Creating platforms by hosting rivals

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

March 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

1 Introduction

Recently a lot of attention has been given to multi-sided platforms such as those operated by Airbnb, Amazon, Expedia, Facebook and Tencent, to name a few. In part, this reflects that many of the most valuable companies in the world 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, even if they don't realize it. The most straightforward way for a product company to do so is by inviting third parties to sell their products or services on top of the original

[*]MIT Sloan School of Management. E-mail: ahagiu@mit.edu
[†]Toulouse School of Economics. E-mail: bruno.jullien@tse-fr.eu.
[‡]Department of Economics, National University of Singapore. E-mail: jwright@nus.edu.sg

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 benefits 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), but bans others. Salesforce's Force.com allows customers to purchase CRM apps that directly compete with Salesforce's CRM product. 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 (e.g. payroll management).

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. A recent trend in the industry is for "big box gyms" to rent 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.[1] Country clubs work in a similar way, sometimes hosting third-parties that provide classes or specialized services

[1] https://www.newyorksportsclubs.com/cyc

to the clubs' members (e.g. swimming or tennis coaching) where these were previously (or sometimes still are) provided by the clubs. Other examples include Amazon's shift from being a reseller to opening up to third-party sellers of competing products (outside of its core products, like books), and cable TV providers which in many countries have allowed Netflix to sell to their subscribers through their own platforms, even though Netflix competes with the cable companies' video-on-demand services.

We provide a simple model which captures some of the key tradeoffs that arise when a traditional 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. In the gym example, 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 or by a specialist firm (i.e. Cyc). 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.

A key advantage of not hosting a rival in this setup is that it allows the multiproduct firm M to price discriminate between the consumers who want both products and those who just want A. In the equilibrium of our baseline setting, because of sufficiently high shopping costs, consumers who want both products buy them from M, rather than buying A from M and B from S. As a result, M can always extract the maximum amount from consumers who just want A by raising the price of A to their maximum willingness to pay and lowering the price of B to make sure consumers who want both products don't want to switch to S to buy B instead. Hosting eliminates the additional shopping cost consumers would face if they wanted to buy A from M and B from S. Essentially, with hosting, M unbundles the A and B products. This eliminates the ability of M to price discriminate across the two types of consumers, which is a key disadvantage of hosting that our model highlights.

By unbundling through the removal of the additional shopping costs, 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. On the other hand, hosting allows M to potentially gain by only selling to consumers who want both products and increasing its price on product A because shopping costs are now taken care of by the surplus offered by S's superior version of B. If there are enough such consumers, M gives up on selling to consumers who just want A and extracts more from consumers who want both products. In this sense, the presence of shopping costs allow M to gain from hosting by turning a substitute into a complement.

Our discussion up until this point has ignored two obvious additional factors in the hosting decision. One is that hosting may entail significant fixed costs, which one or both parties must cover. For instance, hosting a specialized cycling class in a gym may require re-arranging and customizing the space with the relevant equipment and brand, updating software systems for scheduling and reservations to include the specialized class, etc. 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. Obviously, these costs shift the tradeoff towards non-hosting. On the other hand, the hosting and hosted firm may be able to enter a financial contract that allows for payments between the firms. As a baseline we assume that the two firms can use a lump-sum payment and negotiate costlessly ex-ante, thus deciding whether to sign a hosting arrangement based on whether doing so increases their joint profits. Such a lump-sum payment means that the additional profit that S obtains under hosting (reflecting its quality advantage in product B) becomes relevant to the hosting decision, thus expanding the region of parameters where hosting dominates. In particular, hosting may be jointly preferred even if M continues to sell to consumers who just want A, since competition for product B may be less intense when firms compete within the market rather than for the market. This happens when the additional surplus offered by S's version of B is large relative to shopping costs, since then M cannot extract much from B without hosting but S can extract this full additional surplus when it is hosted. On the other hand, when M prefers to focus on buyers who want both products under hosting, our results show that hosting is jointly preferred when there is a higher fraction of consumers who want both products, when the additional surplus offered by S's version of B is higher, and when the surplus consumers get from buying A net of shopping costs is small.

We also provide several extensions of this simple benchmark, exploring how they shift the tradeoffs analyzed in the baseline model. In particular, we explore what happens when (i) there is correlation (positive or negative) in consumers' valuations across products A and B, (ii) consumers view M and S as offering horizontally rather than vertically differentiated versions of product B, (iii) there can be more than one specialist hosted on M, (iv) product B is an add-on good to product A, and (v) M can monitor and charge for transactions on its platform, so use a two-part tariff rather than just a lump-sum fee.

2 Related literature

Our paper relates to several strands of literature.

Some existing papers consider the tradeoff between the platform business model and more traditional alternatives. For example, Hagiu and Wright (2015a) considers an intermediary that chooses between functioning as a reseller (by purchasing products from suppliers and selling them to buyers) and functioning as a marketplace (in which suppliers sell their products directly to buyers). Similarly, Hagiu and Wright (2015b and 2017) consider a service firm that chooses between employing professionals to provide the service to its customers on terms completely controlled by the firm, or enabling independent professionals to provide the service directly to customers on terms controlled by these professionals. Other papers have similarly compared the "agency" business model in which "suppliers" set prices vs. the traditional wholesale pricing model in which the retailer sets prices (e.g. Abhishek et al., 2016, Johnson, 2017). In all these papers, the key distinction 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 key 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). For instance, Hagiu and Wright (2015a) study Amazon's choice to function as a re-seller or a marketplace for a given product, while ignoring the fact that in reality Amazon competes as a reseller against the independent sellers of the same product present on Amazon's marketplace. Thus, this strand of literature does not address the issue of hosting rivals, the central theme of the current paper.

To some extent, the platform as modeled in the current paper can be viewed as a vertically integrated firm that uses the upstream input (product A in our model) to offer downstream products (the various versions of product B). The vertically integrated firm can consider letting rival downstream firms (S in our model) access its upstream facility (i.e. sell on/through product A). 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 (See Rey and Tirole, 2007, for a summary). The literature focuses on the role of the wholesale tariff, which may be linear (e.g., Ordover et al., 1990, Chen, 2001, Sappington, 2005, Ordover and Shaffer, 2007, Hoeffler and Schmidt, 2008, and Bourreau at al., 2011) or non-linear (e.g., Hart and Tirole, 1990, O'Brien and Shaffer, 1992). Our setting is different in several respects. First, the platform is not essential but allows consumers to save on shopping costs, so the hosted firm can still sell the good outside the platform. Second, 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. Although we also discuss the case of two-part wholesale tariffs, most of our analysis assumes no payment or only a fixed transfer between the downstream competitor and the multiproduct firm.

Since the platform allows consumers to save on shopping costs, products sold on the platform are complementary to one another (buying one product raises the value of the other product). On the other hand, products sold outside the platform are simply competing with the same product on the platform. This links our work to the literature on the choice of compatibility between system components. From this perspective, hosting product B_S offered by the specialist S bears some similarity with making B "compatible" with product A. As in Matutes and Regibeau (1988), "compatibility" allows the platform to relax competition and raise the price of good A. A key difference is that the platform in our model does not constitute a "system" in the sense that products offered on the platform can be consumed independently and can also be offered outside the platform—in this context, compatibility simply raises the value of joint consumption.

One can also interpret hosting in our model as unbundling the sale of products A and B by the platform (when the platform cannot charge a variable transfer fee to the hosted rival). As a result, the tradeoffs we identify between hosting and non-hosting are related to the tradeoffs between bundling and unbundling at a very general level (see Whinston, 1990).

Finally, our model setup is quite closely related to Chen and Rey (2012), who provide a new theory of loss-leader pricing in a setting where consumers choose between one-stop shopping and multi-stop shopping. A key difference is that in our paper consumers differ in their valuations for product B, whereas Chen and Rey assume consumers differ in their shopping costs. Furthermore, for the most part, they focus on the case when there is a competitive fringe of specialist firms. Most importantly, Chen and Rey do not consider the possibility of hosting. We have attempted to analyze hosting in their framework, but because of heterogeneous shopping costs, the case in which there is a single specialist firm (which we have focused on) is not very tractable in their framework.

3 Benchmark 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 A is monopolized by M while B can be supplied by M or S. We normalize both firms' costs to zero.

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 0). We call these monoproduct consumers "A-type" consumers. A fraction $\lambda_B > 0$ (which equals $1 - \lambda_A$) of consumers are multiproduct consumers who 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 $u_A > 0$. The B-types value B_M at $u_B > 0$ and B_S at $u_S = u_B + \Delta u_B$, where $\Delta u_B \geq 0$. Thus, B-types view B_S as superior to B_M.

All consumers face a shopping cost $\sigma > 0$ of 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 shopping"), they will incur σ twice. Consumers can always purchase an outside option which gives them a payoff normalized to 0. Throughout the paper we assume that

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

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

$$\Delta u_B \leq \sigma,$$

i.e. that the shopping cost exceeds the added value of S's product B_S. This ensures that without hosting, M can have an advantage in selling product B, provided it makes it attractive for consumers to want to buy A. We will note later what happens if $\Delta u_B > \sigma$.

3.1 Without hosting

In this section, we characterize the equilibrium that arises when S sells directly in competition with M.

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

Because the shopping cost σ outweighs S's added value in product B (i.e. Δu_B), there is a unique equilibrium outcome in which M makes all the sales.[2] 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 the prices are $p_A^* = u_A - \sigma$, $p_B^* = \sigma - \Delta u_B$, $p_S^* = 0$. The A-type consumers always purchase A, and the B-type consumers all buy A and B_M from M. Profits are $\pi_M^* = u_A - \sigma + \lambda_B (\sigma - \Delta u_B)$ 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 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, 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 $u_A - \Delta u_B$.

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., $p_A + p_B = u_A - \Delta u_B > u_A - \sigma$). Note M achieves this outcome by charging $p_A^* = u_A - \sigma$ and $p_B^* = \sigma - \Delta u_B$. This represents 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 $p_B = \sigma - \Delta u_B$ as in Proposition 1 or (ii) $u_A - \sigma < p_A \leq u_A$ and $p_B = u_A - \Delta u_B - p_A$, 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 constrains M's price p_A and eliminates the range $u_A - \sigma < p_A \leq u_A$.

[2] In principle, there are other equilibria in which $p_S^* < 0$. However, setting such prices is always weakly dominated for S (it is never profitable). Therefore, we rule out such equilibria.

3.2 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. We still allow S to sell directly, at price p_S. Meanwhile, let $\widehat{p}_S$ denote the price S charges when it sells B_S through M.

We assume there is a fixed cost of hosting, denoted $F > 0$. 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. 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 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 directly is less profitable for S than selling through M.

By removing the additional shopping cost for consumers who want to buy A and B_S, hosting puts both firms on an even playing field. Since $\Delta u_B > 0$, in equilibrium S always wins the competition in the B market, and sells B_S to all B-types at $\widehat{p}_S = \Delta u_B$, 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. Given this, 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

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

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.[3] A formal proof of the proposition follows as a special case of the proof of Proposition 13 in which the transaction fee τ is set to zero.

Proposition 2 *In the baseline model with hosting, there are two cases to consider. If $\lambda_A < \frac{\sigma}{u_A}$, the equilibrium prices are $p_A^* = u_A$, $p_B^* = 0$, and $\widehat{p}_S^* = p_S^* = \Delta u_B$. 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 $\pi_S^* = \lambda_B \Delta u_B$. If $\lambda_A > \frac{\sigma}{u_A}$, the equilibrium prices are $p_A^* = u_A - \sigma$, $p_B^* = 0$, and $\widehat{p}_S^* = p_S^* = \Delta u_B$. 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 $\pi_S^* = \lambda_B \Delta u_B$.*

The margin Δu_B that S obtains on B-types reflects that under hosting, with no shopping cost disadvantage, S has a competitive advantage of Δu_B in selling B, which it can fully extract. By

[3] Under this setting there are multiple equilibria, reflecting that there are also equilibria in which $p_B^* < 0$ given that M does not sell B_M in equilibrium, and equilibria in which $p_S^* < \widehat{p}_S^*$ given that S does not sell B_S directly in equilibrium. Among these different 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) that they would prefer to change if some consumers actually purchased from them (i.e. off the equilibrium path). We follow a similar approach whenever there are multiple equilibria in the various extensions we consider.

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.

Note that under hosting, M would not want to bundle A and B_M. The A-types get no additional value from the bundle, while for B-types, they all want to buy A provided $p_A + p_B \leq u_A$, and given this, they also want to buy B_S provided $\widehat{p}_S \leq \Delta u_B$. With bundling, M chooses between setting $p_A + p_B = u_A - \sigma$, selling the bundle to all consumers, or setting $p_A = u_A$ and sells the bundle to B-type consumers only, but in either case it obtains no more profit from bundling than in the equilibrium described in the proposition above. Moreover, if there were some positive marginal cost associated with M providing B_M, bundling would make M strictly worse off since it provides B_M to A-types, who are not paying anything for it.

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. Thus, the only way M would not offer B_M is if it could commit *ex-ante* to not offer it. This may not be possible in practice if it requires M to write a contract specifying that it will not compete with S. 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 M could commit to remove B_M, it has no unilateral incentive to do so. Consider the possible multiple equilibria that can arise in this case from the fact that both M and S are now monopolists, for A and B_S respectively. The equilibria are defined by $p_A^* = u_A - \omega\sigma$ and $\widehat{p}_S^* = u_B + \Delta u_B - (1 - \omega)\,\sigma$, where either $\omega = 1$ (and M sells A to all consumers) or $0 \leq \omega \leq \frac{\sigma - \lambda_A u_A}{\sigma - \lambda_A \sigma}$ (and M sells A only to B-types, where the upper bound on ω ensures that M does not want to lower its price p_A so as to also sell to A-types). Then M's profit is either $u_A - \sigma$ or $\lambda_B\,(u_A - \omega\sigma) \leq \lambda_B u_A$. Clearly, M does not do any better in this equilibrium compared to the case it sells B_M as well.[4]

3.3 Comparison of hosting with non-hosting

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.

[4]Later we will explore whether M would want to commit to remove B_M from a joint profit perspective. To the extent this gives S a monopoly over B-types, it would seem to always increase joint profits in our baseline model. However, this does not seem realistic in practice, since other firms would then want to enter to sell B. Therefore, we explore this question in Section 4.3, where we allow for multiple specialists.

Proposition 3 *Consider the baseline model. If* $\lambda_A < \frac{\sigma}{u_A}$*, then hosting is preferred by* M *iff* $\Delta u_B > \frac{\lambda_A(u_A-\sigma)+F}{1-\lambda_A}$*. If* $\lambda_A > \frac{\sigma}{u_A}$*, then non-hosting is preferred by* M*.*

The proposition is illustrated in the left-hand panel of Figure 1. It is easiest to interpret this result when there are almost no A-types, i.e. when $\lambda_A \to 0$. Then, hosting allows M to increase its profit by Δu_B, so hosting is always preferred provided Δu_B exceeds the fixed cost F of hosting. The gain of Δu_B that M obtains from hosting comes from a gain of σ on the A product and a loss of $\sigma - \Delta u_B$ 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 gives M a gain of σ on the A product. This is the sense in which the presence of shopping costs allows M to gain by turning a substitute into a complement. On the other hand, under hosting M no longer extracts $\sigma - \Delta u_B$ from its sale of B_M—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.

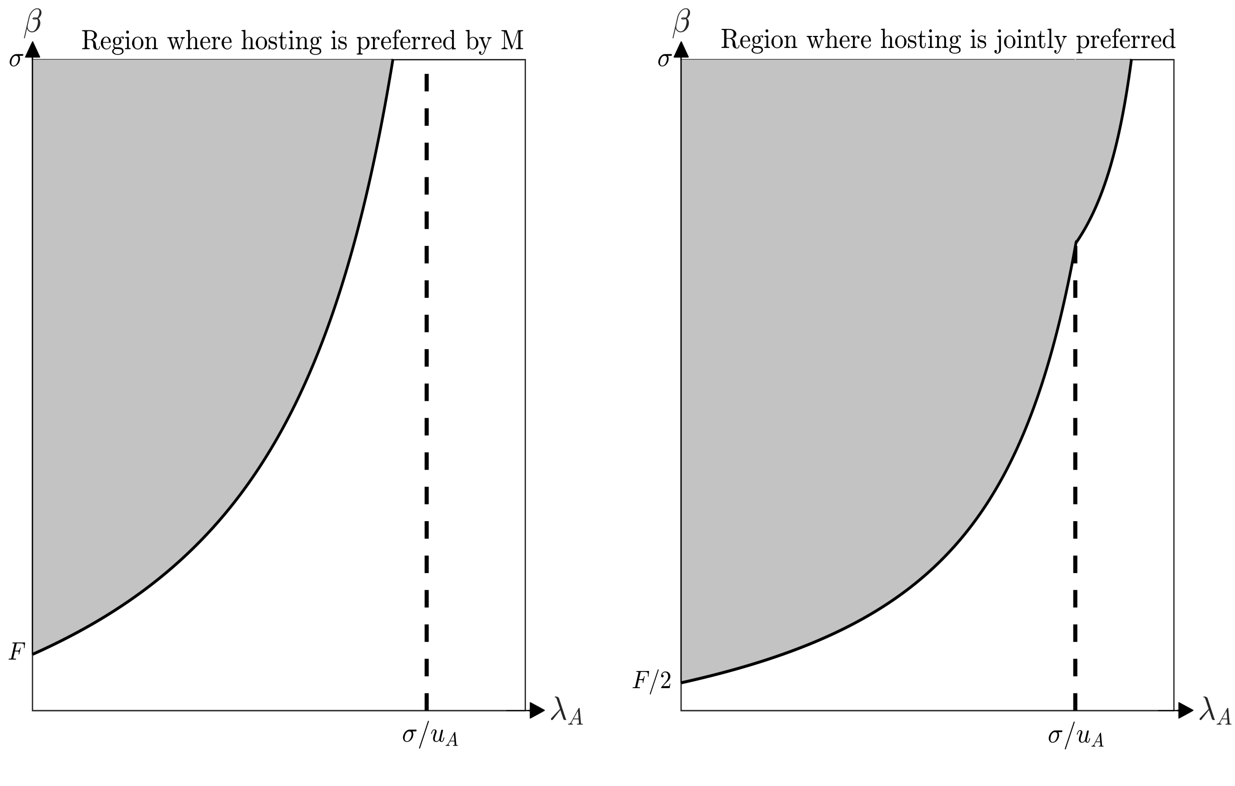

Figure 1: Regions where hosting is preferred (shaded) and where non-hosting is preferred (non-shaded)

If we just focus on extraction of surplus on product B, note that before hosting, M extracts $\sigma - \Delta u_B$, which reflects its efficiency advantage on B given it already sells A, while after hosting, it extracts nothing on product B. Thus, M is worse off by $\sigma - \Delta u_B$ as a result of hosting. However, without A types, hosting always allows M to extract more surplus from product A (i.e. its price is σ

higher) as consumers' shopping cost is now covered by the surplus they get from buying B_S, and this always dominates the loss in surplus extracted on product B.

Now consider what happens when there are some (but not too many) A-types, so $\lambda_A < \frac{\sigma}{u_A}$. 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 from product A under hosting may no longer dominate the negative effect of hosting on its profit in the B market. Put differently, the introduction of A-types constrains M's ability to extract more for product A from B-types under hosting since A-types don't get the extra surplus generated from hosting S's superior version of product B, thus raising the possibility that M is 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 $\sigma - \Delta u_B$) is large, i.e. when Δu_B 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 doesn't want to give up on selling to the A-types, and so there is no gain on A-types to offset the loss on B-types. In this case, hosting always lowers M's profit.

An alternative way to understand this logic is to note that M does better if it can price discriminate between A-type and B-type consumers. This is achieved by charging a low enough price for product A so as to cover the shopping cost incurred by A-type consumers, given that their shopping cost cannot be covered by the surplus obtained on another product. M can then compensate for this lower p_A when selling to B-types by charging a higher p_B. Without hosting, since M sells the bundle to B-types, it can do exactly this, so it always extracts the maximum amount from A-types. With hosting, because of unbundling, M can no longer achieve this form of price discrimination, which is the cost of hosting.

3.3.1 Joint profits

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, and negotiate costlessly. 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 exploring the impact of hosting on the firms' joint profit.[5] In this baseline setting, we obtain the following result.

Proposition 4 *Consider the baseline model. If* $\lambda_A < \frac{\sigma}{u_A}$*, then hosting is jointly preferred iff* $\Delta u_B > \frac{\lambda_A(u_A-\sigma)+F}{2(1-\lambda_A)}$*. If* $\lambda_A > \frac{\sigma}{u_A}$*, then hosting is jointly preferred iff* $\Delta u_B > \frac{\sigma}{2} + \frac{F}{2(1-\lambda_A)}$*.*

Proposition 4 is illustrated in the right hand panel of Figure 1. Comparing the two panels, it is clear that the tradeoff between hosting and non-hosting in terms of joint profits is similar to one when

[5]Later we will explore what happens when M can monitor sales and so charge S a per-transaction fee as well.

we focused only on M's profit. The difference is that now under hosting we must add S's profit, which is $\lambda_B \Delta u_B$. As figure 1 illustrates, this expands the region of Δu_B where hosting dominates.

Focusing once again on the extraction of the surplus created by product B, note that before hosting, M extracts $\sigma - \Delta u_B$, which reflects its efficiency advantage on B given that it already sells A, while after hosting S extracts Δu_B. Thus, the condition for joint profit on product B to be higher after hosting is for $\Delta u_B > \sigma - \Delta u_B$, or $\Delta u_B > \frac{\sigma}{2}$. Thus, there is now a sense in which the intensity of competition on product B may either increase or decrease as a result of hosting, with competition relaxed if S's efficiency advantage is high relative to the M's shopping cost advantage. In this case, which arises if $\Delta u_B > \frac{\sigma}{2}$, hosting dominates provided F is sufficiently small (regardless of λ_A).[6] However, even when $\Delta u_B < \frac{\sigma}{2}$, if there are few A-types (i.e. λ_A is low), hosting always allows M to extract more surplus from product A (the price p_A is higher by σ) as consumers' shopping cost is now covered by the surplus they get from buying B_S. This effect dominates the negative effect of increased competition in product B, so hosting can still dominate.

As before, the existence of A-types constrains M's ability to extract the additional surplus from product A from B-types, and so leads to the possibility that the additional surplus extracted from product A under hosting may no longer dominate the negative effect of hosting on joint profits in the B market which arises when Δu_B is small relative to σ.

3.3.2 Consumer surplus and welfare

We can also evaluate the effect of hosting on consumer surplus and welfare, which comes from a straightforward comparison of the equilibria defined in Propositions 1 and 2. The results are summarized in the following proposition.

Proposition 5 *Consider the baseline model. If $\lambda_A < \frac{\sigma}{u_A}$, hosting lowers consumer surplus, but increases total welfare if and only if $\Delta u_B > \frac{\lambda_A(u_A - \sigma) + F}{1 - \lambda_A}$. If $\lambda_A > \frac{\sigma}{u_A}$, hosting increases consumer surplus and increases total welfare if and only if $\Delta u_B > \frac{F}{1-\lambda_A}$.*

Thus, 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 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 Δu_B compared to when they were buying A and B_M without hosting. Other than the fixed cost F, the only other downside

[6] This result extends to the case we ruled out by assumption, namely $\Delta u_B > \sigma$, 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 $\Delta u_B > \sigma$, hosting is jointly preferred if $\lambda_B \sigma + \max\{\sigma - \lambda_A u_A, 0\} > F$.

of hosting occurs when M stops selling to A-types, which happens when Δu_B 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 $\lambda_A < \frac{\sigma}{u_A}$ and $\frac{\lambda_A}{2(1-\lambda_A)}(u_A - \sigma) < \Delta u_B < \frac{\lambda_A}{1-\lambda_A}(u_A - \sigma)$. 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 $\Delta u_B < \frac{\sigma}{2}$.

4 Extensions

In this section we consider several extensions of the baseline model.

4.1 Correlation in consumers' valuations across products

So far we have 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. 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).

4.1.1 Negative correlation

Suppose $0 < \alpha \leq \sigma$, so A-types are willing to pay α more for product A than are B-types.[7] 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 6 *If $\lambda_A < \frac{\sigma-\alpha}{u_A}$, then hosting is jointly preferred iff $\Delta u_B > \frac{\lambda_A(u_A+\alpha-\sigma)+F}{2(1-\lambda_A)}$. If $\lambda_A > \frac{\sigma-\alpha}{u_A}$, then hosting is jointly preferred iff $\Delta u_B > \frac{\sigma-\alpha}{2} + \frac{F}{2(1-\lambda_A)}$.*

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 Δu_B when M still sells to A-types (this

[7] The restriction that $\alpha \leq \sigma$ is used to ensure there is always a pure strategy equilibrium. Otherwise, there can be a range of intermediate values of λ_A for which only a mixed strategy equilibrium exists.

occurs for large α), but dominates for a smaller range of Δu_B 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.

4.1.2 Positive correlation

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

Proposition 7 *If* $\lambda_A < -\frac{\alpha}{u_A-\sigma}$*, then hosting is jointly preferred iff* $\Delta u_B > \frac{F}{2(1-\lambda_A)}$*. If* $-\frac{\alpha}{u_A-\sigma} < \lambda_A < \frac{\sigma-\alpha}{u_A}$*, then hosting is jointly preferred iff* $\Delta u_B > \frac{\lambda_A(u_A-\sigma+\alpha)+F}{2(1-\lambda_A)} + \frac{\alpha}{2}$*. If* $\lambda_A > \frac{\sigma-\alpha}{u_A}$*, then hosting is jointly preferred iff* $\Delta u_B > \frac{\sigma}{2} + \frac{F}{2(1-\lambda_A)}$*.*

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 A-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.

4.2 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 $u_B + \Delta u_B$ 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). In the appendix we prove the following proposition.

Proposition 8 *Suppose there is horizontal differentiation for product B, with the mismatch parameter t satisfying* $\max\left(\frac{\sigma-\Delta u_B}{3}, \frac{\Delta u_B}{3}\right) < t < \frac{2u_B}{3} + \min\left\{\frac{\sigma-\Delta u_B}{9}, \frac{\Delta u_B}{3}\right\}$. *When $\lambda_A < \frac{\sigma}{u_A}$, hosting is jointly preferred iff $\Delta u_B > \frac{\sigma}{2} + \frac{9t(\lambda_A u_A - \sigma + F)}{2\sigma(1-\lambda_A)}$. When $\lambda_A > \frac{\sigma}{u_A}$, hosting is jointly preferred iff $\Delta u_B > \frac{\sigma}{2} + \frac{9tF}{2\sigma(1-\lambda_A)}$.*

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 $\lambda_A < \frac{\sigma}{u_A}$, 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 $\lambda_A > \frac{\sigma}{u_A}$.

4.3 Competing specialist firms

Consider the variation from our baseline model in which there are multiple (two or more) identical specialist firms S: each sells B_S, which gives B-type consumers a utility of $u_B + \Delta u_B$. This case allows us to understand how competition among specialists changes our results, as well as to explore an additional strategic decision of M—whether to host one or multiple specialists. Finally, it provides a more realistic environment in which to explore whether M should give up selling B_M under hosting, since the pricing of S on M will now be constrained by competing specialists (either because they are also hosted on M or because they will still try to sell directly to consumers).

The non-hosting equilibrium remains unchanged: it involves B_S being priced at marginal cost (i.e., at zero), both in the baseline model and here. We therefore focus on the case with hosting, which can now involve hosting one or two firms. We will consider each scenario in turn.

If M hosts one specialist firm, the equilibrium in Proposition 2 for $\lambda_A < \frac{\sigma}{u_A}$, in which $p_A^* = u_A$, $p_B^* = 0$ and $\widehat{p}_S^* = \Delta u_B$, no longer applies. This is because the non-hosted specialists can profitably attract B-type consumers, who would get $u_B + \Delta u_B - \sigma - p_S$ buying directly, rather than $u_A + u_B + \Delta u_B - p_A^* - \widehat{p}_S^* - \sigma = u_B - \sigma$ buying A from M and B_S from the hosted specialist. On the other hand, when $\lambda_A > \frac{\sigma}{u_A}$, the equilibrium in Proposition 2 involves $p_A^* = u_A - \sigma$, so the non-hosted specialists cannot compete away consumers from the hosted specialist. Indeed, in this case B-types obtain u_B buying A from M and B_S from the hosted specialist, whereas they would get $u_B + \Delta u_B - \sigma - p_S$ buying directly, and the latter utility is less than u_B for all non-negative p_S (recall that $\Delta u_B < \sigma$). Thus, in this case the previous equilibrium outcome continues to apply. We obtain the following characterization of the equilibrium.

Proposition 9 *When there are multiple competing specialist firms and only one is hosted, there are two cases to consider. If $\lambda_A < \frac{\sigma}{u_A}$, there is a unique equilibrium outcome in which the prices are $p_A^* = u_A$, $p_B^* = 0$, $\widehat{p}_S^* = 0$ and $p_S^* = 0$ (the last price is that of the non-hosted specialists who sell* *directly). The A-type consumers do not purchase, while the B-type consumers all buy A and B_S through M. Profits are $\pi_M^* = \lambda_B u_A$ and $\pi_S^* = 0$. If $\lambda_A > \frac{\sigma}{u_A}$, there is a unique equilibrium in which the prices are $p_A^* = u_A - \sigma$, $p_B^* = 0$ and $\widehat{p}_S^* = \Delta u_B$. The A-type consumers purchase A, and the B-type consumers all buy A and B_S through M. Profits are $\pi_M^* = u_A - \sigma$ and $\pi_S^* = \lambda_B \Delta u_B$ (for the hosted specialist).*

Alternatively, M can host two or more specialists, who will always compete their price down to marginal cost, implying $\widehat{p}_S^* = 0$. In this case, the existence of any remaining outside specialists is irrelevant. Since M makes no profit on product B, its pricing decision for A remains exactly as before. We obtain the following characterization of the equilibrium.

Proposition 10 *When there are competing specialist firms and two or more are hosted, there are two cases to consider. If $\lambda_A < \frac{\sigma}{u_A}$, there is a unique equilibrium outcome in which the prices are $p_A^* = u_A$, $p_B^* = 0$, $\widehat{p}_S^* = 0$. The A-type consumers do not purchase, while the B-type consumers all buy A and B_S through M. Profits are $\pi_M^* = \lambda_B u_A$ and $\pi_S^* = 0$. If $\lambda_A > \frac{\sigma}{u_A}$, there is a unique equilibrium in which the prices are $p_A^* = u_A - \sigma$, $p_B^* = 0$, $\widehat{p}_S^* = 0$. The A-type consumers purchase A, and the B-type consumers all buy A and B_S through M. Profits are $\pi_M^* = u_A - \sigma$ and $\pi_S^* = 0$.*

The results in the last two propositions imply that the joint gain from hosting (i.e. the gain for M and the hosted specialist firms) is the same regardless of whether M hosts one or multiple specialists when M gives up on A-types (i.e. when $\lambda_A < \frac{\sigma}{u_A}$). However, when M keeps serving A-types under hosting (i.e. when $\lambda_A > \frac{\sigma}{u_A}$), both M's profits and joint profits are higher when M hosts one rather than multiple specialists. Thus, the joint choice between hosting and non-hosting becomes as follows.

Proposition 11 *Suppose there are competing specialist firms. If $\lambda_A < \frac{\sigma}{u_A}$, then hosting one or multiple specialists leads to the same joint profits, and either option is jointly preferred to non-hosting iff $\Delta u_B > \frac{\lambda_A(u_A-\sigma)+F}{1-\lambda_A}$. If $\lambda_A > \frac{\sigma}{u_A}$, then hosting a single specialist is jointly preferred to non-hosting iff $\Delta u_B > \frac{\sigma}{2} + F$, while hosting multiple specialists is always dominated.*

Finally, we can explore whether joint profits (i.e. M's profits plus the profits of any hosted specialists) would be higher if M could commit to remove B_M under hosting. Clearly, such a commitment has no value when M hosts multiple specialists, since they make zero profits. Thus, the only interesting scenario is when M hosts one specialist and can commit to permanently remove B_M. In the resulting equilibrium, the remaining outside specialists (one or several) will price at zero and the hosted specialist S must price so that consumers prefer buying A and B_S on M to either buying A on M and B_S directly, or just buying B_S directly. This requires $\widehat{p}_S \leq \min\{\sigma, u_A - p_A\}$. It is then easily seen that the equilibria are defined by $p_A^* = u_A - \omega\sigma$ and $\widehat{p}_S^* = \omega\sigma$, where either $\omega = 1$ and M sells A to all consumers, or $0 \leq \omega \leq \max\left\{\frac{\sigma - \lambda_A u_A}{\sigma - \lambda_A \sigma}, 0\right\}$ and M sells A only to B-types (the upper bound on ω is defined so that M does not want to lower its price for A to sell to A-types as well).

Thus, M's profit is either $u_A - \sigma$ or $\lambda_B (u_A - \omega\sigma) \leq \lambda_B u_A$, and the hosted S's profit is $\lambda_B \omega\sigma$, so joint profits are either $u_A - \sigma + \lambda_B \sigma$ or $\lambda_B u_A$ depending on whether the equilibrium with $\omega = 1$ is selected, or the equilibrium with $0 \leq \omega \leq \max\left\{\frac{\sigma - \lambda_A u_A}{\sigma - \lambda_A \sigma}, 0\right\}$. Note that $u_A - \sigma + \lambda_B \sigma > \lambda_B u_A$ given $u_A > \sigma$, so joint profits are maximized by selecting the equilibrium with $\omega = 1$. Compare this to the joint profit obtained under hosting when M keeps selling B_M, which are $\lambda_B u_A$ when $\lambda_A < \frac{\sigma}{u_A}$ and $u_A - \sigma + \lambda_B \Delta u_B$ when $\lambda_A > \frac{\sigma}{u_A}$. This implies if $\lambda_A < \frac{\sigma}{u_A}$, removing B_M either does not change joint profits or increases it (if the equilibrium with $\omega = 1$ is selected), and if $\lambda_A > \frac{\sigma}{u_A}$, then removing B_M either decreases joint profits (if the equilibrium with $0 \leq \omega \leq \frac{\sigma - \lambda_A u_A}{\sigma - \lambda_A \sigma}$ is selected) or increases it (if the equilibrium with $\omega = 1$ is selected). Thus, depending on λ_A and equilibrium selection, we can have that removing B_M makes no difference to joint profits, increases them, or decreases them. However, if we select the equilibrium that maximizes joint profits ($\omega = 1$), then removing B_M unambiguously increases joint profits and hosting unambiguously dominates non-hosting.

4.4 Add-on goods

Consider the variation from our baseline model in which product B is an add-on good. Specifically, suppose B-types only get the value u_B from B_M or $u_B + \Delta u_B$ from B_S if they also purchase A. This changes the previous analysis by removing the option of one-stop shopping for B at either firm. As a result, A becomes like an essential input, which consumers must purchase to get any value from B.

Proposition 12 *Suppose product B is an add-on good. If $\lambda_A < \frac{u_B + \Delta u_B - \sigma}{u_B + \Delta u_B - \sigma + u_A - \sigma}$, hosting is jointly preferred iff $\Delta u_B > \frac{F}{1 - \lambda_A}$. If $\frac{u_B + \Delta u_B - \sigma}{u_B + \Delta u_B - \sigma + u_A - \sigma} < \lambda_A < \frac{u_B}{u_B + u_A - \sigma}$, hosting is jointly preferred iff $\Delta u_B > \frac{\lambda_A (u_A - \sigma) - \lambda_B (u_B - \sigma) + F}{2(1 - \lambda_A)}$. If $\lambda_A > \frac{u_B}{u_B + u_A - \sigma}$, hosting is jointly preferred iff $\Delta u_B > \frac{\sigma}{2} + \frac{F}{2(1 - \lambda_A)}$.*

If M is willing to give up on selling to A-types, it is then able to extract the net surplus B-types obtain from buying product B. This is achieved by setting $p_A > u_A$. Thus, M and S together extract the full consumer surplus from B-types, both with and without hosting. As a result, the condition for hosting to be jointly preferred is just that the added value of S's product, $\lambda_B \Delta u_B$, exceeds the fixed cost of hosting F. This explains the first result in Proposition 12. For high enough λ_A, M always prefers to sell to A-types, in which case p_A is constrained to $u_A - \sigma$, and the tradeoff remains the same as in the benchmark case. This explains the third result in Proposition 12. For intermediate values of λ_A, M prefers to sell to A-types without hosting but not with hosting, which explains the more complicated cutoff for Δu_B under which hosting is jointly preferred in the second result in Proposition 12.

4.5 Linear fees

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

The timing remains as before: after the contract is specified (including linear and lump-sum fees), the two firms set their prices simultaneously, taking the linear 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.

We obtain the following result.

Proposition 13 *Consider the baseline model in which in the hosting contract, M specifies a per-transaction fee to charge S, and a lump-sum transfer. Hosting is jointly preferred if and only if*

$$\Delta u_B \geq \frac{F}{2(1-\lambda_A)}.$$

Comparing Proposition 13 with Proposition 4, it is straightforward to see that the possibility of using per-transaction fees shifts the tradeoff towards hosting. In fact, 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 transaction fees. The reason is that the fee τ gives M another instrument, allowing price discrimination between A-types and B-types whenever M chooses to keep selling to A-types. This gives M the same benefit as non-hosting, in being able to price discriminate, but now with the higher surplus offered to consumers thanks to S's superior product, which M can extract. Meanwhile, M can use p_B to control any double marginalization problem that could otherwise arise with linear fees. Nevertheless, with $F > 0$, the key tradeoff determining whether hosting will be chosen is still increasing in Δu_B and decreasing in λ_A as in our benchmark setting.

One reason why linear fees may not be as effective in shifting the tradeoff towards hosting as suggested by Proposition 13 is that S may not be able (or willing) to price discriminate across the two channels under hosting (selling on M and selling directly). To capture this realistic issue, suppose S also continues to offer B_S directly, in competition with one or more other specialists who offer B_M in the outside market for a large set of consumers who just want to buy B through a specialist. The competitive direct price set by S is $p_S = \Delta u_B$, which allows it to capture all consumers in the outside market, reflecting the usual asymmetric Bertrand logic (given it offers surplus of $u_B + \Delta u_B$, and the rival specialists offer just u_B).

Assume S is not willing to set $\widehat{p}_S > p_S$ because doing so could have negative spillovers on the larger outside market. As evidence of this, Cyc sets the same price for its classes throughout all of its New York City locations, including at NYSC. Under this assumption, if S wanted to raise $\widehat{p}_S$ above Δu_B, then it also would have to raise p_S above Δu_B. But this would mean giving up on the profit made on the outside market, which we assume is more important for S than the profit it can make on M. Thus, S will not want to set $\widehat{p}_S > \Delta u_B$. Furthermore, even though S is free to lower both prices below Δu_B, it will not want to do so because that would result in strictly lower profits. As a result, S always sets $\widehat{p}_S = p_S = \Delta u_B$, regardless of the linear fee τ charged by M. This means that τ simply acts to transfer revenues from S to M, just like a lump-sum fee. Thus, in this case the equilibrium joint profit with hosting is the same as in our benchmark case. However, non-hosting is now more profitable, because M is competing with a firm that charges $p_S = \Delta u_B$ for outside sales rather than $p_S = 0$ in the benchmark case. We obtain the following result.

Proposition 14 *Suppose that due to competition in the outside market, S sets $p_S = \Delta u_B$ and due to concerns regarding price discrimination, S cannot set $\widehat{p}_S > p_S$. If $\lambda_A < \frac{\sigma}{u_A}$, then hosting is jointly preferred iff $\Delta u_B > \frac{\lambda_A(u_A-\sigma)+F}{(1-\lambda_A)}$. If $\lambda_A > \frac{\sigma}{u_A}$, then non-hosting is always jointly preferred.*

Interestingly, the tradeoff is now identical to the unilateral incentive of M to adopt hosting in the benchmark model with no transfers (i.e. Proposition 3).

Another scenario in which linear fees may not be as effective in shifting the tradeoff towards hosting as suggested by Proposition 13 is when some of M's consumers are actually captive with respect to product B_M in the absence of hosting (perhaps due to lack of information about alternatives outside M), but hosting opens up the possibility of them buying from S. For example, many New York Sports Club members may not know about the existence of Cyc before it is hosted, but would learn about Cyc if it is hosted. Consider the extreme case in which all B-types fit this characterization. Without hosting, such consumers are captive, and M maximizes its profit by setting $p_A = u_A - \sigma$ and $p_B = u_B$, fully capturing their surplus. This gives M a profit of $\pi_M = u_A - \sigma + \lambda_B u_B$, with S obtaining no profit. Assuming linear fees can be used, the equilibrium with hosting (in which all consumers become aware of B_S offered by S) is the same as that characterized in the proof of Proposition 13, in which joint profits are $u_A - \lambda_A\sigma + \lambda_B\Delta u_B - F$. In this scenario, hosting is jointly preferred iff

$$\Delta u_B > u_B - \sigma + \frac{F}{1-\lambda_A}.$$

Thus, for hosting to be jointly preferred, the efficiency benefit of B_S over B_M must be significantly larger than what is required by Proposition 13. Moreover, even if there are no fixed costs associated with hosting (i.e. $F = 0$), M may prefer not to host S.

A less extreme version of the previous scenario would be to assume that B-types know about the existence of B_S before it is hosted, but do not realize that it is superior to B_M. Thus, before hosting, B-type consumers act as if $\Delta u_B = 0$, whereas hosting makes them aware that $\Delta u_B > 0$. In this case, without hosting, given $p_S = 0$, M maximizes its profit by setting $p_A = u_A - \sigma$ and $p_B = \sigma$, which yields profits $\pi_M = u_A - \sigma + \lambda_B\sigma$ and $\pi_S = 0$. Assuming linear fees can be used, the equilibrium with hosting (in which all consumers become aware that S offers higher utility) is the same as that characterized in the proof of Proposition 13, in which joint profits are $u_A - \lambda_A\sigma + \lambda_B\Delta u_B - F$. We obtain that hosting is jointly preferred iff

$$\Delta u_B > \frac{F}{1-\lambda_A}.$$

While the tradeoff shifts towards non-hosting when compared to Proposition 13, joint profits are still higher with hosting whenever $F = 0$. This shows that the previous, more dramatic form of learning about B_S is required in order for non-hosting to sometimes be jointly preferred in the absence of fixed costs of hosting.

5 Conclusion

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. 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.

We plan to extend the paper to explore other fundamental reason why hosting may not be profitable even if the firms can use per-transaction fees and/or the hosting firm M can commit to remove its version of B. One further possibilities includes that hosting may subject each party to a hold-up problem to the extent they each need to incur some non-recoverable fixed costs of setting up so that S can sell through M. After incurring these costs, there is the possibility for one of the parties to renegotiate the contract, which ex-ante may make hosting harder to achieve depending on the initial bargaining situation. We also plan to consider what happens if instead of the shopping cost being incurred by consumers, each firm needs to incur a cost of reaching customers. Then hosting can be modeled as the elimination of the additional cost of reaching customers on the platform. Thus, if a specialist wants to reach customers outside the platform, it incurs the cost, whereas if it wants to reach them on top of the platform, that cost is eliminated. Similarly, we plan to explore the case in which instead of saving the entire shopping cost σ under hosting, consumers only save some fraction of σ, where the hosting firm can control the fraction. The natural question that arises is whether M would ever want to choose to reduce σ by some fraction rather than by nothing (non-hosting) or by 100% (full hosting).

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] Bourreau, M., J. Hombert, J. Pouyet, N. Schutz (2011) "Upstream competition between vertically integrated firms," *Journal of Industrial Economics*, 59, 677–713.

[3] Chen, Y. (2001) "On vertical mergers and their competitive effects," *RAND Journal of Economics*, 32, 667–685.

[4] Chen, Z. and P. Rey (2012) "Loss Leading as an Exploitative Practice," *American Economic Review*, 102(7), 3462–3482.

[5] Hagiu, A. and J. Wright (2015a) "Marketplace or Reseller?" *Management Science*, 61(1), 184–203.

[6] Hagiu and Wright (2015b) "Multi-sided Platforms," *International Journal of Industrial Organization*, 43, 162–174

[7] Hart, O. and J. Tirole (1990) "Vertical integration and market foreclosure," *Brookings Papers on Economic Activity*, 205–276.

[8] Johnson, J. (2017) "The Agency Model and MFN Clauses" *Review of Economic Studies*, 84(3), 1151–1185.

[9] O'Brien, D. P. and G. Shaffer (1992) "Vertical control with bilateral contracts," *RAND Journal of Economics*, 23, 299–308.

[10] Ordover, J. A., G. Saloner, and S.C. Salop (1990) "Equilibrium vertical foreclosure," *American Economic Review*, 80, 127–142.

[11] Ordover, J. A. and G. Shaffer (2007) "Wholesale access in multi-firm markets: when is it profitable to supply a competitor?" *International Journal of Industrial Organization* 25, 1026–1045.

[12] 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.

[13] Matutes, C. and P. Regibeau (1988) "Mix and match: product compatibility without network externalities." *RAND Journal of Economics*, 221–234.

[14] Sappington, D. (2005) "On the irrelevance of input prices for the make-or-buy-decision," *American Economic Review*, 95, 1631–1638.

[15] Whinston, M. (1990) "Tying, foreclosure, and exclusion," *American Economic Review*, 80(4), 837–859.

6 Appendix

6.1 Proof or 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 $v_B^* = u_A + u_B - p_A^* - p_B^* - \sigma = u_B + \Delta u_B - \sigma > 0$ 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 $u_B - p_B^* - \sigma = u_B + \Delta u_B - 2\sigma$, 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 $u_B + \Delta u_B - p_S - \sigma$. Given $u_A > \Delta u_B$, in the proposed equilibrium we must have $p_S = 0$, and $p_A + p_B = u_A - \Delta u_B$, giving M a profit of $\lambda_B (u_A - \Delta u_B)$. 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 $p_B = \sigma - \Delta u_B$ 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.

6.2 Proof of Proposition 6

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 $p_A^* = u_A + \alpha - \sigma$, $p_B^* = \sigma - \alpha - \Delta u_B$, $p_S^* = 0$. The A-type consumers always purchase A, and the B-type consumers all buy A and B_M from M. Profits are $\pi_M^* = u_A - \sigma + \lambda_A \alpha + \lambda_B (\sigma - \Delta u_B)$ 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 $\sigma - \Delta u_B > 0$ 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 $\lambda_A < \frac{\sigma-\alpha}{u_A}$, the selected equilibrium involves the prices $p_A^* = u_A$, $p_B^* = 0$, and $\widehat{p}_S^* = \Delta u_B$, the A-type consumers do not purchase, while the B-type consumers all buy A and B_S through M, and profits are $\pi_M^* = \lambda_B u_A - F$ and $\pi_S^* = \lambda_B \Delta u_B$; (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 $\widehat{p}_S^* = \Delta u_B$, the A-type consumers always purchase A, and the B-type consumers all buy A and B_S through M, and profits are $\pi_M^* = u_A + \alpha - \sigma - F$ and $\pi_S^* = \lambda_B \Delta u_B$.

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

6.3 Proof of Proposition 7

Consider first what happens without hosting. The equilibrium prices with non-hosting must satisfy $p_B^* \leq \sigma - \Delta u_B$ in order for B-type consumers to prefer buying B_M to B_S, and $p_A^* + p_B^* \leq u_A - \Delta u_B$ 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 $p_A^* = u_A + \alpha - \sigma$, $p_B^* = \sigma - \Delta u_B$, $p_S^* = 0$, with A-types still purchasing, and B-types buying the bundle from M, with profits being $\pi_M^* = u_A + \alpha - \lambda_A \sigma - \lambda_B \Delta u_B$ 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 $p_A^* + p_B^* = u_A - \Delta u_B$, $p_S^* = 0$, with B-types buying the bundle from M, with profits being $\pi_M^* = \lambda_B (u_A - \Delta u_B)$ 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 6, in which there are two cases:

- H-1 If $\lambda_A < \frac{\sigma - \alpha}{u_A}$, profits are $\pi_M^* = \lambda_B u_A - F$ and $\pi_S^* = \lambda_B \Delta u_B$.
- H-2 If $\lambda_A > \frac{\sigma - \alpha}{u_A}$, profits are $\pi_M^* = u_A + \alpha - \sigma - F$ and $\pi_S^* = \lambda_B \Delta u_B$.

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 $\lambda_B (u_A + \Delta u_B) - F$ under hosting with $\lambda_B (u_A - \Delta u_B)$ without hosting.

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

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

6.4 Proof of Proposition 8

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 $\pi_M = p_A + \lambda_B p_B \left(\frac{1}{2} + \frac{p_S - p_B + \sigma - \Delta u_B}{2t}\right)$ and $\pi_S = \lambda_B p_S \left(\frac{1}{2} - \frac{p_S - p_B + \sigma - \Delta u_B}{2t}\right)$. The equilibrium involves $p_A^* = u_A - \sigma$, $p_B^* = t + \frac{\sigma - \Delta u_B}{3}$, $p_S^* = t - \frac{\sigma - \Delta u_B}{3}$, $\pi_M^* = u_A - \sigma + 2t\lambda_B \left(\frac{1}{2} + \frac{\sigma - \Delta u_B}{6t}\right)^2$ and $\pi_S^* = 2t\lambda_B \left(\frac{1}{2} - \frac{\sigma - \Delta u_B}{6t}\right)^2$. It is straightforward to check that our assumptions on t imply S can cover its costs 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 $\lambda_A < \frac{\sigma}{u_A}$, M will set $p_A = u_A$ and sell only to B-types. In either case, the equilibrium involves $p_B^* = t - \frac{\Delta u_B}{3}$ and $\widehat{p}_S^* = t + \frac{\Delta u_B}{3}$. As a result, if M sets $p_A = u_A$, profit are $\pi_M = \lambda_B u_A + 2t\lambda_B \left(\frac{1}{2} - \frac{\Delta u_B}{6t}\right)^2 - F$ and $\pi_S^* = 2t\lambda_B \left(\frac{1}{2} + \frac{\Delta u_B}{6t}\right)^2$, while if M sets $p_A = u_A - \sigma$, profits are $\pi_M^* = u_A - \sigma + 2t\lambda_B \left(\frac{1}{2} - \frac{\Delta u_B}{6t}\right)^2 - F$ and $\pi_S^* = 2t\lambda_B \left(\frac{1}{2} + \frac{\Delta u_B}{6t}\right)^2$. 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 $p_S < \widehat{p}_S^*$ 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 $p_S < \widehat{p}_S^* - \sigma$. 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 $p_S < \widehat{p}_S^* - (u_A - p_A)$. In both cases, S could attract more additional consumers by lowering $\widehat{p}_S$ instead of p_S by the given amount. The fact it doesn't want to (i.e. that $\widehat{p}_S^*$ is the equilibrium level of p_S) implies it also cannot be better off lowering p_S below $\widehat{p}_S^*$.

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

6.5 Proof of Proposition 12

Since good B has no value without good A, without hosting B-types face two viable options (as well as the choice of not purchasing anything):

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

B-types will prefer to buy A and B_M over A and B_S, whenever $p_B \leq p_S + \sigma - \Delta u_B$. Since $\sigma > \Delta u_B$, in equilibrium we have $p_S^* = 0$ and $p_B^* = \sigma - \Delta u_B$. Note M now has two options:

1. set $p_A = u_A - \sigma$, sell to all consumers to obtain $\pi_M = u_A - \sigma + \lambda_B (\sigma - \Delta u_B)$.
2. set $p_A > u_A - \sigma$, so sell only to B-types. In this case, given that $p_B^* = \sigma - \Delta u_B$, M will sell A and B_M provided consumers' total surplus is non-negative, i.e. $p_A + p_B^* \leq u_A + u_B - \sigma$, or equivalently $p_A \leq u_A + u_B + \Delta u_B - 2\sigma$. So M sets $p_A = u_A - (2\sigma - (u_B + \Delta u_B))$, yielding profits of $\lambda_B (u_A + u_B - \sigma)$.

Thus, the equilibrium without hosting is defined as follows (in both cases, $p_B^* = \sigma - \Delta u_B$, $p_S^* = 0$ and $\pi_S^* = 0$):

NH-1 If $\lambda_A < \frac{u_B + \Delta u_B - \sigma}{u_B + \Delta u_B - \sigma + u_A - \sigma}$, M only sells to B-types and $\pi_M^* = \lambda_B (u_A + u_B - \sigma)$.

NH-2 If $\lambda_A > \frac{u_B + \Delta u_B - \sigma}{u_B + \Delta u_B - \sigma + u_A - \sigma}$, M sells to both types of consumers and $\pi_M^* = u_A - \sigma + \lambda_B (\sigma - \Delta u_B)$.

Now consider the case with hosting. Consider first an equilibrium in which M does not sell to A-types, so $p_A > u_A - \sigma$. B-types have two feasible options:

- buy A and B_M from M, obtaining net utility $u_A - p_A + u_B - p_B - \sigma$
- buy A from M and B_S through M, obtaining net utility $u_A - p_A + u_B + \Delta u_B - \widehat{p}_S - \sigma$

As in the benchmark case, we obtain the equilibrium $p_B^* = 0$ and $\widehat{p}_S^* = \Delta u_B$. Then p_A is set to extract the entire remaining surplus from B-types, i.e. $p_A = u_A + u_B - \sigma$, and $\pi_M^* = \lambda_B (u_A + u_B - \sigma) - F$ and $\pi_S^* = \lambda_B \Delta u_B$.

On the other hand, if M sets p_A to sell to A-types, then we get the standard equilibrium prices $p_A^* = u_A - \sigma$, $p_B^* = 0$, and $\widehat{p}_S^* = \Delta u_B$. 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 $\pi_S^* = \lambda_B \Delta u_B$.

Thus, the equilibrium with hosting is defined as follows (in both cases $p_B^* = 0$, $\widehat{p}_S^* = \Delta u_B$ and $\pi_S^* = \lambda_B \Delta u_B$):

H-1 If $\lambda_A < \frac{u_B}{u_B + u_A - \sigma}$, M only sells to B-types and $\pi_M^* = \lambda_B (u_A + u_B - \sigma) - F$.

H-2 If $\lambda_A > \frac{u_B}{u_B + u_A - \sigma}$, M sells to both types of consumers and $\pi_M^* = u_A - \sigma - F$.

Note that $\frac{u_B}{u_B+u_A-\sigma} \geq \frac{u_B+\Delta u_B-\sigma}{u_A-\sigma+u_B+\Delta u_B-\sigma}$ since $\sigma \geq \Delta u_B$, so if NH-1 holds then H-1 holds, and if H-2 holds then NH-2 holds. Comparing joint profits under hosting with non-hosting, we can therefore conclude:

- If $\lambda_A < \frac{u_B+\Delta u_B-\sigma}{u_B+\Delta u_B-\sigma+u_A-\sigma}$, hosting is jointly preferred iff $\Delta u_B > \frac{F}{1-\lambda_A}$.
- If $\frac{u_B+\Delta u_B-\sigma}{u_B+\Delta u_B-\sigma+u_A-\sigma} < \lambda_A < \frac{u_B}{u_B+u_A-\sigma}$, hosting is jointly preferred iff $\Delta u_B > \frac{\lambda_A(u_A-\sigma)-\lambda_B(u_B-\sigma)+F}{2(1-\lambda_A)}$.
- If $\lambda_A > \frac{u_B}{u_B+u_A-\sigma}$, hosting is jointly preferred iff $\Delta u_B > \frac{\sigma}{2} + \frac{F}{2(1-\lambda_A)}$.

6.6 Proof of Proposition 13

Consider the case with hosting. Consider solving for the equilibrium in the second stage first, for a given τ. We can assume $\tau \leq \sigma$, otherwise S would always prefer to sell directly and M would never get paid τ. So it doesn't make sense to set $\tau > \sigma$.

Consider first an equilibrium in which M does not sell to A-types, so $u_A - \sigma < p_A \leq u_A$. Note this also implies B-types will never buy just A from M. B-types have three options:

- buy A and B_M from M, obtaining net utility $u_A - p_A + u_B - p_B - \sigma$
- buy A from M and B_S through M, obtaining net utility $u_A - p_A + u_B + \Delta u_B - \widehat{p}_S - \sigma$
- buy B_S from S directly, obtaining net utility $u_B + \Delta u_B - p_S - \sigma$

Given that S can offer a higher surplus than M under option 2 compared to option 1, in equilibrium it must be that S wins sales of B. And since M must sell A in equilibrium, the only equilibrium possibility is that B-types buy A from M and B_S through M. This means we must have $\widehat{p}_S \leq p_B + \Delta u_B$ and $\widehat{p}_S \leq p_S + u_A - p_A$. If the first inequality is strict, then S could profitably increase $\widehat{p}_S$ and p_S by the same amount. If the second inequality is strict, then M could profitably increase p_A. Thus, we must have

$$\widehat{p}_S = p_B + \Delta u_B \quad (1)$$
$$\widehat{p}_S = p_S + u_A - p_A. \quad (2)$$

In this equilibrium, S makes profits $\lambda_B (\widehat{p}_S - \tau)$ and must not have an incentive to slightly decrease p_S and get all B-types to just buy from it directly. This means we must also have $\widehat{p}_S - \tau \geq p_S$, which combined with the second equality above implies

$$p_A \leq u_A - \tau.$$

We must also have $p_B \leq \tau$ in this equilibrium, otherwise M could profitably take over sales of B by slightly decreasing p_B. And of course, $p_B \geq \tau - \Delta u_B$, so that S makes a non-negative margin in equilibrium (it cannot be forced to sell at a loss). Finally, M must not want to decrease p_A all the way to $u_A - \sigma$ and also serve A-types. This implies

$$\lambda_B p_A \geq u_A - \sigma.$$

Taking these conditions together, M's equilibrium prices must satisfy

$$\begin{aligned} \tau - \Delta u_B &\leq p_B \leq \tau \\ \frac{u_A - \sigma}{1 - \lambda_A} &\leq p_A \leq u_A - \tau. \end{aligned}$$

Thus, this equilibrium exists if and only if

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

If this condition is satisfied, then any (p_A, p_B) satisfying the inequalities above can sustain this equilibrium, where $\widehat{p}_S$ and p_S are completely determined by (1)-(2). We can then choose the equilibrium with the highest combined profits, which is also the one with the highest profits for M, i.e.

$$\begin{aligned} p_A &= u_A - \tau \\ p_B &= \tau \\ \widehat{p}_S &= \Delta u_B + \tau \\ p_S &= \Delta u_B \end{aligned}$$

Profits are $\lambda_B u_A$ for M and $\lambda_B \Delta u_B$ for S.

Now consider the other possible equilibrium, in which M does sells to A-types, so $p_A \leq u_A - \sigma$. This means M sells A to all consumers and in equilibrium we must have $p_A = u_A - \sigma$. Type B consumers have three options:

- buy A and B_M from M, obtaining net utility $u_B - p_B$
- buy A from M and B_S through M, obtaining net utility $u_B + \Delta u_B - \widehat{p}_S$
- buy A from M and B_S from S directly, obtaining net utility $u_B + \Delta u_B - p_S - \sigma$

In equilibrium, M cannot win sales of B since that would require $\widehat{p}_S = \tau$, $p_S = 0$ and $p_B = \tau - \Delta u_B$, implying M's equilibrium profits would be $\lambda_B (u_A - \sigma + \tau - \Delta u_B)$. However, M could then profitably deviate by setting some very large p_B and keeping $p_A = u_A - \sigma$, which leads B-types to buy A from M and B_S through M (because $\tau \leq \sigma$) resulting in profits $\lambda_B (u_A - \sigma + \tau)$ for M. Furthermore, $\tau \leq \sigma$ implies that in equilibrium it can't be that B-types buy A from M and B_S from S directly (option 3 above). If this were the case, then we would have to have $p_S + \sigma \leq \widehat{p}_S$ and $p_S = p_B + \Delta u_B - \sigma$, so S's profits would be $\lambda_B (p_B + \Delta u_B - \sigma)$. But then S could profitably deviate to some very high p_S and

$\widehat{p}_S = p_B + \Delta u_B$, yielding profits $\lambda_B (p_B + \Delta u_B - \tau)$. Thus, in equilibrium it must be that B-types buy A from M and B_S through M. This implies we must have

$$\widehat{p}_S = p_B + \Delta u_B \leq p_S + \sigma.$$

We must also have $p_B \leq \tau$ in this equilibrium, otherwise M could profitably take over sales of B by slightly decreasing p_B. And we must also have $p_B \geq \tau - \Delta u_B$, so that S makes a non-negative margin in equilibrium (it cannot be forced to sell at a loss). Finally, M must not want to increase p_A above $u_A - \sigma$ and only serve B-types. The best such deviation for M is to set p_A such that

$$u_A - p_A + u_B + \Delta u_B - \widehat{p}_S - \sigma = u_B + \Delta u_B - p_S - \sigma,$$

provided the solution in p_A is below u_A. So the best deviation is $p_A = u_A - \max\{0, \widehat{p}_S - p_S\}$. The profits in this deviation are $\lambda_B (u_A + \tau - \max\{0, \widehat{p}_S - p_S\})$ and we want this to be lower than M's profits in the proposed equilibrium, $u_A - \sigma + \lambda_B \tau$. Thus, we need

$$\begin{aligned} \lambda_B u_A - \lambda_B \max\{0, \widehat{p}_S - p_S\} &\leq u_A - \sigma \\ \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} &\leq \max\{0, \widehat{p}_S - p_S\}. \end{aligned}$$

Taking these conditions together, this equilibrium must have

$$\begin{aligned} p_A &= u_A - \sigma \\ \widehat{p}_S &= p_B + \Delta u_B \leq p_S + \sigma \\ \tau - \Delta u_B &\leq p_B \leq \tau \\ \frac{\sigma - \lambda_A u_A}{1 - \lambda_A} &\leq \max\{0, \widehat{p}_S - p_S\}. \end{aligned}$$

Since $\frac{\sigma - \lambda_A u_A}{1 - \lambda_A} \leq \sigma$, this equilibrium always exists. As in the previous case, we can focus on the equilibrium with the highest joint profits (M's profits are uniquely determined and equal to $u_A - \sigma + \lambda_B \tau$), which is

$$\begin{aligned} p_A &= u_A - \sigma \\ p_B &= \tau \\ \widehat{p}_S &= \Delta u_B + \tau \\ p_S &= \Delta u_B + \tau - \sigma. \end{aligned}$$

Profits are $u_A - \sigma + \lambda_B \tau$ for M and $\lambda_B \Delta u_B$ for S.

Thus, if $\sigma \leq \lambda_A u_A$, then only the second type of equilibrium exists for any $\tau \in [0, \sigma]$. In this case, M maximizes its profits by setting $\tau = \sigma$, leading to profits $u_A - \lambda_A \sigma$ for M and $\lambda_B \Delta u_B$ for S, with joint profits equal to $u_A - \lambda_A \sigma + \lambda_B \Delta u_B$.

Suppose now $\sigma > \lambda_A u_A$. In this case, for $\tau \in \left[0, \frac{\sigma - \lambda_A u_A}{1-\lambda_A}\right]$, there are two possible equilibria: one in which profits are $\lambda_B u_A$ for M and $\lambda_B \Delta u_B$ for S, and the other in which profits are $u_A - \sigma + \lambda_B \tau$ for M and $\lambda_B \Delta u_B$ for S. Since $\tau \leq \frac{\sigma - \lambda_A u_A}{1-\lambda_A}$, the former is the more profitable equilibrium both for M and from a joint profit perspective, so we can choose to focus on it. For $\tau \in \left[\frac{\sigma - \lambda_A u_A}{1-\lambda_A}, \sigma\right]$, there is once again a unique equilibrium, in which profits are $u_A - \sigma + \lambda_B \tau$ for M and $\lambda_B \Delta u_B$ for S. Thus, with this equilibrium selection M maximizes its profits by once again choosing $\tau = \sigma$, which leads to the profits $u_A - \lambda_A \sigma$ for M and $\lambda_B \Delta u_B$ for S, with joint profits equal to $u_A - \lambda_A \sigma + \lambda_B \Delta u_B$.

In summary, M will always choose $\tau = \sigma$, yielding profits $u_A - \lambda_A \sigma - F$ for M and $\lambda_B \Delta u_B$ for S, with joint profits equal to $u_A - \lambda_A \sigma + \lambda_B \Delta u_B - F$. Comparing this with the joint profit without hosting (recall that was $u_A - \lambda_A \sigma - \lambda_B \Delta u_B$), we obtain the result in the text of the proposition.

6.7 Proof of Proposition 14

Consider what happens to the proof of Proposition 13 above if we impose the constraint that $p_S = \Delta u_B$ and $\widehat{p}_S \leq \Delta u_B$. The fact that $p_S = \Delta u_B$ implies the equilibrium without hosting will involve $p_A = u_A - \sigma$ and $p_A + p_B = u_A$ since M only needs to leave the surplus $u_B - \sigma$ for B-types in order for them not to buy from S directly. This implies $p_B = \sigma$, so that without hosting the equilibrium profits are $\pi_M^* = u_A - \sigma + \lambda_B \sigma$ and $\pi_S^* = 0$. With hosting, the constraint $\widehat{p}_S \leq \Delta u_B$ implies that any $\tau > 0$ will not increase $\widehat{p}_S$, but simply transfer profits from S to M. Thus, from the perspective of maximizing joint profit, the equilibrium under hosting leads to the same joint profit as the hosting equilibrium with $\tau = 0$ in Proposition 2. If $\lambda_A < \frac{\sigma}{u_A}$, joint profits are $\lambda_B (u_A + \Delta u_B) - F$, while if $\lambda_A > \frac{\sigma}{u_A}$, joint profits are $u_A - \sigma + \lambda_B \Delta u_B - F$. Thus, if $\lambda_A < \frac{\sigma}{u_A}$ hosting is jointly profitable if $\Delta u_B > \frac{\lambda_A (u_A - \sigma) + F}{1 - \lambda_A}$, while if $\lambda_A > \frac{\sigma}{u_A}$ non-hosting is always jointly profitable.