearly $p_A + \tau \leq u_A$. In this candidate equilibrium, M makes profits $u_A - \sigma + \lambda_B \tau$. Note that in this case M does not want to deviate by setting p_B slightly below $\hat{p}_S - \Delta = \tau - \Delta$ and selling B_M instead of getting τ from the hosted specialists. The only remaining condition for ensuring this is an equilibrium is that M does not want to increase p_A and only serve B types. The best such deviation is to set $p'_A = u_A - \tau$, leading to deviation profits $\lambda_B u_A$. This deviation is not profitable if and only if $\tau \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$.

Next consider the equilibrium when M does not sell to A types. Then $p_A = u_A - \tau \geq u_A - \sigma$. In this candidate equilibrium, M makes profits $\lambda_B u_A$. This is an equilibrium if and only if M does not want to deviate by decreasing p_A to $u_A - \sigma$ and serve both A types and B types. This condition is equivalent to $\tau \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$.

Thus, summarizing case (iii):

- If $\tau \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, then there is a unique equilibrium in which M sells to A types by setting $p_A = u_A - \sigma$ and joint profits are $u_A - \sigma + \lambda_B \tau$
- If $\tau \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, then there is a unique equilibrium in which M does not sell to A types involving $p_A = u_A - \tau$ and joint profits are $\lambda_B u_A$. Once again, joint profits are weakly increasing in τ and maximized for $\tau = \sigma$, when they are equal to $u_A - \lambda_A \sigma$.

Consequently, joint profits are weakly increasing in τ up to $\tau = \sigma$ in all three scenarios that can occur after uncertainty is realized. This implies that from an ex-ante perspective (before uncertainty is realized), M will set $\tau = \sigma$. Thus, when $k \geq 1$ specialists are hosted, expected joint profits are

$$\begin{aligned} E[\pi_M + \pi_S] &= (1 - \theta)^k (u_A - \sigma + \lambda_B (\sigma - \Delta)) + \left(1 - (1 - \theta)^k\right) (u_A - \lambda_A \sigma) \\ &= u_A - \lambda_A \sigma - (1 - \theta)^k \lambda_B \Delta. \end{aligned}$$

Thus, joint profits are increasing in k , so from a joint profit perspective, M wants to host all n available specialists in order to maximize the probability of having at least one specialist of high quality on the platform. In this case, hosting is preferred to non-hosting if and only if

$$u_A - \lambda_A \sigma - (1 - \theta)^n \lambda_B \Delta - F > u_A - \sigma + \lambda_B (\sigma - \Delta),$$

which is equivalent to

$$\Delta > \frac{F}{(1 - \lambda_A)(1 - (1 - \theta)^n)}.$$

Finally, if M cannot charge monitor the hosted specialists' sales under hosting (Proposition 10), then the analysis above applies by imposing $\tau = 0$. The outcomes in the three cases above are now as follows:

- Case (i): none of the hosted specialists turn out to offer $u_B + \Delta$. In this case, all specialists make zero profit. If $\Delta \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, then M sells A to both types and joint profits are $u_A - \sigma$. If $\Delta < \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, then M sells to B types only by setting $p_A = u_A - \Delta$ and joint profits are $\lambda_B (u_A - \Delta)$. In short, joint profits are $\max \{u_A - \sigma, \lambda_B (u_A - \Delta)\}$.
- Case (ii): only one of the hosted specialists turns out to offer $u_B + \Delta$. In this case, if $\Delta \leq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, then only the equilibrium in which M sells exclusively to B types exists and joint profits are $\lambda_B u_A$. If $\sigma \leq \lambda_A u_A$, then only the equilibrium in which M sells to both A and B types exists and joint profits are $u_A - \sigma + \lambda_B \Delta$. If $\Delta \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} \geq 0$, then both types of equilibria exist, so joint profits are either $\lambda_B u_A$ or $u_A - \sigma + \lambda_B \Delta$. And since $\Delta \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A}$, we know that $u_A - \sigma + \lambda_B \Delta \geq \lambda_B u_A$, so we assume the two firms coordinate on the equilibrium with higher joint profits, i.e. $u_A - \sigma + \lambda_B \Delta$. Thus, to summarize this case: joint profits are $\max \{\lambda_B u_A, u_A - \sigma + \lambda_B \Delta\}$.
- Case (iii): two or more hosted specialists turn out to offer $u_B + \Delta$. In this case, if $\sigma \leq \lambda_A u_A$, then joint profits are $u_A - \sigma$; if $\sigma \geq \lambda_A u_A$, then joint profits are $\lambda_B u_A$. In short, joint profits are $\max \{u_A - \sigma, \lambda_B u_A\}$.

Consequently, expected joining profits are:

$$E[\pi_M + \pi_S] = (1 - \theta)^k \max \{u_A - \sigma, \lambda_B (u_A - \Delta)\} + k(1 - \theta)^{k-1} \theta \max \{\lambda_B u_A, u_A - \sigma + \lambda_B \Delta\} \\ + \left(1 - (1 - \theta)^k - k\theta(1 - \theta)^{k-1}\right) \max \{u_A - \sigma, \lambda_B u_A\}.$$

There are therefore three cases:

- If $\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} \geq \Delta$, then joint profits are

$$E[\pi_M + \pi_S] = (1 - \theta)^k \lambda_B (u_A - \Delta) + \left(1 - (1 - \theta)^k\right) \lambda_B u_A.$$

They are increasing in k , so joint profits are maximized by hosting all available specialists, i.e. $k^* = n$. In this case, hosting is jointly preferred to non-hosting if and only if

$$\Delta > \frac{\lambda_A (u_A - \sigma) + F}{(1 - \lambda_A) (1 - (1 - \theta)^n)}.$$

- If $\Delta \geq \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} \geq 0$, then joint profits are

$$E[\pi_M + \pi_S] = (1 - \theta)^k (u_A - \sigma) + k(1 - \theta)^{k-1} \theta (u_A - \sigma + \lambda_B \Delta) + \left(1 - (1 - \theta)^k - k(1 - \theta)^{k-1} \theta\right) \lambda_B u_A \\ = k(1 - \theta)^{k-1} \theta (\lambda_A u_A - \sigma + \lambda_B \Delta) - (1 - \theta)^k (\sigma - \lambda_A u_A) + \lambda_B u_A.$$

In this case, the optimal number of specialists to host from a joint profit perspective is

$$k^* = \frac{(1 - \theta) (\sigma - \lambda_A u_A)}{\theta (\lambda_A u_A - \sigma + \lambda_B \Delta)} - \frac{1}{\ln(1 - \theta)}$$

and hosting is preferred to non-hosting if and only if

$$\Delta > \frac{\sigma(1 - \lambda_A) - \left(1 - (1 - \theta)^{k^*} - k^*(1 - \theta)^{k^*-1}\theta\right)(\sigma - \lambda_A u_A) + F}{\left(1 + k^*(1 - \theta)^{k^*-1}\theta\right)(1 - \lambda_A)}$$

- If $\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} \leq 0$, then joint profits are

$$\begin{aligned} E[\pi_M + \pi_S] &= k(1 - \theta)^{k-1}\theta(u_A - \sigma + \lambda_B \Delta) + \left(1 - k(1 - \theta)^{k-1}\theta\right)(u_A - \sigma) \\ &= k(1 - \theta)^{k-1}\theta\lambda_B \Delta + u_A - \sigma. \end{aligned}$$

In this case, the optimal number of specialists to host from a joint profit perspective is

$$k^* = -\frac{1}{\ln(1 - \theta)}$$

and hosting is preferred to non-hosting if and only if

$$\Delta > \frac{\sigma(1 - \lambda_A) + F}{\left(1 + k^*(1 - \theta)^{k^*-1}\theta\right)(1 - \lambda_A)}.$$

Online Appendix

This online appendix contains the formal details for various results noted in the main text, and the proofs behind some of the results and claims in the main text.

A Correlation in consumers' valuations across products

In the benchmark model, we assumed both types of consumers value A the same. We now explore what happens when the two types of consumers place different values on product A in the case without variable fees. Specifically, we assume B -types continue to value product A at u_A , but A -types value it at $u_A + \alpha$. We will consider both the case when α is positive (i.e. there is *negative* correlation between the values different types of consumers place on products A and B) and the case α is negative (i.e. there is *positive* correlation between the values different types of consumers place on products A and B).

A.1 Negative correlation

Suppose $0 < \alpha \leq \sigma$, so A -types are willing to pay α more for product A than are B -types. This captures the idea that there are some consumers who value A highly and do not need B (e.g. they may be serious body builders who go to the gym only to use the weightlifting equipment and have no time for cycling), while others are interested in both A and B , but value A relatively less (e.g. they go to the gym to for a variety of workouts). Comparing joint profits under hosting and without hosting, we obtain the following proposition.

Proposition 12 *If $\lambda_A \leq \frac{\sigma - \alpha}{u_A}$, then hosting is jointly preferred iff $\Delta > \frac{\lambda_A(u_A + \alpha - \sigma) + F}{2(1 - \lambda_A)}$. If $\lambda_A > \frac{\sigma - \alpha}{u_A}$, then hosting is jointly preferred iff $\Delta > \frac{\sigma - \alpha}{2} + \frac{F}{2(1 - \lambda_A)}$.*

Proof. Consider first what happens without hosting. Note if M charges a price of p_A , A -type consumers will buy A provided $p_A \leq u_A + \alpha - \sigma$. The choices of B -types is the same as our previous analysis with $\alpha = 0$. Recall M competes by selling both A and B_M to B -types. This allows it to increase its price to A -types to their maximum willingness to pay (now $u_A + \alpha - \sigma$), while still giving exactly the same surplus to B -types as before. Thus, without hosting there is a unique equilibrium outcome in which the prices are $p_A^* = u_A + \alpha - \sigma$, $p_B^* = \sigma - \alpha - \Delta$, $p_S^* = 0$. The A -type consumers always purchase A , and the B -type consumers all buy A and B_M from M . Profits are $\pi_M^* = u_A - \sigma + \lambda_A \alpha + \lambda_B (\sigma - \Delta)$ and $\pi_S^* = 0$. The proof follows the same steps as the proof of Proposition 1. As before, M cannot do better deviating. Note this remains true even if $p_B^* < 0$. If M sets $p_A = u_A - \sigma$ and sets a high p_B to induce multi-stop shopping, it will be worse off, since M would give up α on A -types and $\sigma - \Delta > 0$ on B -types. Moreover, the same alternative possibilities for equilibria can be ruled out using the same arguments as before, since M always does better setting the maximum price possible to sell to the A -types and adjusting p_B so as to compete with S . This logic, also rules out any equilibrium with a price $u_A - \sigma < p_A < u_A + \alpha - \sigma$.

With hosting a similar tradeoff arises to before (i.e. whether to sell to all consumers or just B -types), except now the benefit of keeping A -type consumers is greater given they are willing to pay for product A . As before M has two options. Either it can set $p_A = u_A + \alpha - \sigma < u_A$ and sell A to all consumers, obtaining $\pi_M = u_A + \alpha - \sigma$, or set $p_A = u_A$ and sell A only to B -types, obtaining $\pi_M = \lambda_B u_A$. Then, we find (i) if $\lambda_A \leq \frac{\sigma - \alpha}{u_A}$, the selected equilibrium involves the prices $p_A^* = u_A$, $p_B^* = 0$, and $\widehat{p}_S^* = \Delta$, the A -type consumers do not purchase, while the B -type consumers all buy A and B_S through M , and profits are $\pi_M^* = \lambda_B u_A$ and $\pi_S^* = \lambda_B \Delta$; (ii) if $\lambda_A > \frac{\sigma - \alpha}{u_A}$, the selected equilibrium involves the prices $p_A^* = u_A + \alpha - \sigma$, $p_B^* = 0$, and $\widehat{p}_S^* = \Delta$, the A -type consumers always purchase A , and the B -type consumers all buy A and B_S through M , and profits are $\pi_M^* = u_A + \alpha - \sigma$ and $\pi_S^* = \lambda_B \Delta$.

The proposition follows by comparing the joint profit worked out above under hosting with joint profit under non-hosting, and taking into account the fixed cost of hosting F . ■

The tradeoff is similar to before, but there are some changes to note. The non-hosting profit extracted by M from B -types is not affected by α , since M can price discriminate: this means M 's profit is just higher by the additional α obtained from A -types. By contrast, under hosting, the fact that $\alpha > 0$ means A -types are less of a constraint on the amount that M can extract from B -types since A -types are willing to pay more for A . This improves the profitability of hosting, unless M no longer wants to serve A -types under hosting, in which case M gives up more by hosting.

Consistent with this logic, a comparison of the regions under which hosting makes the firms jointly better off shows that hosting dominates for a larger range of Δ when M still sells to A -types (this occurs for large α), but dominates for a smaller range of Δ when M stops selling A -types (this occurs for small α). In the extreme case when $\alpha = \sigma$, the shopping cost is offset by the extra benefit that A -types get from product A , so A -types do not constrain at all the amount that M can extract from B -types even if it cannot price discriminate. Thus, apart from the fixed costs of hosting, hosting always dominates when $\alpha = \sigma$ as there is no other cost to hosting.

A.2 Positive correlation

Suppose instead that $\alpha < 0$, so B -type consumers are willing to pay more for both products than A -type consumers. Comparing joint profits under hosting and without hosting, we obtain the following proposition.

Proposition 13 *If $\lambda_A < -\frac{\alpha}{u_A - \sigma}$, then hosting is jointly preferred iff $\Delta > \frac{F}{2(1 - \lambda_A)}$. If $-\frac{\alpha}{u_A - \sigma} \leq \lambda_A \leq \frac{\sigma - \alpha}{u_A}$, then hosting is jointly preferred iff $\Delta > \frac{\lambda_A(u_A - \sigma + \alpha) + F}{2(1 - \lambda_A)} + \frac{\alpha}{2}$. If $\lambda_A > \frac{\sigma - \alpha}{u_A}$, then hosting is jointly preferred iff $\Delta > \frac{\sigma}{2} + \frac{F}{2(1 - \lambda_A)}$.*

Proof. Consider first what happens without hosting. The equilibrium prices with non-hosting must satisfy $p_B^* \leq \sigma - \Delta$ in order for B -type consumers to prefer buying B_M to B_S , and $p_A^* + p_B^* \leq u_A - \Delta$ in order for B -type consumers to prefer buying A and B_M instead of just B_S , and $p_A^* \leq u_A + \alpha - \sigma$ if M sells to A -types or $p_A^* \leq u_A$ if M just sells to B -types. The new equilibria are characterized by:

NH-1 If $\lambda_A (\sigma - u_A) \leq \alpha < 0$ (or equivalently, $\lambda_A \geq -\frac{\alpha}{u_A - \sigma}$), then $p_A^* = u_A + \alpha - \sigma$, $p_B^* = \sigma - \Delta$, $p_S^* = 0$, with A -types still purchasing, and B -types buying the bundle from M , with profits being $\pi_M^* = u_A + \alpha - \lambda_A \sigma - \lambda_B \Delta$ and $\pi_S^* = 0$.

NH-2 If $\alpha < \lambda_A (\sigma - u_A) < 0$ (or equivalently, $\lambda_A < -\frac{\alpha}{u_A - \sigma}$), then M gives up on selling to A -types, $u_A - \sigma \leq p_A^* \leq u_A$ and $p_A^* + p_B^* = u_A - \Delta$, $p_S^* = 0$, with B -types buying the bundle from M , with profits being $\pi_M^* = \lambda_B (u_A - \Delta)$ and $\pi_S^* = 0$.

With hosting, the previous analysis with $\alpha \geq 0$ still holds, so the profit is defined in the proof of Proposition 12, in which there are two cases:

H-1 If $\lambda_A \leq \frac{\sigma - \alpha}{u_A}$, profits are $\pi_M^* = \lambda_B u_A$ and $\pi_S^* = \lambda_B \Delta$.

H-2 If $\lambda_A > \frac{\sigma - \alpha}{u_A}$, profits are $\pi_M^* = u_A + \alpha - \sigma$ and $\pi_S^* = \lambda_B \Delta$.

Note that $u_A + \alpha > \sigma$ (which is required for A -types to be willing to participate) implies the threshold $-\frac{\alpha}{u_A - \sigma}$ is smaller than the threshold $\frac{\sigma - \alpha}{u_A}$. Therefore, we have three cases when comparing the joint profits under hosting with non-hosting.

- If $\lambda_A < -\frac{\alpha}{u_A - \sigma}$, then NH-2 and H-1 apply, so we can compare $\lambda_B (u_A + \Delta) - F$ under hosting with $\lambda_B (u_A - \Delta)$ without hosting.
- If $-\frac{\alpha}{u_A - \sigma} \leq \lambda_A \leq \frac{\sigma - \alpha}{u_A}$, then NH-1 and H-1 apply, so we can compare $\lambda_B (u_A + \Delta) - F$ under hosting with $u_A + \alpha - \lambda_A \sigma - \lambda_B \Delta$ without hosting.
- If $\lambda_A > \frac{\sigma - \alpha}{u_A}$, then NH-1 and H-2 apply, so we can compare $u_A + \alpha - \sigma + \lambda_B \Delta - F$ under hosting with $u_A + \alpha - \lambda_A \sigma - \lambda_B \Delta$ without hosting.

The proposition follows by comparing the joint profit worked out above under hosting with joint profit under non-hosting. ■

The previous logic and tradeoff still apply. This suggests that $\alpha < 0$ tightens the constraint coming from A -types in the hosting equilibrium, thus making hosting less profitable. On the other hand, this also means that M loses less when it stops selling to A -types, which tends to make hosting more profitable. Finally, there is a novel effect when $\alpha < 0$: under non-hosting p_B is now constrained by competition in B (previously this constraint was not binding so M could adjust p_A and p_B to extract the maximum surplus from B -types). This limits M 's ability to price discriminate, which previously was the key benefit provided by non-hosting. If α is sufficiently negative, then M no longer serves A -types under non-hosting, so in this case, if $F = 0$, then hosting always dominates. If M keeps selling to A -types under non-hosting, M 's limited ability to benefit from price discrimination shifts the tradeoff in favor of hosting.

B Horizontal differentiation with respect to product B

Our results do not depend crucially on the assumption that B -type consumers are all the same. Consider the variation from our baseline model in which B -type consumers have heterogeneous tastes over products B_M and B_S . Specifically, suppose B -type consumers value B_M and B_S at u_B and $u_B + \Delta$ respectively, less their individual mismatch cost. Their mismatch cost is tx if purchasing B_M and $t(1 - x)$ if purchasing B_S for a consumer located at x , where consumers have x drawn from $U[0, 1]$. Thus, we model heterogeneous tastes using the standard Hotelling model of horizontal product differentiation. Other than this, we retain the assumptions of our baseline specification, and add a condition on the mismatch parameter t so that the market for B is always covered (t is not too high) and a condition on t so that both firms obtain positive markets shares in equilibrium both with and without hosting (t is not too low). Then we obtain the following proposition.

Proposition 14 *Suppose there is horizontal differentiation for product B , with the mismatch parameter t satisfying $\max(\frac{\sigma - \Delta}{3}, \frac{\Delta}{3}) < t < \frac{2u_B}{3} + \min\{\frac{\sigma - \Delta}{9}, \frac{\Delta}{3}\}$. When $\lambda_A \leq \frac{\sigma}{u_A}$, hosting is jointly preferred iff $\Delta > \frac{\sigma}{2} + \frac{9t(\lambda_A u_A - \sigma + F)}{2\sigma(1 - \lambda_A)}$. When $\lambda_A > \frac{\sigma}{u_A}$, hosting is jointly preferred iff $\Delta > \frac{\sigma}{2} + \frac{9tF}{2\sigma(1 - \lambda_A)}$.*

Proof. First consider the case without hosting. Note that $p_A \leq u_A$ otherwise M never sells A . We can also rule out M setting p_A such that $u_A - \sigma < p_A \leq u_A$, so A -types do not buy A . Suppose there is an equilibrium with this property. In this case B -types would not get a positive surplus from just buying A from M . Therefore, they either buy A and B_M from M or just B_S from S . It is straightforward to check that M will always prefer to set $p'_A = u_A - \sigma$ so as to sell to the A -types, and adjust the price for p_B to sell the bundle A and B_M at the same joint price $p_A + p_B$ as in the proposed equilibrium, which it can always do by setting a higher price for p_B .

Given $p_A \leq u_A - \sigma$, we know A -types will purchase and B -types who prefer to buy B_S will choose to multi-stop shop rather than one-stop shop at S . In this case, M does best setting $p_A = u_A - \sigma$, and the two firms' respective profits are $\pi_M = p_A + \lambda_B p_B \left(\frac{1}{2} + \frac{p_S - p_B + \sigma - \Delta}{2t}\right)$ and $\pi_S = \lambda_B p_S \left(\frac{1}{2} - \frac{p_S - p_B + \sigma - \Delta}{2t}\right)$. The equilibrium involves $p_A^* = u_A - \sigma$, $p_B^* = t + \frac{\sigma - \Delta}{3}$, $p_S^* = t - \frac{\sigma - \Delta}{3}$, $\pi_M^* = u_A - \sigma + 2t\lambda_B \left(\frac{1}{2} + \frac{\sigma - \Delta}{6t}\right)^2$ and $\pi_S^* = 2t\lambda_B \left(\frac{1}{2} - \frac{\sigma - \Delta}{6t}\right)^2$. It is straightforward to check that our assumptions on t imply S can earn a non-negative profit at these prices, both firms get some share of the B market, and the market is covered, and moreover that there is no profitable deviation for either firm.

Now suppose S is hosted by M . For the standard reasons, if $\lambda_A > \frac{\sigma}{u_A}$, M will set $p_A = u_A - \sigma$ and sell A to everyone, while if $\lambda_A \leq \frac{\sigma}{u_A}$, M will set $p_A = u_A$ and sell only to B -types. In either case, the equilibrium involves $p_B^* = t - \frac{\Delta}{3}$ and $\hat{p}_S^* = t + \frac{\Delta}{3}$. As a result, if M sets $p_A = u_A$, profit are $\pi_M = \lambda_B u_A + 2t\lambda_B \left(\frac{1}{2} - \frac{\Delta}{6t}\right)^2$ and $\pi_S^* = 2t\lambda_B \left(\frac{1}{2} + \frac{\Delta}{6t}\right)^2$, while if M sets $p_A = u_A - \sigma$, profits are $\pi_M^* = u_A - \sigma + 2t\lambda_B \left(\frac{1}{2} - \frac{\Delta}{6t}\right)^2$ and $\pi_S^* = 2t\lambda_B \left(\frac{1}{2} + \frac{\Delta}{6t}\right)^2$. Our assumption on t ensures the market is covered, both firms get some share of the B market, and there is no profitable deviation for each firm. Note checking that there is no profitable deviation also requires checking that S would never want to set $p_S < \hat{p}_S^*$ to induce some multi-stop shopping or some buyers to one-stop shop at S . Doing so will not attract any consumers to multi-stop shop unless $p_S < \hat{p}_S^* - \sigma$. Since all B -type consumers

buy A , to the extent they get some surplus from buying A , getting consumers to one-stop shop at S instead of at M will also require $p_S < \hat{p}_S^* - (u_A - p_A)$. In both cases, S could attract more additional consumers by lowering \hat{p}_S instead of p_S by the given amount. The fact it doesn't want to (i.e. that \hat{p}_S^* is the equilibrium level of p_S) implies it also cannot be better off lowering p_S below \hat{p}_S^* .

The proposition follows by comparing the joint profit worked out above under hosting with joint profit under non-hosting, taking into account the fixed cost of hosting F . ■

Note that the right-hand side in the tradeoff is always increasing in λ_A and F , which is consistent with the logic of the baseline model, namely that hosting is less likely for high λ_A and high F . If $\lambda_A \leq \frac{\sigma}{u_A}$, the right-hand side in the tradeoff is also increasing in u_A and decreasing in σ , which is also consistent with the logic in the baseline model. On the other hand, if $\lambda_A > \frac{\sigma}{u_A}$, the right-hand side in the tradeoff may be increasing or decreasing in σ , whereas in the baseline model it was always increasing. Finally, note the right-hand side of the tradeoff can be increasing or decreasing in the degree of product differentiation t when $\lambda_A < \frac{\sigma}{u_A}$ but is always increasing in the degree of product differentiation when $\lambda_A > \frac{\sigma}{u_A}$.

C Consumer surplus and welfare

We evaluate the effect of hosting on consumer surplus and welfare in the benchmark case without variable fees, which comes from a straightforward comparison of the equilibria defined in Propositions 1 and 2. The results are summarized in the following proposition.

Proposition 15 *Consider the baseline model. If $\lambda_A \leq \frac{\sigma}{u_A}$, hosting lowers consumer surplus, and it increases total welfare if and only if $\Delta > \frac{\lambda_A(u_A - \sigma) + F}{1 - \lambda_A}$. If $\lambda_A > \frac{\sigma}{u_A}$, hosting raises consumer surplus, and it increases total welfare if and only if $\Delta > \frac{F}{1 - \lambda_A}$.*

The only parameter region where consumers are better off with hosting is the region in which M would individually prefer not to host. The reason is that hosting constrains the ability of M to extract profit from product A because M can no longer price discriminate. Only when this constraint is sufficiently important can consumer surplus be higher. While hosting increases competition over product B , it may not increase overall competition for the benefit of consumers when both products are taken into account.

It is intuitive that hosting increases total welfare by eliminating the additional shopping cost for B -type consumers to get A and B_S . This gives B -types an additional utility of Δ compared to when they were buying A and B_M without hosting. Other than the fixed cost F , the only other downside of hosting occurs when M stops selling to A -types, which happens when Δ is not very high. In this case, welfare can be lower with hosting even in the absence of any fixed cost (i.e. $F = 0$).

Thus, it is possible that hosting is jointly profitable but leads to lower total welfare. This happens when $\lambda_A \leq \frac{\sigma}{u_A}$ and $\frac{\lambda_A}{2(1 - \lambda_A)}(u_A - \sigma) < \Delta < \frac{\lambda_A}{1 - \lambda_A}(u_A - \sigma)$. Conversely, it is possible that hosting is not jointly profitable but leads to higher total welfare: this happens when $\lambda_A > \frac{\sigma}{u_A}$ and $\Delta < \frac{\sigma}{2}$.