Creating platforms by hosting rivals *

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

March 2019

Abstract

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

JEL classification: D4, L1, L5

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

1 Introduction

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

*We thank Patrick Rey and Alex White for very helpful comments, as well as participants at 44th EARIE in Maastricht, the Online Platform Competition Conference at the University of Florida Gainesville, 16th ZEW Conference on ICT in Mannheim, and the 2018 platform research symposium at Boston University. Gunhaeng Lee and Tat-How Teh provided excellent research assistance. The financial support of the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 670494) and Singapore Ministry of Education Social Science Research Thematic Grant (MOE2017-SSRTG-023) is gratefully acknowledged.

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

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

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

However, in many cases, existing product (or service) companies have the potential to become (multi-sided) platforms too. The most straightforward way for a product company to do so is by inviting third parties to sell their products or services on top of the original product the company already sells. Two well-known and successful examples are Apple's iPhone and Salesforce's customer relationship management (CRM) software. After launching the iPhone in June 2007 as a stand-alone product running exclusively Apple produced apps, Apple quickly realized it would benefit from the creativity of third-party developers. As a result, in July 2008 the company opened up the iPhone to third-parties (including those producing rival apps) and created the App Store. Salesforce was founded in 1999 as a seller of CRM software products to small-to-medium-size businesses. In 2005, the company created a platform (Force.com) and an app marketplace (AppExchange) around its offering, which allows third-party software developers to build and sell other software to Salesforce's CRM customers. Today there are over one million Force.com registered developers and over 2,500 apps offered on AppExchange.

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

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

Similarly, consider Intuit, the seller of QuickBooks, which is the leading software for accounting, financial management and tax compliance for small businesses in the United States. Over the past seven 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. Today, the platform offers QuickBooks customers around 1,400 apps.[1] Some of these products compete with features already included in QuickBooks: for instance, third-party payroll management apps such as TimeTracker and TimeRewards are direct substitutes for Intuit's own Tsheets. Intuit has also created QuickBooks Capital, a marketplace where both Intuit and a number

[1]See "How Intuit Reinvents Itself" in Fortune, October 2017, http://fortune.com/2017/10/20/how-intuit-reinvents-itself/

of selected third-party lenders can offer loans to QuickBooks customers. Intuit negotiates attractive rates with the third-party lenders and makes it easier for QuickBooks customers to apply for loans directly from QuickBooks by directly providing the relevant information to lenders.[2]

In many countries Cable and satellite TV providers have allowed Netflix to sell to their subscribers through their own platforms, even though Netflix competes with the cable companies' video-on-demand services.[3] In the financial sector, a European company named "Deposit Solutions" is facilitating the movement of banking to a platform model. Deposit Solutions provides a software infrastructure that allows banks to offer third party deposit products to their own customers through their existing accounts. As an example, Deutsche Bank offers its German account holders the chance to access fixed deposits of selected rival banks through its "ZinsMarkt".

Gyms provide a more "physical" example. Recently, some "big box gyms" have begun renting out space in their facilities to specialty studios, where the latter can offer classes to the gym's members. For instance, the New York Sports Club (NYSC) hosts cycling classes offered by Cyc Fitness, a boutique cycling studio, within several of the NYSC's gym locations in New York City.[4] Country clubs work in a similar way, sometimes hosting third-parties that provide classes or specialized services to the clubs' members (e.g. swimming or tennis coaching), where these were previously (or sometimes still are) provided by the clubs.

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

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

By eliminating the additional shopping cost consumers incur when they wish to buy A from M and B from S, hosting eliminates M's ability to price discriminate across the two types of consumers. Essentially, when S is hosted, M unbundles the A and B products, so hosting turns competition for the market into competition *within* the market. This means M can no longer make a profit by selling B, given that S offers a superior version of B and the firms now compete on a level playing field in B. On the other hand, hosting allows M to potentially gain by raising its price on product A because

[2]See https://quickbooks.intuit.com/capital/

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

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

shopping costs are now taken care of by the surplus offered by S's superior version of B. In this sense, hosting allows M to gain by turning a substitute into a complement. This logic extends to the joint profit perspective, which becomes relevant when the two firms can make lump-sum transfers to each other under hosting. Taking into account both firms' profits expands the region of parameters where hosting dominates, but the tradeoff and its underlying logic remain.

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

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

Finally, we discuss the factors that determine the optimal number of specialist to host when there are multiple competing specialists.

2 Related literature

Our paper relates to several strands of literature.

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

Somewhat closer is the literature discussing a firm's decision to open itself up to third-party developers. The key issues that this literature has focused on are how much technology to share with platform participants (Boudreau, 2010, Parker and Van Alstyne, 2018, Niculescu et al., 2018), whether to make a piece of software open source (August et. al, 2013 and 2018), and whether a firm should become a “platform” and collaborate with producers of complementary products (Mantovani and Ruiz-Aliseda, 2016). In these papers, the main benefit of opening up to third-parties is to encourage innovation and contributions by outside developers or complementors, who are not rivals of the platform. The current paper is different in that it shows “opening up” can make sense even if the third-parties are rivals.

While our focus is on the question whether a firm should host a rival, a related but different question is whether, conditional on hosting, the platform would want to offer (and possibly bundle) its own inferior version of the hosted product, even though no one will buy it in equilibrium. This is the focus of Carlton et al. (2010), who show that bundling the inferior version of the complementary product reduces consumer's willingness to pay for that product and thereby allows the platform to extract higher profits from the sale of its primary product. Another key difference is that in Carlton et al. (2010) the two products are strict complements, whereas in our model, they are independent and only become complements due to shopping costs.

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

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

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

At a high level, the profitability of helping a rival by hosting them seems related to the literatures on raising rivals' costs (Salop and Scheffman, 1983) and second sourcing (Farrell and Gallini, 1988). However, hosting in our paper is closer to *reducing* a rival's costs rather than raising them. And the mechanism which makes hosting profitable in our model does not rely on a commitment to have more competition, which is the main driver behind second sourcing.

Our paper is also related to the literature on compatibility between system components sold by different firms. In particular, Matutes and Regibeau (1988) have shown that compatibility between components from various firms may relax competition between systems. In our model, hosting also makes it easier for consumers to buy from multiple firms, but the products are independent. Hosting then intensifies competition on one good while relaxing the competitive pressure on the other one.

Finally, our paper belongs to a burgeoning economics literature which explores the implications of competition with multiproduct firms when, due to shopping costs, consumers have a demand for one-stop shopping. In this literature, like our paper, different products sold at the same firm become de facto complementary to one another due to consumers wanting to save on shopping costs. Important early contributions to this literature are Lal and Matutes (1989, 1994), and Chen and Rey (2012). Lal and Matutes (1989) study the pricing equilibrium in a game with two competing stores, each selling two products, and in which a fraction of consumers have an inherent preference to buy both products from the same store (in order to save on shopping costs). Lal and Matutes (1994) builds on the same model by allowing consumers to be imperfectly informed about prices (stores can advertise some of their prices but not others). It shows that imperfect information about prices can lead to loss-leader pricing strategies. Chen and Rey (2012) takes this analysis one step further, showing that even under perfect information about prices, loss-leader pricing can be a profitable strategy for a large store competing with a smaller one (i.e. a store offering a narrower range of products). None of these papers allows for a firm (store) to "host" the competitor's product(s), which is the focus of our paper. More recent works such as those by Zhou (2014) and Rhodes and Zhou (2018) have focused on the implications of one-stop shopping for pricing, bundling and product range decisions in search contexts. We are the first to focus on the implications of one-stop shopping for hosting a rival's product.

3 Model

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

The total measure of consumers is normalized to one. Among them, there are two types. A fraction $\lambda_A > 0$ of consumers just want to purchase one unit of A and are not interested in B (i.e. they value both versions of B at zero). In the gym example, these are consumers who just want to use the gym's core facilities and are not interested in specialized classes. We call these monoproduct consumers "A-type" consumers. The remaining fraction $\lambda_B > 0$ (which equals $1 - \lambda_A$) of consumers obtain value from both A and B, and so are interested in purchasing one unit of each. We call these multiproduct consumers "B-type" consumers. All consumers value product A at $u_A > 0$.[5] The B-types value B_M at $u_B > 0$ and B_S at $u_S = u_B + \Delta$, where $\Delta \geq 0$. Thus, B-types view B_S as superior to B_M.[6] Note that A and B are not complementary products—B-types can consume one or both of them without changing the utilities derived from each.

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

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

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

$$\Delta \leq \sigma,$$

i.e. the shopping cost exceeds the added value of S's product B_S. This implies that without hosting, M can have an advantage in selling product B, provided it makes it attractive for consumers to want to buy A. In Section 4.4, we briefly discuss what happens in the less interesting case in which $\Delta > \sigma$.

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

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

4 Benchmark results

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

4.1 Without hosting

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

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

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

Because the shopping cost σ outweighs S's added value in product B (i.e. Δ), there is a unique equilibrium outcome in which M makes all the sales.[7] Formal proofs for this result and others are provided in the appendix.

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

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

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

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

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

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

4.2 Hosting

Now suppose S is hosted by M, meaning B-types can buy B_S from S through M without incurring the additional shopping cost σ. In the gym example, this means that the NYSC now hosts cycling classes offered by Cyc Fitness, so a NYSC member interested in Cyc's classes does not have to go to a separate Cyc studio. We still allow S to sell directly, at price p_S (Cyc Fitness did not stop offering classes in its studio after being hosted by the NYSC).[8] 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 $F \geq 0$. For instance, hosting a specialized cycling class in a gym may require re-arranging and customizing the space with the relevant equipment and branding, as well as updating software systems for scheduling and reservations to include the specialized class. Similarly, for a bank there could be significant system costs (software, compliance, training) of allowing rival providers to sell their term deposits to its customers. Any negotiating and legal costs associated with writing a hosting contract would also be included in F, as should anticipated costs of integrating systems and employees more generally. In practice, both M and S may incur such costs. Since throughout most of the paper we will focus on the solution in which a lump-sum transfer can be made between the two firms (i.e. through a fixed fee), it will make no difference which firm actually incurs the fixed costs of hosting. Thus, for convenience, we will assume F is always incurred by M.

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

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

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

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

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.[9]

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

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

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

Under hosting, if M could engage in third-degree price discrimination, it would want to charge $u_A - \sigma$ to A-types and u_A to B-types (indeed, the B-types' shopping costs are now covered by the surplus offered by B_S). However, given that S now competes and wins the market for B on the platform created by M, such price discrimination is no longer possible. This drives a tradeoff between hosting and non-hosting, which we will explore in the next section.[10]

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

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

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

monopoly over product B), clearly M can do better by offering B_M.[11] Thus, the only way M would not offer B_M is if it could commit *ex-ante* to not offer it. This would always increase joint profits in our baseline model. However, this may not be possible in practice if it requires M to write a contract specifying that it will not compete with S on M's platform. Indeed, this type of contract would likely raise antitrust concerns because it could be viewed as a form of collusion. A commitment to remove B_M may therefore require a technological commitment, which may not always be feasible. Even if such a commitment is feasible, it may not be jointly profitable once firms take into account the realistic possibility that other firms would then want to enter to sell B. This is indeed the situation when there are multiple specialists competing, which we consider in Section 5.2.

4.3 Unilateral incentive to host

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

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

It is easiest to interpret this result when there are almost no A-types, i.e. when $\lambda_A \to 0$. Then, hosting allows M to increase its profit by Δ, so hosting is preferred provided Δ exceeds the fixed cost F of hosting. The gain of Δ that M obtains from hosting comes from a gain of σ on the A product and a loss of $\sigma - \Delta$ on the B product (which is smaller). Indeed, hosting allows M to charge u_A for A instead of $u_A - \sigma$, because shopping costs are now taken care of by S through the surplus obtained from B_S. This is the sense in which hosting S (and thereby eliminating the shopping cost necessary to access B_S) allows M to gain by turning a substitute into a complement. This captures the fundamental benefit of hosting a rival of superior quality for M: it increases the value derived by consumers from visiting M. In our simple model, M can only capture this value by raising its price p_A (because demand of each type of consumer is inelastic), but in reality, this could simply result in higher demand even if M does not raise its price.

On the other hand, under hosting M no longer extracts $\sigma - \Delta$ from its sale of B_M as sales of B are now made by S. Thus, turning competition for the market into competition within the market means that M gives up on its profit in the B market. Put differently, hosting unbundles the products and levels the playing field in product B competition, which means M can no longer make a profit on B. In contrast, S can now extract the profit $\lambda_B \Delta$, selling B_S to B-types, which is why it strictly prefers hosting.

Now consider what happens when there are some (but not too many) A-types, so $\lambda_A \leq \frac{\sigma}{u_A}$. In this case, if M charges u_A for A instead of $u_A - \sigma$, it loses the A-types, who no longer purchase. This means the additional surplus extracted by M from product A under hosting may no longer dominate

[11] This is reminiscent of the results in Carlton et al. (2010).

the negative effect of hosting on M's profit in the B market. Put differently, the presence of A-types constrains M's ability to extract more from product A under hosting since A-types do not care about the extra surplus generated from hosting S's superior version of product B. Thus, M can be worse off under hosting even without taking into account the fixed costs of hosting. This happens when the loss of B-type sales under hosting (which recall is equal to $\sigma - \Delta$) is large, i.e. when Δ is small.

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

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

The logic of our results does not depend on the stark prediction that hosting is unilaterally profitable for M only if A-types stop purchasing. To show this, in Section C of the Online Appendix, we provide an extension of our baseline model to allow for A-types to have elastic demand (i.e. different willingness to pay for product A). We show that by increasing the willingness to pay of B-types, hosting allows M to raise its price for product A and therefore can be unilaterally profitable even though only some A-types stop purchasing.

4.4 Joint incentives to host without monitoring

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

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

Comparing Proposition 4 with Proposition 3, it is clear that the tradeoff between hosting and non-hosting in terms of joint profits is similar to the one when we focused only on M's profit. The only difference is that now under hosting we must add S's profit $\lambda_B \Delta$, which expands the region of parameter values for which hosting dominates. Specifically, if $\Delta > \frac{\sigma}{2}$, then hosting raises joint profits, so in this case it always dominates non-hosting for F sufficiently small.[12] This extends to the case

[12]The impact of hosting on consumer surplus and welfare is explored in Section D of the Online Appendix.

we ruled out by assumption, namely $\Delta > \sigma$, i.e. S's efficiency advantage is so high that it more than offsets the shopping cost advantage of M. In that case, aside from the fixed cost F, there is no downside to hosting given that M does not sell B either way (with or without hosting). In particular, it is easily shown that, when $\Delta > \sigma$, hosting is jointly preferred if $\lambda_B\sigma + \max\{\sigma - \lambda_A u_A, 0\} > F$.

The tradeoff between hosting and non-hosting in Proposition 4 shifts towards hosting when Δ increases and towards non-hosting when λ_A, F or u_A increase. When σ increases, the tradeoff shifts towards hosting if $\lambda_A \leq \frac{\sigma}{u_A}$ and towards non-hosting if $\lambda_A > \frac{\sigma}{u_A}$. To understand why the tradeoff is non-monotonic in the level of the shopping cost σ, note that without hosting, because M is able to price discriminate, the amount it collects is lowered by the shopping cost for the A-types only (i.e. by $\lambda_A\sigma$). On the other hand, under hosting, if M wants to sell to A-types, it has to lower its price by σ on all consumers. This means when there are relatively many A-types, so they are served by M under hosting, a higher σ shifts the tradeoff in favor of non-hosting. On the other hand, when there are relatively few A-types, so M gives up selling to them under hosting, a higher σ shifts the tradeoff in favor of hosting.

4.5 Joint incentives to host with monitoring

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

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

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

- *If $\lambda_A \leq \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, then equilibrium prices $p_A^* = u_A$, $p_B^* = \min\{\sigma, u_B - \sigma\}$ and $\widehat{p}_S^* = \min\{\sigma, u_B-\sigma\} + \Delta$ yield the highest joint profit that can be achieved in the hosting equilibrium, equal to $\lambda_B(u_A + \Delta + \min\{\sigma, u_B - \sigma\}) - F$. Hosting is jointly preferred if and only if $\Delta > \frac{F-(\sigma-\lambda_A u_A+(1-\lambda_A)\min\{0,u_B-2\sigma\})}{2(1-\lambda_A)}$.*

- *If $\lambda_A > \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, then equilibrium prices $p_A^* = u_A - \sigma$, $p_B^* = \sigma$ and $\widehat{p}_S^* = \sigma + \Delta$ yield the highest joint profit that can be achieved in the hosting equilibrium, equal to $u_A - \sigma + \lambda_B(\sigma + \Delta) - F$. Hosting is jointly preferred if and only if $\Delta > \frac{F}{2(1-\lambda_A)}$.*

Comparing Proposition 5 with Proposition 4, it can be shown that the possibility of using a variable fee unambiguously shifts the tradeoff towards hosting. Hosting S and charging it a variable fee $\tau = \sigma$ allows M to preserve the same competitive edge as without hosting while relaxing the competitive pressure on good A.[13] Thus, joint profits under hosting increase as a result of using variable fees. It turns out the use of variable fees increases joint profits under hosting such that we now need $F > 0$ for the tradeoff to be non-trivial: if $F = 0$, then hosting is always jointly preferred when M can charge a variable fee. The reason is that the variable fee τ is another instrument that M can use to price discriminate between A-types and B-types whenever M chooses to keep selling to A-types: τ allows M to extract the increase in surplus offered to B-types by S's superior product. Thus, τ makes up for the loss of the ability to price discriminate using the price p_B for B_M. Meanwhile, M can use p_B to control any double marginalization problem that would otherwise arise with variable fees.

As in the case without variable fees, the tradeoff between hosting and non-hosting in Proposition 5 shifts towards hosting when Δ increases and towards non-hosting when λ_A, F increase. The effect of an increase in u_A is to shift the tradeoff towards non-hosting if $\lambda_A \leq \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, but otherwise it has no effect on the tradeoff. This is the same as the case without variable fees, except the cutoff value of λ_A for which an increase in u_A shifts the tradeoff towards non-hosting is higher (i.e. when $\lambda_A \leq \frac{\sigma}{u_A}$). The only qualitative differences compared to the case without variable fees are the tradeoffs with respect to increases in σ and u_B. Recall without variable fees an increase in σ shifted the tradeoff towards hosting if $\lambda_A \leq \frac{\sigma}{u_A}$ and towards non-hosting if $\lambda_A > \frac{\sigma}{u_A}$. With variable fees, in the parameter range $\lambda_A < \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, an increase in σ similarly shifts the tradeoff towards hosting except if $u_B < 2\sigma$ and $\lambda_A > \frac{1}{2}$ in which case an increase in σ shifts the tradeoff towards non-hosting. On the other hand, in the parameter range $\lambda_A > \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, the tradeoff no longer changes when σ increases. The reason is that now, under hosting, M can use the variable fee to extract an additional margin of σ from B-types, which is the same margin S could extract under non-hosting. Finally, without variable fees, u_B had no impact on the tradeoff because it was always competed away by Betrand competition between S and M on the platform. With variable fees, if $u_B < 2\sigma$ and $\lambda_A \leq \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, then the tradeoff shifts towards hosting if u_B increases (otherwise, u_B has no effect on the tradeoff). The reason is that in this case, under hosting, u_B is small enough that S and M together extract the entire surplus from B-type consumers.

5 Extensions

In this section we consider two important extensions of the baseline model. These address two limitations of the benchmark model. The first is that hosting always dominates once firms are allowed sufficient contracting instruments, absent any fixed cost of hosting. The second is that we only allowed for one specialist, but in reality firms may often host multiple specialists.

To address the first limitation, we allow some consumers to have imperfect information about the specialist's existence, and show that hosting specialists may not always dominate, even if variable fees can be used and even if there are no fixed costs associated with hosting. To address the second

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

limitation, we discuss the reasons why a firm may want to host multiple specialists.

5.1 Hosting as information

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

In the case without hosting, given some of M's customers don't know about S's offering, they will be willing to pay a higher price for M's second-rate product than they would if they knew about the substitute specialist. This allows M to set a higher price for its version of B in the case without hosting, but at the expense of losing informed consumers to S. Thus, if the fraction η of uninformed consumers is small, M does not find it profitable to target uninformed consumers exclusively by setting a high price, thereby giving up selling B_M to informed consumers. Therefore the equilibrium is the same as with fully informed consumers in this case. On the other hand, if the fraction of uninformed consumers is sufficiently high, starting from this pure-strategy equilibrium, M would prefer to increase p_B all the way to u_B in order to extract the entire surplus of uninformed consumers, thus giving up on selling to the informed B-types. However, this cannot be an equilibrium since S would best respond by setting a higher price, which would then entice M to once again try selling to informed B-types. Instead, there is a unique mixed strategy equilibrium where both M and S randomize over the price of their respective versions of B, while M still sells A to all consumers by setting $p_A = u_A - \sigma$. In the parameter range on which this mixed strategy equilibrium prevails, both firms obtain higher expected profits than in the benchmark case, reflecting that the presence of uninformed consumers relaxes price competition. In contrast, by promoting the specialist, hosting removes the friction that prevented S from reaching all consumers, thereby intensifying competition.

Taking into account this effect, and that the profits under hosting are unchanged compared to the full information case (given in the hosting equilibrium all B-type consumers shop at M and so become informed of S), we can then compare joint profits under hosting with non-hosting. The full details (and proofs) are in Section E of the Online Appendix, which includes the results for the case in which the firms cannot use a variable fee and the case in which they can. Here we focus on the result for the most interesting case in which the firms can use a variable fee, given this was the case in which previously hosting always dominated in the absence of any fixed cost of hosting.

Proposition 6 *Suppose a fraction η of B-type consumers are uninformed of S's existence under non-hosting, but they become informed about S under hosting if they visit M. When variable fees can be used:*

- *If $\eta \leq \frac{\sigma-\Delta}{u_B}$, the conditions for hosting to be jointly preferred are identical to those in Proposition 5.*
- *If $\eta > \frac{\sigma-\Delta}{u_B}$ and $\lambda_A \leq \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, hosting is jointly preferred if and only if*

$$\Delta > \frac{F-(\sigma-\lambda_A u_A)}{\eta(1-\lambda_A)} + \frac{(2-\eta)(\eta u_B-\sigma)-\min\{0,u_B-2\sigma\}}{\eta}.$$

- *If $\eta > \frac{\sigma-\Delta}{u_B}$ and $\lambda_A > \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, hosting is jointly preferred if and only if*

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

The first case in Proposition 6 is identical to the benchmark analysis (see Proposition 5) since the equilibrium analysis with hosting is unchanged and the equilibrium analysis without hosting is also unchanged when there are not many uninformed consumers, as explained above. In the second and third cases in Proposition 6, we show in Section E of the Online Appendix that the tradeoff unambiguously shifts in favor of non-hosting relative to the benchmark in Proposition 5. In particular, in this case non-hosting can be jointly preferred even though M can use variable fees and even if there is no fixed cost associated with hosting (i.e. if $F=0$), provided η is high enough.

5.2 Multiple specialists

Thus far, we have focused on the case of a single specialist S. Obviously, in reality there may be several competing specialists and M has a choice of how many of them to host. For instance, while Salesforce hosts several third-party apps for consumer surveys (e.g. Survey Monkey, GetFeedback, QuestionPro) and data management (e.g. CongaGrid, GridBuddy) that compete with functionality included in Salesforce's CRM product, there are other commonly used survey apps (e.g. SurveyGizmo, KeySurvey, Praiseworthy) and data management apps (e.g. SAS) that are not hosted on AppExchange. Nor does Salesforce's platform host any competitor to its core CRM functionality, such as Pipeliner CRM.[14] Similarly, on its QuickBooks platform, Intuit hosts some third-party payroll management apps (e.g. TimeTracker and TimeRewards) that compete with Intuit's own Tsheets, but other prominent alternatives to Tsheets (e.g. Hubstaff, Toggl, Freckle) are not hosted.[15]

Suppose there are multiple, identical competing specialists in our benchmark setting and the fixed cost of hosting is independent of the number of hosted specialists. When no transfers are feasible

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

[15] See https://blog.hubstaff.com/tsheets-alternatives/

between the firms, M is weakly better off hosting two or more specialists rather than one. This is because the hosted specialists compete away their value added Δ, which means M can extract all the surplus from B-types created by hosting if it chooses not to serve A-types. If M chooses to serve A-types, then hosting one or more specialists yields the same profits since M cannot charge more than $u_A - \sigma$ for product A. When lump-sum transfers are feasible but variable fees are not, hosting one specialist is weakly better from a joint profit perspective than hosting two or more specialists. The reason is that now, if M chooses not to serve A-types, then M and the hosted specialist(s) together always extract the entire hosting surplus from B-types. However, if M chooses to serve A-types, joint profits are higher with one hosted specialist because it can extract $\lambda_B \Delta$ from B-types, whereas competing hosted specialists extract nothing. When both fixed transfers and variable fees are feasible, then joint profits are the same regardless of how many specialists are hosted. The reason is that M can use the variable fee τ to offset competition among specialists on its platform.[16] Thus, in our benchmark setting, there is no joint gain from hosting more than one specialist.

Richer settings can make hosting multiple specialists jointly profitable. First, if the specialists are horizontally differentiated, being able to give consumers access to as many different specialists on M will raise consumers' willingness to pay for A. This leads M to want to host as many specialists as possible. Another reason for doing so is network effects: inviting multiple differentiated specialists (that appeal to different consumers) attracts more consumers, and these additional consumers make it more attractive for specialists to be hosted by the platform, and so on. Third, if there is uncertainty over how the specialists will perform when hosted, then hosting as many specialists as possible helps maximize the chance that at least one specialist performs well.[17]

On the other hand, there may also be reasons for restricting the number of hosted specialists that are not captured in our benchmark model. Two obvious reasons are capacity constraints (e.g. a gym can only host a limited number of specialist fitness studios) and the cost of dealing with each additional specialist (e.g. writing and enforcing hosting contracts if they cannot be standardized, managing technological integrations and customer support). A more interesting reason arises when the specialists need to make relationship-specific investments. In that case, having more specialists increases competition among them, which reduces each individual specialist's incentive to invest in the first place. In such contexts, the platform may want to restrict the number of specialists so as to balance the benefits of product diversity and lower prices for consumers with the specialists' investment incentives.

6 Managerial implications

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

Regardless of whether firms can make use of fixed transfers or variable fees based on monitoring transactions on the platform, our results imply that the multiproduct firm M should host a rival when

[16] We prove these claims formally in the Online Appendix, at the end of Section F.

[17] We formally analyze this case in Section F of the Online Appendix.

the fraction of consumers who value both the core-product and the specialist product is high (i.e. there are many B-type consumers), when the added value of the specialist's version of the product is high (i.e. Δ is high), when the utility that consumers get from the core product is low (i.e. u_A is low), and when the fixed cost of hosting is low (i.e. F is low). If there are sufficiently many consumers who value both products, then hosting should be preferred when consumers benefit more from one-stop shopping (i.e. when σ is high). This helps explain for example why big box gyms (like the New York Sports Club) are increasingly willing to host specialty fitness studios (like the Cyc) under revenue sharing contracts: many users are interested in both standard gym amenities (e.g. weight equipment) and specialized classes (e.g. cycling), and there is a clear benefit from having both collocated. However, if there are a lot of consumers who are interested in the core product A only and M cannot monitor transactions and charge a variable fee, the tradeoff between hosting and non-hosting actually shifts towards non-hosting when shopping costs are high.[18]

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

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

Finally, our results showed that M's decision to host rivals also depends on the availability of multiple competing specialists. The existence of two or more competing specialists makes it less likely that M will host: indeed, hosting is no longer sufficient to avoid head-to-head competition because competition will continue to exist outside of the platform. In this case, hosting many specialists raises the value of the platform but it also intensifies on-platform competition, so it should be desirable provided the platform has sufficient instruments to extract the resulting surplus. This seems to be the case for Apple's iPhone App Store: if Apple hosts rival app developers for a product category, then it is open to all such qualified apps. Clearly, Apple has instruments to extract surplus from these rivals, such as taking a cut on all app purchases (including in-app purchases), as well as via taking a share of their advertising revenues. On the other hand, if the platform does not have enough pricing instruments or pricing power to extract surplus from the hosted firms, and shopping costs are

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

sufficiently low, firms should limit how many specialists are hosted in order to avoid having multiple high-quality specialists on the platform that compete away the additional surplus they offer. For instance, on its QuickBooks Capital lending marketplace for its QuickBooks small business customers, Intuit only allows a limited number of selected lenders in order to avoid aggressive price competition (i.e. competition on loan rates).

7 Future directions

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

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

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

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

References

[1] Abhishek, V., K. Jerath and Z.J. Zhang (2016) "Agency Selling or Reselling? Channel Structures in Electronic Retailing," *Management Science*, 62(8), 2259–2280.

[2] August, T., H. Shin and T.I. Tunca (2013) "Licensing and competition for services in open source software," *Information Systems Research*, 24(4), 1068-1086.

[3] August, T., H. Shin and T. I. Tunca (2017) "Generating Value Through Open Source: Software Service Market Regulation and Licensing Policy," *Information Systems Research*, 29(1), 186-205.

[4] Boudreau, K. (2010) "Open platform strategies and innovation: Granting access vs. devolving control," *Management Science*, 56(10), 1849-1872.

[5] Carlton, Dennis W., J.S. Gans, and M. Waldman (2010) "Why tie a product consumers do not use?" *American Economic Journal: Microeconomics* 2(3): 85-105.

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

[7] Eisenmann, T., G. Parker and M. Van Alstyne (2011) "Platform envelopment," *Strategic Management Journal*, 32, 1270-1285.

[8] Farrell, J., and N. T. Gallini (1988) "Second-sourcing as a commitment: monopoly incentives to attract competition," *Quarterly Journal of Economics*, 103(4), 673-694.

[9] Hagiu, A. and D. Spulber (2013) "First-Party Content and Coordination in Two-Sided Markets," *Management Science*, 59(4), 933-949.

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

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

[12] Hagiu and Wright (2018) "Controlling versus Enabling," *Management Science*, forthcoming.

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

[14] Lal, R. and C. Matutes (1989) "Price Competition in Multimarket Duopolies," *RAND Journal of Economics*, 20(4), 516-537.

[15] Lal, R. and C. Matutes (1994) "Retail Pricing and Advertising Strategies," *Journal of Business*, 67(3), 345-370.

[16] Li, Z. and A. Agarwal (2017) "Platform Integration and Demand Spillovers in Complementary Markets: Evidence from Facebook's Integration of Instagram," *Management Science*, 63(10), 3438-3458.

[17] Mantovani, A. and F. Ruiz-Aliseda (2016) "Equilibrium innovation ecosystems: the dark side of collaborating with complementors," *Management Science*, 62(2), 534-549.

[18] Matutes, C. and P. Regibeau (1988). "Mix and match: product compatibility without network externalities," *RAND Journal of Economics*, 19(2), 221-234.

[19] Niculescu, M. F., D.J. Wu and L. Xu (2018). "Strategic intellectual property sharing: Competition on an open technology platform under network effects," *Information Systems Research*, forthcoming.

[20] Parker, G. and M. Van Alstyne (2018) "Innovation, openness, and platform control," *Management Science*, 64(7), 3015-3032.

[21] 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.

[22] Rhodes, A. and J. Zhou (2018) "Consumer Search and Retail Market Structure", *Management Science*, forthcoming.

[23] Salop, S. C. and D. T. Scheffman (1983), "Raising rivals' cost," *American Economic Review Papers and Procceedings*, 73(2), 267-271.

[24] Whinston, M. (1990) "Tying, Foreclosure, and Exclusion," *American Economic Review*, 80(4), 837-59.

[25] White, A. (2013) "Search engines: Left side quality versus right side profits," *International Journal of Industrial Organization*, 31(6), 690-701.

[26] Zhou, J. (2014) "Multiproduct Search and the Joint Search Effect," *American Economic Review*, 104(9), 2918-2939

[27] Zhu, F. and Q. Liu (2018) "Competing with Complementors: An Empirical Look at Amazon.com," *Strategic Management Journal*, forthcoming.

8 Appendix

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

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

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

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

8.2 Proof of Propositions 2 – 5

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

Consider the case with hosting. We start by solving for the equilibrium in the second stage for a given τ. If $\tau > \sigma$, then S prefers to sell directly instead of through M, and then the outcome is the same as under non-hosting. Thus, we focus on $\tau \leq \sigma$. The simpler case with $\tau = 0$, which establishes Proposition 2, will be discussed at the end. In equilibrium, S must sell B_S to all B types through M because this is how it can offer the highest utility for the B product. There are two cases depending on p_A. (Throughout this proof and subsequent proofs characterizing the hosting equilibrium we ignore the fixed cost of hosting F, since it is irrelevant for the analysis, but we do consider it when determining the tradeoff with non-hosting.)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Note that $p'_A \geq u_A - \sigma$ because $p_S + \sigma \geq p_B + \Delta$. Deviation profits are therefore $\lambda_B (u_A + \min\{0, p_S - p_B - \Delta\} + \tau)$, whereas M's equilibrium profits are $u_A - \sigma + \lambda_B \tau$. For the deviation not to be profitable, we need

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

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

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

$$u_A - p_A + u_B + \Delta - \widehat{p}_S - \sigma \geq \max\{u_A - p_A + u_B - p_B - \sigma, u_B + \Delta - p_S - \sigma\}. \quad (1)$$

In this case we must also worry about consumers' non-negative utility constraint

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

If this constraint is not binding, as in the previous case, the condition (1) must hold with equality in equilibrium, so

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

and for the same reason as in the previous case, we must have $u_A - p_A + p_S \geq p_B + \Delta$. If the non-negative utility constraint is binding, S cannot raise $\widehat{p}_S$ and sell on the platform, irrespective of the value of p_B. So we can set $p_B = \widehat{p}_S - \Delta$ (the maximal possible price for B_M) without affecting equilibrium conditions. Thus in both cases, we have

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

Then the non-negative utility constraint is equivalent to

$$u_A - p_A + u_B - p_B - \sigma \geq 0. \tag{2}$$

Also, given $\tau \leq \sigma$, as in the previous case, S does not want to deviate by setting p_S slightly below $p_B + \Delta - \sigma$ and a sufficiently high $\widehat{p}_S$, such that B types prefer to buy B_S from S outside M. Combined with the fact that S makes no sales at p_S in equilibrium, this implies that joint equilibrium profits are maximized by setting p_S sufficiently high that it places no constraint on other equilibrium prices.

Furthermore, as in the previous case, we must have

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

Equilibrium profits are $\lambda_B(p_A + \tau)$ for M and $\lambda_B(p_B + \Delta - \tau)$ for S. Clearly, M benefits from increasing p_A as much as possible subject to $p_A \leq u_A$ and (2),[19] which means $p_A = u_A + \min\{0, u_B - p_B - \sigma\}$. Thus, maximal equilibrium profits are $\lambda_B(u_A + \min\{0, u_B - p_B - \sigma\} + \tau)$ for M and $\lambda_B(p_B + \Delta - \tau)$ for S. If $u_B \geq \tau + \sigma$, then M's profits are $\lambda_B(u_A + \tau)$ for all $p_B \leq \tau$ and since S's profits are increasing in p_B, we set p_B to the highest possible value in order to maximize joint profits, i.e. $p_B = \tau$. Joint profits are then $\lambda_B(u_A + \tau + \Delta)$. If $u_B < \tau + \sigma$, then joint profits are increasing in p_B for $p_B \leq u_B - \sigma$ and constant in p_B for $u_B - \sigma \leq p_B \leq \tau$. Again, we choose p_B to maximize joint profits subject to the constraints we have determined so far, so without loss of generality from a joint profit perspective, $p_B = u_B - \sigma$ (note that among the prices p_B that maximizes joint profits, this is the price that yields the highest profits for M). In this case, joint profits are $\lambda_B(u_A + u_B + \Delta - \sigma)$. Combining the two cases, $p_A = u_A$, $p_B = \min\{\tau, u_B - \sigma\}$, profits for M are $\lambda_B(u_A + \tau)$ and profits for S are $\lambda_B(\min\{0, u_B - \sigma - \tau\} + \Delta)$.

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

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

Since joint profits are increasing in τ for both the equilibrium in which M sells to both types of consumers and the equilibrium in which M just sells to B types, M will set the highest possible τ compatible with hosting, which is $\tau = \sigma$. Then the equilibrium with M selling to both A types and B types yields joint profits $u_A - \sigma + \lambda_B(\sigma + \Delta)$, whereas the equilibrium with M selling to B types only yields joint profits $\lambda_B(u_A + \Delta + \min\{\sigma, u_B - \sigma\})$ whenever it exists, i.e. whenever $\sigma \geq \lambda_A u_A$. The latter equilibrium has higher joint profits whenever

$$\lambda_A < \frac{\sigma + \min\{0, u_B - 2\sigma\}}{u_A + \min\{0, u_B - 2\sigma\}} \leq \frac{\sigma}{u_A}.$$

[19]Indeed, as noted above, p_S can be set sufficiently high such that the constraint $u_A - p_A + p_S \geq p_B + \Delta$ is never binding.

In conclusion:

- If $\lambda_A \leq \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, then the highest equilibrium joint profit under hosting is $\lambda_B\left(u_A+\Delta+\min\{\sigma,u_B-\sigma\}\right)$. Comparing with the equilibrium joint profits without hosting $u_A-\sigma+\lambda_B(\sigma-\Delta)$, and taking into account the fixed cost of hosting F, hosting is preferred if and only if $\Delta > \frac{F-(\sigma-\lambda_A u_A+(1-\lambda_A)\min\{0,u_B-2\sigma\})}{2(1-\lambda_A)}$.
- If $\lambda_A > \frac{\sigma+\min\{0,u_B-2\sigma\}}{u_A+\min\{0,u_B-2\sigma\}}$, then the highest equilibrium joint profit under hosting is $u_A-\sigma+\lambda_B(\sigma+\Delta)$. In this case, hosting is preferred if and only if $\Delta > \frac{F}{2(1-\lambda_A)}$.

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

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

8.3 Proof of Proposition 6

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

The result then follows from the comparison between the hosting profits given in Proposition 5 and the profits without hosting derived in Proposition 11 in the Online Appendix.