Regulating Platform Fees under Price Parity

Renato Gomes*
Andrea Mantovani†

PRELIMINARY AND INCOMPLETE - PLEASE DO NOT CIRCULATE

First Version: ...
This Version: 29/11/2019

Abstract

Online market places, such as Amazon's, or online travel agencies (such as Booking.com), greatly expand consumer information about market offers, but also raise firms' marginal costs by charging high commissions (which in some cases reach 25%). To prevent showrooming, platforms adopted price parity clauses, which restrict sellers' ability to offer lower prices in alternative sales channels. Whether to uphold, reform, or ban price parity has been at the center of the policy debate, but so far little consensus has emerged. In this paper, we investigate a natural alternative to lifting price parity; namely, we study how to optimally cap platforms' commissions. The optimal cap reflects the pigouvian precept according to which the platform should not charge fees greater than the externality that its presence generates on other market participants. Employing techniques from extreme-value theory, we are able to express the optimal cap in terms of observable quantities. In an application to online travel agencies, we find that current average fees are welfare increasing only if platforms at least double consumers' consideration sets (relative to alternative ways of gathering information online). This suggests that, in some markets, regulation capping commissions should bind if optimally set.

JEL Classification: D83, L10, L41.

Keywords: platforms, price parity, regulation, commission caps, extreme value theory.

*Toulouse School of Economics, CNRS, University of Toulouse Capitole: renato.gomes@tse-fr.eu.
†University of Bologna, Department of Economics: a.mantovani@unibo.it

1 Introduction

Most platforms operate under the agency model, according to which sellers are free to set prices, but are charged by the platform a commission per transaction. For instance, on Amazon Marketplace, professional sellers pay on average 13 percent per sale, whereas in Booking.com the average fee is 20 percent. On top of that, dominant platforms often employ contractual restrictions and business practices that impose additional constraints to client sellers. Some of these practices have been investigated by different antitrust authorities, with a recent focus on the adoption of price parity clauses. These covenants restrict sellers' ability to charge lower prices on alternative sales channels.[1]

Price parity clauses are widespread in the e-commerce and lodging sectors, but have also been applied to other industries such as entertainment, insurance, digital goods, and payment systems. Platforms claim price parity is essential to the business, as it avoids *showrooming*, whereby consumers use the platform to find the relevant information about the seller, but then switch to the direct channel to obtain a discount. By contrast, regulators and consumer associations often regard price parity as the source (or, at best, a reinforcer) of the platforms' market power.

In the past years, EU national competition authorities have reached important decisions on price parity clauses adopted by Amazon and by leading online travel agencies (OTAs), such as Booking.com and Expedia. Currently, all types of PPCs are forbidden in France, Italy, Belgium, and Austria, whereas in Germany they are prohibited for Amazon, Hotel Reservation System (HRS) and Booking.com, and in the UK for Amazon. In the US, Amazon recently decided to remove PPCs from its contracts with third-party marketplace sellers.[2]

Yet, it is not entirely clear that a ban, or voluntary removal, of price parity actually produces tangible results. For one, sellers might still practice price parity so to remain in good terms with the platform. Moreover, in some countries, such as France, the law forbids the imposition of price parity, but allows it if voluntarily accepted by sellers. In many preferred partner programs (PPPs) created by OTAs, price parity is the counterpart for top listing sellers. As joining PPP's is a voluntary action, such programs are often a legal way to bypass the ban.

All in all, whether one should uphold, reform, or ban price parity has been at the center of the policy debate, but so far little consensus has emerged. In this paper, we investigate a natural alternative to restrict price parity; namely, we study how to optimally cap platforms' commissions.

In our baseline model, we consider a monopolist platform which imposes price parity on sellers and charges them a fee per sale. Our starting observation is that platforms greatly expand consumer information about market offers, therefore augmenting each seller's potential demand (which is the

[1] While *wide* price parity prevents sellers from posting a lower price on *any* alternative sales channel, under *narrow* price parity, sellers may differentiate prices across platforms, but cannot charge a lower price on the direct sales channel.

[2] Senator Richard Blumenthal reportedly sent in December 2018 influential letters to the Justice Department and the Federal Trade Commission demanding investigations into Amazon's contracts with marketplace sellers.

set of consumers who consider the seller when making a purchasing decision). Conversely, if a seller does not join the platform, his potential demand consists of a much smaller set of consumers: those who know the seller from other sources (such as friends, or a previous purchase). Platforms also add convenience to transactions, what makes them the preferred sales channels by consumers.

The platform leverages on consumers' (lack of) information and sellers' head-to-head competition to levy high commissions. The equilibrium fee is chosen to leave each firm indifferent between (i) delisting from the platform, facing a much reduced potential demand, but competing with lower marginal costs than all other firms (who face the platform's commission), and (ii) remaining in the platform, enjoying a much expanded potential demand, but competing with all other firms under no marginal cost advantage.

Crucially, the platform's market power stems from the following contractual externality: delisting from the platform sends the deviating firm to a world where all consumers who *still* consider the firm *also* consider *all* competing firms operating in the market. This makes the deviating firm face a degree of competition (among its remaining potential customers) much higher than in a world where no platform is available. This reduces profits outside of the platform, and induces firms to accept paying high commissions. On the aggregate, firms might pay in commissions substantively more than the profit gain generated by the platform's service. Consumers may also be hurt; as prices may increase more than the gain from enjoying better market information.

In light of this market failure, we consider regulation capping the platform's fee per sale. If the regulator is utilitarian (assigning equal weights to the consumers' surplus, firms' profits, and the plaform's profits), the optimal cap takes a familiar form: It lets the platform charge no more than the expected externality it imposes on other market participants. This pigouvian rule has similarities to the "tourist test" regulation adopted in the payment industry;[3] the main difference being that the contribution to welfare generated by the platform is essentially informational in our setting, as it enables consumers to realize purchases of much higher (match) value.

Measuring the (expected) contribution to welfare imputable to the expansion on consumers' consideration sets is typically challenging, as it requires fine knwoledge of consumers' match values. To circunvent this difficulty, we apply techniques from extreme-value theory to express, in the context of random utility models, variations in consumer surplus as a function of (more easily) observable quantities. Namely, we show that, in large markets, the consumer's informational gain from considering more firms is asymptotically equivalent to the firms' markup multiplied by the rate of expansion on consumer consideration sets.

Remarkably, this approximation result may be a useful tool for policy-makers, as we illustrate in

[3] According to which the merchant fee cannot exceed the merchant's convenience benefit of a card payment. Under this condition, the consumer's choice of payment instrument imposes no negative externality on the merchant, which implies the payment system's aggregate price (across consumers and merchants) is no more than the aggregate benefit of a card payment. See

an application to online travel agencies. We find that current average fees are welfare increasing only if platforms at least double consumers' consideration sets (relative to alternative ways of gathering information online). In light of other sources of information easily available in internet (e.g., Google), this suggests that regulation capping commissions might bind in some markets, if optimally set.

1.1 Related Literature

Price parity clauses, also known as "most-favoured nation" (MNF) clauses, came back to the fore in the economic debate. In the traditional wholesale setting, a relatively large theoretical literature emphasized the role of MFN agreements as a commitment device not to price discriminate between retailers (see, *e.g.*, Schnitzer, 1994; Besanko and Lyon, 1993). A recent stream of literature focused instead on the price parity clauses practiced by platforms, where the contractual relationship usually follows the "agency model": Sellers decide the final price on the platform, which charges a commission rate per transaction.

The majority of these papers emphasizes the anticompetitive effect of price parity clauses. Edelman and Wright (2015) examine consumers' decision to either purchase directly or through a platform, who may invest to provide a non-pecuniary benefit to consumers. They show that, under price parity, the platform over-invests in the provision of the non-pecuniary benefit, diverting consumers from direct channels. This may lead to an increase in final prices, and a decrease in social welfare. In turn, Boik and Corts (2016) consider a unique supplier who commercializes her entire production through two differentiated platforms. They find that price parity clauses lead to higher commissions, which, in turn, increase final prices and prevent entry by low-cost competitors. Johnson (2017) compares the wholesale and the agency models in a framework with multiple suppliers and retailers. He also studies the effects of price parity clauses, obtaining results in line with Boik and Corts (2016).

In Wang and Wright (2019) platforms provide both search and intermediation services. Consumers positively value these services, but can decide to free-ride if direct purchasing is allowed. In this context, price parity typically hurts consumers, except when it is essential for the viability of the platform. EXPAND

A discording view is offered by Johansen and Vergé (2017), who find that price parity is pro-competitive in some circumstances. They propose a model where consumers are aware of all firms in the market, multi-home across all platforms, and are also able to purchase directly with firms. Consumers perceive all these options as (horizontally) differentiated, what generates market power to platforms even in the absence of price parity. When this clause is imposed in either its wide or narrow forms, firms become more prone to delisting, which reduces average costs. There are instances where the participation constraint is so tight that commissions decrease (relative to unrestricted pricing), benefiting consumers and firms. In such cases, the Pigouvian regulatory cap proposed here is slack, as the platform (practicing price parity) generates a Pareto improvement.

These contributions take a competition policy perspective, lending support to banning price parity, or to favoring its wide or narrow version. By contrast, we study optimal regulation, which maintains price parity, but restricts the platform's ability to levy high commissions. The regulatory approach is arguably more flexible than its competition policy counterpart: As we show, lifting price parity is akin to capping the platform commission at an inefficiently low level.

At the empirical level, Hunold *et al.* (2018) used metasearch data to show that hotels in Germany expanded room availability on different platforms and increased the number of sales channels after the Bundeskartellamt prohibited *Booking.com* from using price parity clasues in 2015 They also showed that hotels charged the lowest price on the direct channel more often in Germany than in countries that did not abolish such clauses. Cazaubiel *et al.* (2018) obtain a dataset from two major hotel chains in Scandinavia with prices, volumes and sales channels between 2012 and 2016. They estimated the degree of substitution between *Booking.com* and Expedia, and hotels' own websites, by considering a boycott against Expedia led by hotels between 2012 and 2014. Finally, Mantovani *et al.* (2019) provided quasiexperimental evidence on the full removal of price parity clauses in France in 2015 and in Italy in 2017 for hotels listed on *Booking.com*. They revealed a relatively limited effect in the short run followed by a significant reduction in room prices in the medium run. They also found that hotels affiliated with chains decreased their prices more than independent hotels, both in the short and medium run. In this paper we use data taken from different sources in order to estimate the optimal level of the commission fee from a social welfare standpoint.

2 Model and Preliminaries

Consider an economy populated by N firms, indexed by $j \in \mathcal{N} \equiv \{1, \ldots, N\}$, and a unit-mass continuum of consumers $\mathcal{I} \equiv [0, 1]$ with single-unit demands. A consumer's gross utility from firm j's product is given by $\hat{v}_j = v_j + z_j$, where v_j is the vertical component of preferences common to all consumers (for instance, the number of stars of a hotel), while z_j is the consumer-specific match value of firm j (for instance, the hotel's proximity to some location of interest). We assume that, for each consumer, $z \equiv (z_1, \ldots, z_N)$ is a draw (iid across consumers) from a symmetric distribution G with support contained on $\mathbb{R}^N_+$ and density g. Each firm j faces a constant marginal cost c_j per consumer served.

We say that a firm (call it j) belongs to the consideration set of a consumer if he/she observes the pair $(\hat{v}_j, p_j)$, where p_j is the price charged by firm j. Consumers can only transact with firms in their consideration sets. Not buying from any firm generates a zero payoff to consumers.

Consumers are heterogeneous on their consideration sets. We describe this heterogeneity by means of a *consideration profile* $\sigma : 2^{\mathcal{N}} \to \mathcal{B}[0, 1]$, which maps each subset of firms (contained in the power set $2^{\mathcal{N}}$) into the (measurable) set of consumers who consider exactly that set of firms.[4] Because we normalized the mass of consumers to one, it follows that

$$\sum_{s \in 2^{\mathcal{N}}} |\sigma(s)| = 1,$$

where $|\sigma(s)|$ is the Lebesgue measure of the set $\sigma(s)$. The set of consumers whose consideration sets contain firm j (among other firms) then equals

$$d_j[\sigma] \equiv \underset{\{s: j \in s\}}{\cup} \sigma(s),$$

which we call firm j's *potential demand* under the consideration profile σ. We let $d_\emptyset[\sigma] \equiv \sigma(\emptyset)$ be the market's *latent demand*, which comprises all consumers who would be interested in consuming the good, but cannot do it for not knowing any firm in the market.

A particular class of consideration profiles plays an important role in our analysis. We say that the profile σ is *symmetric* if the following condition hold: First, those consumers who possess *some* market information enjoy consideration sets of the same size. Moreover, all firms are considered by the same number of consumers, and therefore have potential demands of the same size. Formally, this means that there exists a number $n \in \mathcal{N}$, which we call the *reach* of σ, such that: (i) $|\sigma(s)| = |\sigma(s')| > 0$ whenever s and s' have n elements, and (ii) $|\sigma(s)| = 0$ whenever $s \neq \emptyset$ has cardinality different from n. While special, this class of consideration profiles offers a tractable way to study changes in consumers' information about market offers.

Before consulting the platform, consumers' consideration sets are described by the profile $\underline{\sigma}$, which captures all the information learnt by consumers outside of the platform (through advertising, travel or shopping guides, friends' recommendations, previous experiences, etc). For simplicity, we assume $\underline{\sigma}$ to be symmetric with reach $\underline{n} < N$.

Once a consumer visits the platform, all firms *listed in the platform* are added to the consideration set of the consumer. For instance, if all firms join the platform, consumer information is then described by the consideration profile $\bar{\sigma}$, which is symmetric with maximal reach N. Accordingly, the platform expands by a factor $\frac{N}{\underline{n}}$ the size of the consideration sets of those consumers who possess *some* market information. The platform also brings to the market the latent demand $d_\emptyset[\underline{\sigma}]$ of consumers that were originally unaware of *any* firm.

Alternatively, suppose all firms join the platform, except for some firm j, which refuses to do it. The consumer information is then described by the consideration profile σ^{-j} such that all consumers that considered firm j under $\underline{\sigma}$ (i.e., before consulting the platform) now consider all firms in the market, whereas those consumers who did not consider firm j under $\underline{\sigma}$ now consider all firms other

[4]We assume that consideration sets are independent of the consumer's profile of match values z. We also assume that $\sigma(s)$ is a borelian subset of $[0,1]$. The collection of such sets is denoted by $\mathcal{B}[0,1]$.

than j. This leads to

$$\sigma^{-j}(\mathcal{N}) = d_j[\underline{\sigma}], \quad \text{and} \quad \sigma^{-j}(s) = \begin{cases} \mathcal{I}/d_j[\underline{\sigma}] & \text{if} \quad s = \mathcal{N}/\{j\} \\ \emptyset & \text{if} \quad s \neq \mathcal{N}, \mathcal{N}/\{j\} \end{cases}.$$

For simpicity, we assume that visiting the platform is costless, so all consumers visit it provided they expect firms to be listed there. Besides providing information, the platform offers a business interface, enabling consumers to finalize transactions with firms. Completing a transaction within the platform generates a convenience benefit $b \geq 0$ to firms, but costs them a fee to be paid to the platform. The platform privately offers firm-specific fees; namely, firm j is asked to pay f_j for each sale within the platform.[5] The platform is profit-maximizing, and operates in a market if it expects its profit to exceed some quasi-fixed cost k, which captures the expenses associated with operating costs, the costs of monitoring firms' compliance to the platform rules, as well as advertising. We assume that the platform's operating cost is private information, being a draw from some distribution Φ, with density ϕ and support on $\mathbb{R}_+$.

We assume in the baseline model that the platform is able to impose price parity, according to which firms have to offer the same prices for transactions either within or outside the platform.[6] As a result, if a firm joins the platform, all of its transaction occur within the platform (as consumers would weakly prefer doing so).

The timing of the model is summarized below:

1. The platform privately observes its cost k, and decides whether to operate or not,
2. In case it operates, the platform privately offers the fee f_j, for each firm $j \in \mathcal{N}$.
3. Firms simultaneously set prices and decide whether (or not) to join the platform,[7]
4. Each consumer makes a purchasing decision considering the firm(s) he is aware of.

Our solution concept is perfect bayesian equilibrium with passive beliefs (for short, equilibrium). That beliefs are passive means that, upon receiving an out-of-equilibrium offer, firms do not change their belief about the fee offered to other firms. Moreover, to simplify matters, we restrict attention to symmetric markets where the expected gains from trade are identical across firms. This amounts to assuming that $\delta \equiv v_j - c_j$ is invariant in j (that is, as quality increases, marginal costs increase by the same amount).

In what follows, we remove subscripts to denote price profiles (i.e., $p \equiv (p_1, \ldots, p_N)$), and write that $p_{-j} \equiv (p_1, \ldots, p_{j-1}, p_{j+1}, \ldots, p_N)$. We use analogous notation for v, c and z.

[5] Add footnote defending private contracting.

[6] We relax this assumption in Section X.

[7] We could alternatively assume that firms choose prices after observing the joining decisions of all other firms. However, assuming prices are chosen simultaneously to joining decisions simplifies the analysis. Both formulations lead to similar results.

Pricing Equilibrium

Consider a symmetric consideration profile σ with reach n. We will now derive the demands faced by firms under σ. To do so, let the profiles (v, p) be such that the market is fully covered, (i.e., $v_k \geq p_k$ for all $k \in \mathcal{N}$), and denote by

$$H_{j,k|s}(x) \equiv \text{Prob}_G \left[z_k - z_j \leq x \quad \text{and} \quad k = \arg\max_{k'} \{v_{k'} + z_{k'} - p_{k'}\} \quad \text{s.t.} \quad k' \in s/\{j\} \right]$$

the probability (induced by the joint G) that consumers' relative match values between firms j and k is less than $x \in \mathbb{R}$, *and* that firm k is j's best competitor in the set s. The demand faced by firm j under σ is then

$$D_j\left(p_j, p_{-j}; \sigma\right) = \sum_{\{k,s:|s|=n,j,k\in s,k\neq j\}} \sigma(s) H_{j|k,s}\left(v_j - p_j - (v_k - p_k)\right).$$

The best response of firm j to the price profile p_{-j} is then

$$P_j\left(p_{-j} \left| \sigma, c_j \right.\right) \equiv \arg\max_{p_j} \quad D_j\left(p_j, p_{-j}; \sigma\right)(p_j - c_j), \tag{1}$$

while an equilibrium p^* is a price profile satisfying $p_j = P_j\left(p_{-j} \left| \sigma, c_j \right.\right)$ for all $j \in \mathcal{N}$.

Before characterizing equilibrium, we shall introduce the following regularity condition, which guarantees the quasi-concavity of firms' best responses.

Assumption 1 ***(regularity)*** *Let $n \geq 2$ and consider the cdf*

$$H^{(n)}(x) \equiv Prob_G\left[z_1 - z_2 \leq x \left| z_2 \geq \max\{z_2, \ldots, z_n\}\right.\right],$$

with density $h^{(n)}(x)$ over $\mathbb{R}$. Then

$$x - \left(\frac{1 - H^{(n)}(x)}{h^{(n)}(x)}\right)$$

is increasing in x.

We say that an equilibrium is *symmetric* if $v_j - p_j \geq 0$ is constant in j.[8] Accordingly, in symmetric equilibria, prices increase one-to-one with the "vertical" quality of a firm (e.g., the number of stars in a hotel). The next lemma characterizes the unique symmetric equilibrium.

Lemma 1 ***(pricing)*** *Suppose that firms compete under the consideration profile σ, assumed to be symmetric with reach $n \geq 2$. Then the unique symmetric equilibrium is such that, for all $j \in \mathcal{N}$,*

$$p_j^* = c_j + \lambda(n), \qquad \textit{where} \qquad \lambda(n) \equiv \frac{1 - H^{(n)}(0)}{h^{(n)}(0)}.$$

[8]This definition also implies the market is fully covered, that is, all consumers buy the good from some firm.

Lemma 1 shows that, in the family of discrete-choice models of this paper, equilibrium prices consist of marginal costs plus the firms' markup under n-sized consideration sets, $\lambda(n)$. For instance, when the platform operates, and all firms join at some symmetric fee f, equilibrium prices are $p_j^* = c_j + f + \lambda(N)$.

The logit model of product differentiation is a special case of our setting when, for each consumer, firm's match values are iid from an extreme value distribution. Another special case is the spokes model of Chen and Riordan (200x), which generalizes Hotelling. In this case, the support of G only contains profiles z for which there exist firms j and k such that $z_j, z_k \geq 0$, $z_j + z_k = 1$ and $z_l = 0$ for all $l \neq j, k$.

3 Laissez-Faire

We now study the equilibrium outcome under a monopolistic platform that is able to impose price parity, and faces no regulation of any kind. The platform's profit from firm j is $D_j f_j$, where D_j is the realized demand of firm j within the platform, and f_j is the fee per sale contracted with firm j. The platform's total profit adds up revenues across all firms that join.

We focus on symmetric equilibria. The unique such equilibrium is characterized in the next proposition.

Proposition 1 ***(equilibrium)*** *There exists a unique symmetric equilibrium. In this equilibrium, all firms join the platform and pay a fee $f^* > b$, which solves*

$$\frac{\lambda(N)}{N} = |d_j[\underline{\sigma}]| \cdot \max_{\Delta p} \left\{ \left(1 - H^{(N)}(\Delta p)\right) (\Delta p + f^* + \lambda(N) - b) \right\}. \tag{2}$$

To understand Proposition 1, consider the behavior of a firm that deviates from the putative equilibrium, and chooses not to join the platform. On the one hand, the firm avoids paying the net fee $f^* - b$. On the other hand, the firm sees it potential demand reduced to $d_j[\underline{\sigma}]$, which is the set of consumers that are aware of the firm's existence *before* visiting the platform.

Taking these effects into consideration, the deviating firm chooses a new price to lure those consumers in $d_j[\underline{\sigma}]$. Letting Δp denote the price adjustment relative to the equilibrium price, the firm's problem is represented in the left-hand side of equation (2). The first-order condition associated with this program is

$$\Delta p - \left(\frac{1 - H^{(N)}(\Delta p)}{h^{(N)}(\Delta p)} \right) + f^* + \lambda(N) - b = 0. \tag{3}$$

As the left-hand side is increasing in Δp (by Assumption 1), it follows that the optimal price adjustment satisfies $\Delta p \leq 0$ (i.e., is a discount) if and only if the net fee $f^* - b$ is positive.

Crucially, this price adjustment increases profit, and the more so the higher is the equilibrium fee f^* incurred by all competitors inside the platform. Condition (2) states that, in equilibrium, the platform chooses its fee to leave each firm indifferent between (i) delisting from the platform, facing a much reduced potential demand, but competing against other firms with inflated marginal costs, and (ii) remaining in the platform, enjoying a much expanded potential demand, but competing with all other firms with no marginal cost advantage.

The platform's equilibrium fee exceeds the convenience benefit of an intermediated transaction: $f^* > b$. To see it formally, note that the right-hand side of (2) is increasing in f^*, and that, at $f^* = b$, the optimal price adjustment is $\Delta p = 0$ (as implied by equation 3). Therefore, at $f^* = b$, the right-hand side of (2) equals

$$d_j[\underline{\sigma}]\frac{\lambda(N)}{N},$$

which is obviously less than the profit from joining the platform, $\frac{\lambda(N)}{N}$. Therefore f^* has to increase above b to render the firms indifferent between particpating or delisting from the platform.

In fact, the equilibrium platform fee may be high enough to actually decrease both firms' profits and consumers' surplus relative a world where no firm joins the platform (as discussed in detail in the next section). What explains the platform's ability to extract more rents from consumers and firms than it generates? Its source of market power lies in the following *contractual externality*: delisting from the platform sends the deviating firm to a world where all consumers that consider the firm (its pre-visit potential demand) also consider *all* competing firms operating in the market. This renders the degree of competition faced by the deviating firm much higher than in a world where no platform is available. This obviosuly reduces profits outside of the platform, and induces firms to accept paying high commissions. Aggregate producer surplus might well decrease relative to the no-platform benchmark.

Corollary 1 ***(comparative statics)*** *Consider two pre-visit consideration profiles, $\underline{\sigma}_0$ and $\underline{\sigma}_1$, and let f_0^* and f_1^* be their respective equilibrium fees once a monopolist platform enters the market. Then $f_0^* \leq f_1^*$ if and only if $d_j[\underline{\sigma}_0] \geq d_j[\underline{\sigma}_1]$.*

According to Corollary 1, the pre-visit information profile affects the equilibrium fee *only* through its potential demand $d_j[\underline{\sigma}]$: as it increases, the equilibrium fee goes down. In particular, as long as the potential demand remains constant, variations in the reach of $\underline{\sigma}$ (which determines the size of pre-visit consideration sets, and therefore the degree of competition among firms), as well as variations in the size of the latent demand, $d_\emptyset[\underline{\sigma}]$, have no effect on the equilibrium fee. This prediction distinguishes our model from alternative theories of aggregator platforms, and can be brought to data if a cross-section of market fees and potential demands are available.

Remark 1 ***(public platform fee)*** *The baseline model assumes that the platform makes each firm a private offer regarding the commission. The equilibrium characterized in Proposotion 1 remains an equilibrium if, alternatively, one assumes that the platform sets a public fee, observable by all firms (before simultaneous joining and pricing decisions are made).*

4 Cap Regulation

In deriving optimal regulation, we have to account for the possibility that the platform may refuse to operate if its ability to extract rents from firms is too limited.[9] The reason is that the gains from tighter regulation have to compared to the losses from having no platform to centralize market information. One main difficult with this comparison is that consumers' information acquisition behavior may change depending on the existence of an aggregator platform. Indeed, in a world where Expedia is avaliable, tourists planning a trip might directly visit this platform, rather than looking for travel guides or using other online sources. Accordingly, the consideration profile $\underline{\sigma}$, which describes consumer information *before* visiting the platform, *but* with knowledge that the platform is available, likely exhibits a small reach $\underline{n}$. By contrast, the "counterfactual" consideration profile $\hat{\sigma}$, which describes consumer information in a world without a platform, arguably exhibits a reach $\hat{n}$ larger than $\underline{n}$. Indeed, the need for information should push consumers to consult sources that would be considered redundant in a world where the platform is available (e.g., regular search engines). As we shall see, regulation crucially depends on conjecturing by how much the platform expands the consideration set of consumers in equilibrium (relative to the counterfactual where no platform is available).

4.1 Mature Markets

Let us consider first the case of a *mature* market, where all consumers (even in the absence of a platform) possess some (though partial) market information. Equivalently, this amounts to assuming that the counterfactual latent demand $d_{\emptyset}[\hat{\sigma}]$ is null.

We consider regulatory interventions consisting of a cap on the platform's fee, akin to what is practiced for payment platforms. We denote this cap by $\bar{f}$, and note that the cap is unconsequential if $\bar{f} > f^*$, but binds otherwise. Therefore, the equilibrium platform fee is $f^r \equiv \min\{\bar{f}, f^*\}$. Because all firms join under this fee, the platform's profit also equals f^r.

Our measure of social welfare combines the surplus derived by firms and consumers, which weight is $\alpha \in (0, 1]$, with that of the platform, which weight is $1 - \alpha \in [0, 1)$. Letting $Z^{1:n}$ denote the first-order statistic out of $n \leq N$ coordinates of the random vector z, the planner's objective is

$$W(\bar{f}) \equiv \int_0^{f^r} \left\{\alpha \left(\delta + \mathbb{E}\left[Z^{1:N}\right] - f^r + b\right) + (1-\alpha)\left(f^r - k\right)\right\} d\phi(k) + (1 - \Phi(f^r))\, \alpha \left(\delta + \mathbb{E}\left[Z^{1:\hat{n}}\right]\right),$$

where the integral describes welfare when the platform's cost realization is low (so that it operates), whereas the second term describes welfare when the platform refuses to operate, as costs are high relative to the regulatory cap. When the platform operates, the aggregate surplus obtained by consumers and firms from each realized sale consists of the gains from trade in the absence of

[9] This is the risk mostly emphasized by participants of industries on the verge of regulation - see XYZ...

informational frictions (that is, under reach N), in addition to the convenience benefit b, but discounted by the platform fee f^r. In turn, the platform's profit is simply $f^r - k$. By contrast, when the platform stays inactive, the surplus of consumers and firms consists of the expected gains from trade under the counterfactual consideration profile $\hat{\sigma}$ (which reach is $\hat{n}$), as described in the second term of $W(\bar{f})$.

The next proposition derives the welfare-maximizing commission cap.

Proposition 2 ***(optimal regulation: mature)*** *Suppose the market is mature ($d_\emptyset[\hat{\sigma}] = \emptyset$), and consider regulation that mandates the platform's commission to satisfy $f \leq \bar{f}$. The welfare-maximizing cap, written as a function of the weight α, implicitly solves*

$$\bar{f}_\alpha = b + \mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] + \frac{\Phi(\bar{f}_\alpha)}{\phi(\bar{f}_\alpha)}\left[\frac{1-2\alpha}{\alpha}\right]. \tag{4}$$

When the planner is utilitarian, that is, $\alpha = \frac{1}{2}$, the welfare-maximizing cap is such that the platform operates if and only if it increases the surplus from consumers and firms.

Consider first the utilitarian case, where $\alpha = \frac{1}{2}$. Here, the third term in the left-hand of side (4) vanishes, and the cap equals the *surplus-neutral fee* $\bar{f}_{\frac{1}{2}}$, which is the convenience benefit added to the informational gain produced by the platform (relative to the counterfactual consideration profile $\hat{\sigma}$). Under this cap, platform entry cannot hurt consumers and firms, as its profit is bounded by the externality it imposes on the other market participants (in the spirit of the pivot mechanism).

As α grows above $\frac{1}{2}$, the planner gives more weight to consumers and firms, setting the cap below the surplus-neutral fee $\bar{f}_{\frac{1}{2}}$. This is optimal because having the firm not operate in some instances is compensated by the increase in the surplus of consumers and firms that a tighter cap generates. Conversely, when α is below $\frac{1}{2}$, the planner gives more weight to the platform vis-à-vis consumers and firms, setting the cap above the surplus-neutral fee $\bar{f}_{\frac{1}{2}}$. In this case, the presence of the platform may hurt the other market participants.

To get a better sense of how optimal regulation compares to current practice, let us for simplicity focus on the utilitarian case, where the cap equals the surplus-neutral fee $\bar{f}_{\frac{1}{2}}$. While being well-grounded in theory, it is hard to implement this cap in practice, as it requires knowledge of the distributions of consumer tastes across firms. To circunvent this difficulty, we proceed by expressing (4), which contains moments of order statistics, in terms of quantities that are arguably easier to observe (or measure); namely, firms' markups and potential demands. To do so, we employ the approximation techniques of Gabaix et al (2016), based on extreme value theory.

Applying these techniques requires however that we specialize the discrete-choice setup employed so far. Specifically, we have to assume that the consumer match values are independent across firms, that is, the taste vector z is composed of N iid realizations of some distribution G_1 with support on $\mathbb{R}_+$. This specification, commonly referred as the random utility model, was first proposed by Perloff and Salop (1985), and has been widely applied thereafter (see Anderson, de Palma, and Thisse 1992).

Proposition 3 ***(approximation)*** *Suppose that the planner is utilitarian and the market is mature ($d_\emptyset[\hat{\sigma}] = 0$). Also assume that consumer match values are independent across firms, and let their cdf G_1 be well-behaved with tail index $\gamma \neq 0$ and support $(\underline{z}, \bar{z})$, where $0 < \underline{z} < \bar{z} \leq \infty$.*[10] *Then, as $\hat{n}$ and N grow large while satisfying $d_j[\hat{\sigma}] = \frac{\hat{n}}{N}$, we obtain that*

$$\bar{f}_{\frac{1}{2}} - b \underset{N}{\sim} \left(\frac{1 - d_j[\hat{\sigma}]}{d_j[\hat{\sigma}]}\right) \Gamma(\gamma + 2)\lambda(N), \tag{5}$$

where $\Gamma(\cdot)$ is the gamma function, and the notation $a_n \underset{n}{\sim} b_n$ means that $\lim_{n\to\infty} \frac{a_n}{b_n} = 1$.

The result above means that, as the market grows large, the surplus-neutral fee $\bar{f}_{\frac{1}{2}}$ behaves like the markup multiplied by the proportional increase in firms' potential demands. Because $\Gamma(\gamma+2) \approx 1$ for $|\gamma|$ small, this suggests the approximation

$$\bar{f}_{\frac{1}{2}} \approx b + \left(\frac{1 - d_j[\hat{\sigma}]}{d_j[\hat{\sigma}]}\right) \lambda(N),$$

which is remarkably precise, as illustrated in the online appendix.[11]

The convenience benefit b (which captures the valued added by the booking and payment interface provided by the platform) should be comensurate to the market rates of online payment gateways (such as Paypal). So we set it to 2%.

Interestingly, the average fee charged by Booking.com in X and Y markets is below the welfare-neutral fee only if the platform expands consumers' consideration sets by Z percent. Although empirical work is needed to better estimate the magnitudes in (5), casual observation suggests that the current fees might be welfare-decreasing, at least in some markets. As argued above, the thought experiment behind the counterfactual consideration profile $\hat{\sigma}$ involves consumers acquiring information (likely online) about the hotels available in a given location in the absence of the platform. Given that many online services (Google, for instance) offer lists and descriptions of hotels as seamlessly as Booking.com and Expedia, it is hard to think that the potential demand $d_j[\hat{\sigma}]$ is too far from one. All in all, it is challenging to find plausible scenarios where $\bar{f}$ is larger than 10%.

An alternative strategy for employing Propisition (3) is too use the equilibrium condition (2) to retrieve the potential demands of firms. Under the same assumptions above,... The difficulty here is that the potential demands implied by the equilibrium condition (2) apply in a world where consumers know they can count on the platform to learn about the available options. By contrast, the consideration profile (and potential demand) relative for regulation is that where consumers have to use other means to get to hotels...

[10]Borrowing the terminology from extreme value theory, that the cdf G_1 is well behaved means that its density is differentiable, that $\lim_{x\to\infty} \frac{1-G_1(x)}{g_1(x)}$ exists in $\bar{\mathbb{R}}$, and that $\gamma = \lim_{x\to\infty} \frac{d}{dx}\left(\frac{1-G_1(x)}{g_1(x)}\right)$ exists and is finite. Examples of cdf's with tail index $\gamma \neq 0$ include Weibull, Fréchet, and Pareto, among others.

[11]For most distributions of interest, the approximation error is less than 10^{-3} provided $10 \leq \hat{n} \leq N$ and $\frac{N}{\hat{n}} \leq x$.

Remark 2 ***(information acquisition costs)*** *The commission cap derived above implicitly assumes that information acquisition is costless for consumers. If the time spent acquiring information is similar with and without an agregator platform, this (arguably strong) assumption is inconsequential in computing the efficient cap. If however the total cost spent searching changed substatially with a platform, we would have to incorporate this variation in the welfare measure, which is often hard for a variety of reasons (for one, the estimation of search costs is notoriously difficult). It is likely however that, if any change occurred, the time spent on information acquisition decreased with the existence of an aggregator platform. In this case,* $f \leq \bar{f}_{\frac{1}{2}}$ *is a sufficient (but not necessary) condition for the platform to increase welfare.*

4.2 Markets with Growth Potential

An important feature of mature markets is that, whenever profit margins decrease as competition intensifies,[12] the presence of the platform is bound to hurt firms' profits. The reason is that, by enlarging consumers' consideration sets, the platform intensifies competition whithout increasing firms's sales. This conclusion is no longer true when the market has growth potential, as captured by the fact that its latent demand $d_\emptyset[\hat{\sigma}]$ is non-empty.

Corollary 2 ***(industry profits)*** *Relative to the no-platform benchmark, firms gain with the presence of a monopolistic platform if and only if*

$$d_\emptyset[\hat{\sigma}] > 1 - \left(\frac{\lambda(N)}{\lambda(\underline{n})}\right).$$

The existence of a platform expands sales (by awakening the latent demand) but intensifies competition. So the industries that are the least competitive in the absence of a platform are the most likely to lose once a platform enters the market. Conversely, firms are better-off with a platform if and only if the latent demand is sufficiently large.

In the latter case, in markets with growth potential, the platform may produce a Pareto improvement, collecting positive profits while rendering both firms and consumers better off. This suggests that, in such markets, the optimal cap regulation should be more lax, allowing for higher commissions.

To investigate this point, let us assume for simplicity that the planner is utilitarian. In order to extend the welfare objective of the previous subsection to markets with growth potential, we have to define the *net* gain obtained by latent consumers from considering all firms in the market (vis-à-vis the situation where they knew no firm). This requires determining their payoff in the absence of a

[12] As shows by Chen and Riordan (200z) and Gabaix et al (2016), profit margins may well increase as more firms compete for consumers. This occurs in our model whenever $\lambda(n)$ is not nonincreasing in n. While possible in theory, price-increasing competition is believed to be more the exception than the rule in most markets of interest.

platform, which is left unmodelled in our partial-equilibrium analysis. We take a conservative stance by assuming that this outside option is the same as that of all other consumers were the platform absent.[13] Under this assumption, we obtain the following welfare objective:

$$\tilde{W}(\bar{f}) \equiv \int_0^{f^r} \left\{\delta + \mathbb{E}\left[Z^{1:N}\right] + d_\emptyset[\hat{\sigma}]\lambda(N) + b - k\right\} d\phi(k) + (1 - \Phi(f^r)) \left(\delta + \mathbb{E}\left[Z^{1:\hat{n}}\right]\right).$$

The next proposition derives the welfare-maximizing cap for markets with growth potential.

Proposition 4 ***(optimal regulation: growth potential)*** *Suppose the planner is utilitarian. Also assume that consumer match values are independent across firms, and let their cdf G_1 have tail index zero.*[14] *Then the welfare-maximizing cap is*

$$\begin{aligned}\tilde{f} &= b + (1 - d_\emptyset[\hat{\sigma}]) \left(\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right]\right) + d_\emptyset[\hat{\sigma}]\lambda(N) \\ &\approx b + \lambda(N)\left(1 - d_\emptyset[\hat{\sigma}]\right)\left(\frac{1 - d_j[\hat{\sigma}]}{d_j[\hat{\sigma}]}\right), \qquad (6)\end{aligned}$$

where the approximation error vanishes as $\hat{n}$ and N grow large.

It is useful to compare equations (5) and (6). Perhaps surprisingly, holding constant the firms' potential demands in the absence of the platform, the welfare-maximizing cap decreases as we move from mature to growing markets. The reason is that, absent the platform, the same potential demand $d_j[\hat{\sigma}]$ implies a higher reach in growing than in mature markets. Equivalently, holding $d_j[\hat{\sigma}]$ fixed, the expansion produced by the platform on consumers consideration sets is higher in mature than in growing markets, which translates into higher caps.

Coming back to the numerical illustration of the last subsection, the welfare-maximing cap would decrease to... provided the latent demand comprises X% of consumers... Again, this cap is significantly below the average fee charged by platforms in most markets...

5 Other Remedies

5.1 Banning Price Parity

TEXT TO BE ADDED

Proposition 5 ***(banning price parity)*** *Absent price parity, in the unique symmetric equilibrium, the platform sets $f = b$. Accordingly, banning price parity is outcome-equivalent to capping the platform fee at an inefficiently low level (as $b < \bar{f}$).*

[13] Explain other possibilities.

[14] That G_1 has tail index zero means that

$$\lim_{x\to\infty} \frac{d}{dx}\left(\frac{1 - G_1(x)}{g_1(x)}\right) = 0.$$

This condition is satisfied by many distributions of interest, such as the normal, Gumbel, extreme value, exponential, gamma, and lognormal, among others.

5.2 Platform Competition

It is tempting to think that platform competition might alleviate market distortions, rendering commissions caps unneeded (or redundant). In this section, we extend our model to allow for multiple platforms, and show that, under natural assumptions, the same equilibrium fee under monopoly also prevails under competition.

To model competition, we consider two platforms, seen as perfect substitutes in the eyes of consumers. As in the baseline model, we assume that the platforms' offers are private, and that firms make simultaneous joining and pricing decisions. The timing of the model is summarized below:

1. Each platform $i \in \{a, b\}$ privately offers the fee f_j^i to each firm $j \in \mathcal{N}$,
2. Firms simultaneously set prices and decide whether to join both, either, or no platform,[15]
3. Each consumer decides which platform to patronize, and makes a purchasing decision among the firms in her consideration set.

Consumers choose a platform without observing neither commissions, the set of listed firms, or prices.[16] As in the monopoly case, platforms expand the market knowledge of consumers, who add to their consideration sets all firms present in the platform they choose. Accordingly, if a firm does not join a platform (say, a), then the only consumers, among the patrons of platform a, who consider that firm are those that already had the firm in their original consideration set.

We study two scenarios, depending on the width of the price parity clause practiced by platforms. Following the literature, we say that *wide* price parity prevails when joining a platform implies that the firm cannot charge a lower price anywhere else. As such, if a firm joins both platforms, it is obliged to set the same price on both platforms, as well as for direct sales. In turn, we say that *narrow* price parity prevails when joining a platform implies that the firm cannot charge a lower price on the direct-sales channel, but is not constrained on the price charged on the other platform (in case the firm is also listed there).

For either scenario, we look for a symmetric perfect bayesian equilibrium.[17] Because we are interested in the impact of competition on equilibrium outcomes, we focus on the symmetric equilibria, therefore discarding all equilibria where the market tips (leading to a de facto monopolistic platform). The next proposition describes our main finding.

[15] As discussed in Section 2, assuming prices are chosen simultaneously to joining decisions simplifies the analysis, but does not affect qualitatively the results.

[16] This is a realistic assumption...

[17] We again assume that out-of-equilibrium beliefs are passive (in that an out-of equilibrium offer by a platform does not affect a firm's beliefs about the offers received by other firms).

Proposition 6 ***(platform competition)*** *Under either narrow or wide price parity, there exists a unique symmetric equilibrium. In this equilibrium, platforms offer all firms the monopolistic fee f^* of equation (2), half of all consumers patronize each platform, and firms join both.*

To understand the result above, let us analyze first the case of a wide price parity. Note that condition (2) implies that, upon being offered the fee f^* in equilibrium, firms are indifferent between joining both platforms and joining no platform.

Consider a platform that deviates and offers to a firm a fee higher than the equilibrium fee. The firm has obviously the option to stay on both platforms and adjust prices upwards, reflecting its unexpectedly high marginal cost. This option is however dominated by quitting both platforms altogether, as, absent this deviation, the firm would be indifferent between joining both platforms or none.

Another option is to quit the deviating platform (call it a), and remain on the non-deviating one (which is b). The price parity clause with platform b however prevents the firm from decreasing prices in its direct-sales channel without doing the same in platform b. Having a lower direct-sales price is desirable for the firm, as avoiding platform a's commission decreases the marginal cost of serving those patrons of platform a who have the firm in their original consideration set. As a consequence of this pricing constraint, the firm prefers quitting both platforms rather than quitting only the deviating one.

Finally, no platform gains by deviating to a fee lower than the equilibrium fee. The reason is that such a deviation does not result in more sales, as consumers have no way to detect the discounted fee. As consequence, lowering fees can only reduce revenues. Because platforms can do no better than offering f^* to each firm, all firms decide to multi-home, and consumers remain indifferent between visiting either platform (and see no gain in visiting both, which is informationally redundant).

The reasoning above applies to the case where the price parity clause is wide. Perhaps surprisingly, the same equilibrium fee is obtained if price parity is narrow. To see why, consider again the case of platform a offering some firm a fee higher than f^*. Unlike in the case of wide price parity, the firm can now increase the price in the deviating platform (and in its direct-sales channel) without raising the price in the non-deviating platform. While valuable for the firm, this extra degree of flexibility is not enough to prevent the firm from quitting both platforms. The reason is that, at the fee f^*, the firm is indifferent between joining both platforms and none. After being offered a higher fee at platform a, joining no platform becomes strictly preferable to remaining on both.

What about delisting from the deviating platform a, but remaining on b? This is again dominated by leaving both platforms. The reason is that, upon leaving platform a, the direct-sales price and the price at platform b are tied exactly as in the case of a broad price parity clause. This again renders discounting direct sales too costly for the firm, which then prefers leaving both platforms.

Finally, decreasing the fee away from f^* is again unprofitable for platforms, as consumer traffic would not react. All in all, moving from wide to narrow price parity fails to foster competition among platforms, and does not reduce equiliubrium fees. Indeed, the arguments above echo the reaction by the German competition authority vis-à-vis Booking.com's proposal to revise its price parity clause.

The stark conclusions of Proposition 6 rely on the inability of platforms to directly signal lower commissions to consumers, and on the assumption that.... This is contrast to Wang and Wright (2019), who propose a model that..., and show that moving from wide to narrow price parity is welfare-increasing. This is also in contrast to Vergé et al (2018), who, by assuming that..., show that even a wide price parity clause can be welfare-enhancing. Of course, it is an empirical question to determine which set of assumptions better describes the market. Our analysis however indicates that there may be good reasons for doubting the effectiveness of competition in these markets, even under milder versions of price parity.

6 Discussion and Conclusions

TO BE ADDED

7 Appendix: Omitted Proofs

Proof of Lemma 1. Consider profiles v and p such that $v_k - p_k$ is invariant to $k \geq 2$, which we denote by μ. Because G is symmetric,

$$D_1(p_1, p_{-1}; \sigma) = d_j[\sigma]\text{Prob}_G[z_2 + (v_2 - p_2) - (v_1 - p_1) \leq z_1 | z_2 \geq \max\{z_2, \ldots, z_n\}],$$

or, equivalently,

$$D_1(p_1, p_{-1}; \sigma) = d_j[\sigma]\left(1 - H^{(n)}(\mu - (v_1 - p_1))\right).$$

Hence

$$\frac{\partial D_1}{\partial p_1}(p_1, p_{-1}; \sigma) = -d_j[\sigma]h^{(n)}(\mu - (v_1 - p_1)).$$

The best response of firm 1 to p_{-1} then solves

$$p_1 = c_1 - \frac{D_1(p_1, p_{-1}; \sigma)}{\frac{\partial D_1}{\partial p_1}(p_1, p_{-1}; \sigma)} = c_1 + \frac{1 - H^{(n)}(\mu - (v_1 - p_1))}{h^{(n)}(\mu - (v_1 - p_1))}.$$

Swapping indexes we obtain the best reply of any other firm.

Any symmetric equilibrium has then to satisfy

$$p_j = c_j + \frac{1 - H^{(n)}(0)}{h^{(n)}(0)} = c_j + \lambda(n),$$

as in the statement of the lemma. This shows that at most one symmetric equilibrium exists.

Existence is assured by the fact that

$$v_j - p_j^*(\sigma, c_j) = v_j - c_j - \lambda(n) = \delta - \lambda(n).$$

Q.E.D.

Proof of Proposition 1. Consider the putative equilibrium fee profile $(f^*, \ldots, f^*)$. To shorten notation, let p_k^* denote $p_k^*(\bar{\sigma}, c_1 + f^* - b)$. If all firms join the platform and pay f^* per transaction, the individual profit of firm 1 (as well as of any other firm) is

$$\frac{1}{N}(p_1^* - c_1 - f^* + b) = \frac{1}{N}\lambda(N), \tag{7}$$

where we employ Lemma 1 to obtain the equality in (7). If firm 1 rejects and decides to do without the platform, it obtains

$$\max_{p_1} \left\{d_j[\underline{\sigma}]\left(1 - H^{(N)}((v_k - p_k^* - (v_1 - p_1))\right)(p_1 - c_1)\right\},$$

where the quantity $v_k - p_k^*((\bar{\sigma}, c_k))$ is invariant in k (by definition of a symmetric equilibrium). Notice that the firms other than 1 do not update their prices after 1 refuses to join, as this decision is simultaneous to price setting. The expression above can be rewritten as

$$\max_{p_1} \left\{d_j[\underline{\sigma}]\left(1 - H^{(N)}((v_k - c_k - f^* - \lambda(N) + b) - (v_1 - p_1))\right)(p_1 - c_1)\right\}$$

$$= \max_{p_1} \left\{ d_j[\underline{\sigma}] \left(1 - H^{(N)}(-f^* - \lambda(N) + b - c_1 + p_1)\right) (p_1 - c_1) \right\}$$

$$= \max_{p_1} \left\{ d_j[\underline{\sigma}] \left(1 - H^{(N)}(p_1 - p_1^*)\right) (p_1 - c_1) \right\}, \tag{8}$$

$$= d_j[\underline{\sigma}] \cdot \max_{\Delta p} \left\{ \left(1 - H^{(N)}(\Delta p)\right) (\Delta p + f^* + \lambda(N) - b) \right\}$$

where we employ Lemma 1 repeatedly and use the fact that $v_k - c_k = \delta$ for all k. Therefore, firm 1 joins the platform if and only if (7) is weakly greater than (8).

If the inequality is strict, the platform can deviate from equilibrium and offer to firm 1 (as well as to any other firm) a fee $f^* + \varepsilon$, $\varepsilon > 0$, such that (7) remains greater than (8). Because, by the envelope theorem, (8) is decreasing in f^*, and (7) is invariant to f^*, we conclude that (7) has to equal (8) in equilibrium. This symmetric equilibrium with full participation is unique because the equality (2) admits a unique solution. Q.E.D.

Proof of Corollary 1. Note that the right-hand side of (2) is increasing in the potential demand $d_j[\underline{\sigma}]$. Moreover, by the envelope theorem, it is also increasing in f^*. Therefore, as $d_j[\underline{\sigma}]$ increases, the equilibrium fee f^* goes down, which establishes the claim. Q.E.D.

Proof of Proposition 2. Follows from differentiating the welfare objective, and noting that the objective is quasi-concave in f^r. Q.E.D.

Proof of Proposition 3. Let as $\hat{n}$ and N grow large while satisfying $d_j[\hat{\sigma}] = \frac{\hat{n}}{N}$, and denote

$$a^{\hat{n}} \equiv \mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] \qquad \text{and} \qquad b^{\hat{n}} \equiv \left((\bar{G}_1)^{-1}\left(\frac{1}{N}\right) - (\bar{G}_1)^{-1}\left(\frac{1}{\hat{n}}\right) \right).$$

Consider first the case where G_1 has unbounded support. Recall that

$$\mathbb{E}\left[Z^{1:n}\right] \underset{n}{\sim} (\bar{G}_1)^{-1}\left(\frac{1}{n}\right)\Gamma(1-\gamma)$$

by Theorem 3 of Gabaix et al (2018).

We want to argue that $a^{\hat{n}} \underset{\hat{n}}{\sim} b^{\hat{n}}$, that is,

$$\lim_{\hat{n}\to\infty} \frac{a^{\hat{n}}}{b^{\hat{n}}} = 1.$$

To this end, note that

$$\left|\frac{a^{\hat{n}}}{b^{\hat{n}}} - 1\right| = \left|\frac{a^{\hat{n}} - b^{\hat{n}}}{b^{\hat{n}}}\right| = \left|\frac{\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] - \Gamma(1-\gamma)\left[(\bar{G}_1)^{-1}\left(\frac{1}{N}\right) - (\bar{G}_1)^{-1}\left(\frac{1}{\hat{n}}\right)\right]}{\Gamma(1-\gamma)\left[(\bar{G}_1)^{-1}\left(\frac{1}{N}\right) - (\bar{G}_1)^{-1}\left(\frac{1}{\hat{n}}\right)\right]}\right|$$

$$= \left|\frac{\left(\frac{\mathbb{E}\left[Z^{1:N}\right]}{(\bar{G}_1)^{-1}\left(\frac{1}{N}\right)\Gamma(1-\gamma)} - 1\right)(\bar{G}_1)^{-1}\left(\frac{1}{N}\right)\Gamma(1-\gamma) - \left(\frac{\mathbb{E}\left[Z^{1:\hat{n}}\right]}{(\bar{G}_1)^{-1}\left(\frac{1}{\hat{n}}\right)\Gamma(1-\gamma)} - 1\right)(\bar{G}_1)^{-1}\left(\frac{1}{\hat{n}}\right)\Gamma(1-\gamma)}{\Gamma(1-\gamma)\left[(\bar{G}_1)^{-1}\left(\frac{1}{N}\right) - (\bar{G}_1)^{-1}\left(\frac{1}{\hat{n}}\right)\right]}\right|$$

$$= \left| \frac{\left(\frac{\mathbb{E}\left[Z^{1:N}\right]}{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\Gamma(1-\gamma)} - 1 \right) - \left(\frac{\mathbb{E}\left[Z^{1:\hat{n}}\right]}{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)\Gamma(1-\gamma)} - 1 \right) \frac{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)}{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)}}{\left[1 - \frac{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)}{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)} \right]} \right|,$$

which numerator we call $c^{\hat{n}}$, and denominator $d^{\hat{n}}$.

Because the tail index of G_1 is γ and it has unbounded support, we know that

$$\frac{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)}{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)} = \frac{\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)}{\left(\bar{G}_1\right)^{-1}\left(\frac{d_j[\hat{\sigma}]}{\hat{n}}\right)} \to \left(d_j[\hat{\sigma}]\right)^{\gamma}.$$

Therefore, $c^{\hat{n}} \to 0$ and $d^{\hat{n}} \to 1 - (d_j[\hat{\sigma}])^{\gamma}$ as $\hat{n} \to \infty$. Consequently, we obtain that $a^{\hat{n}} \underset{\hat{n}}{\sim} b^{\hat{n}}$.

Moreover, note that

$$b^{\hat{n}} = \left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right) - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right) \right) = \frac{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} + O\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)^2,$$

which implies that

$$\frac{b^{\hat{n}} g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)} = 1 + g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right) \frac{O\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)^2}{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)} \to 1$$

as $\hat{n} \to \infty$, since $g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)$ is bounded from above and

$$\lim_{\hat{n}\to\infty} \frac{O\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)^2}{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)} = \lim_{\hat{n}\to\infty} \frac{\left(\frac{1}{\hat{n}}\right)^2 O\left(1 - \frac{\hat{n}}{N}\right)^2}{\left(\frac{1}{\hat{n}}\right)\left(1 - \frac{\hat{n}}{N}\right)} = \lim_{\hat{n}\to\infty} \frac{1}{\hat{n}} \left(1 - d_j[\hat{\sigma}]\right) O(1) \to 0.$$

Therefore,

$$b^{\hat{n}} \underset{\hat{n}}{\sim} \frac{\left(\frac{1}{\hat{n}(N)} - \frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)},$$

which implies that

$$a^{\hat{n}} \underset{\hat{n}}{\sim} \frac{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} = \frac{\left(\frac{1}{d_j[\hat{\sigma}]} - 1\right)}{N g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} \underset{\hat{n}}{\sim} \left(\frac{1}{d_j[\hat{\sigma}]} - 1\right) \Gamma(\gamma+2)\lambda(N),$$

where $N = \frac{\hat{n}}{d_j[\hat{\sigma}]}$.

Consider now the case where G_1 has bounded support with upper limit $\bar{z}$. Recall that

$$\bar{z} - \mathbb{E}\left[Z^{1:n}\right] \underset{n}{\sim} \left(\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{n}\right)\right) \Gamma(1-\gamma)$$

by Theorem 3 of Gabaix et al (2018).

We want to argue that $a^{\hat{n}} \underset{\hat{n}}{\sim} b^{\hat{n}}$, that is,

$$\lim_{\hat{n}\to\infty} \frac{a^{\hat{n}}}{b^{\hat{n}}} = 1.$$

To this end, note that

$$\left|\frac{a^{\hat{n}}}{b^{\hat{n}}}-1\right| = \left|\frac{\bar{z}-\mathbb{E}\left[Z^{1:\hat{n}}\right]-\left(\bar{z}-\mathbb{E}\left[Z^{1:N}\right]\right)-\Gamma(1-\gamma)\left[\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)-\left(\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)\right]}{\Gamma(1-\gamma)\left[\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)-\left(\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)\right]}\right|$$

$$= \left|\frac{\left(\frac{\bar{z}-\mathbb{E}\left[Z^{1:\hat{n}}\right]}{\Gamma(1-\gamma)\left(\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)\right)}-1\right)-\left(\frac{\bar{z}-\mathbb{E}\left[Z^{1:N}\right]}{\Gamma(1-\gamma)\left(\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}-1\right)\frac{\left(\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}{\left(\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)\right)}}{\left[1-\left(\frac{\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)}{\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)}\right)\right]}\right|$$

which numerator we call $c^{\hat{n}}$, and denominator $d^{\hat{n}}$.

Because the tail index of G_1 is γ and it has bounded support, we know that

$$\frac{\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)}{\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)} = \frac{\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)}{\bar{z}-\left(\bar{G}_1\right)^{-1}\left(\frac{d_j[\hat{\sigma}]}{\hat{n}}\right)} \to \left(d_j[\hat{\sigma}]\right)^{\gamma}.$$

Therefore, $c^{\hat{n}} \to 0$ and $d^{\hat{n}} \to 1-(d_j[\hat{\sigma}])^{\gamma}$ as $\hat{n}\to\infty$. Consequently, we obtain that $a^{\hat{n}} \underset{\hat{n}}{\sim} b^{\hat{n}}$.

By the exact same arguments in the unbounded case, one can show that

$$b^{\hat{n}} \underset{\hat{n}}{\sim} \frac{\left(\frac{1}{\hat{n}(N)}-\frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)},$$

which implies that

$$a^{\hat{n}} \underset{\hat{n}}{\sim} \frac{\left(\frac{1}{\hat{n}}-\frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} = \frac{\left(\frac{1}{d_j[\hat{\sigma}]}-1\right)}{Ng_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} \underset{\hat{n}}{\sim} \left(\frac{1}{d_j[\hat{\sigma}]}-1\right)\Gamma(\gamma+2)\lambda(N).$$

Q.E.D.

Proof of Proposition 4. The proof is very similar to that of Proposition 3. TO BE ADDED. Q.E.D.

Proof of Corollary 2. In the no-platform benchmark, the equilibrium profit of each firm is

$$d_j[\underline{\sigma}]\frac{\lambda(\underline{n})}{\underline{n}},$$

whereas, with a monopolistic platform, the equilibrium profit is

$$\frac{\lambda(N)}{N}.$$

Therefore, firms are better-off with platform if and only if

$$\frac{\lambda(N)}{N} > d_j[\underline{\sigma}]\frac{\lambda(\underline{n})}{\underline{n}} \quad \Longleftrightarrow \quad \frac{\lambda(N)}{\lambda(\underline{n})} > d_j[\underline{\sigma}]\frac{N}{\underline{n}} = 1-d_{\emptyset}[\underline{\sigma}],$$

as stated in the lemma. Q.E.D.

Proof of Proposition 5. First, note that, as firms can charge different prices inside and outside of the platform, it is a dominant strategy to join, irrespective of the fee f_j offered by the platform.

Suppose firms set prices $p_j = c_j + \lambda(N)$ for transactions outside the platform, and $\hat{p}_j = c_j + \lambda(N) + f_j - b$ for transactions inside the platform. In light of these prices, consumer buy inside the platform if and only if $f \leq b$. As a result, the platform is constrained to set $f_j \leq b$.

That the pricing rule above is an equilibrium follows directly from Lemma 1, considering firms' marginal costs to be $c_j + f_j - b$ rather than c_j. Q.E.D.

Proof of Proposition 6. Consider first the case of wide price parity. We will first argue that

$$\frac{\lambda(N)+b}{N} > \max_{p_1} \left\{ \frac{d_j[\underline{\sigma}]}{2} \left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1) + \frac{1}{2}\left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1 - f^* + b) \right\}, \tag{9}$$

where f^* is given by equation (2) and p_1^* is a short hand for $p_1^*((\bar{\sigma}, c_1))$. Notice that, by construction of f^*, in the putative equilibrium, firms are indifferent between joining both platforms or neither platform. The inequality in (9) then implies that, in the putative equilibrium, firms strictly prefer joining neither platform than joining one but not the other.

To prove (9), suppose, to obtain a contradiction, that

$$\frac{\lambda(N)+b}{N} \leq \max_{\Delta p} \left\{ \frac{d_j[\underline{\sigma}]}{2} \left(1 - H^{(N)}(\Delta p)\right)(\Delta p + p_1^* - c_1) + \frac{1}{2}\left(1 - H^{(N)}(\Delta p)\right)(\Delta p + p_1^* - c_1 - f^* + b) \right\},$$

where $\Delta p \equiv p_1 - p_1^*$. Employing Lemma 1, we can rewrite this inequality as

$$\frac{\lambda(N)+b}{N} \leq \max_{\Delta p} \left\{ \frac{d_j[\underline{\sigma}]}{2} \left(1 - H^{(N)}(\Delta p)\right)(\Delta p + f^* + \lambda(N)) + \frac{1}{2}\left(1 - H^{(N)}(\Delta p)\right)(\Delta p + \lambda(N) + b) \right\}$$

$$< \max_{\Delta p} \left\{ \frac{d_j[\underline{\sigma}]}{2} \left(1 - H^{(N)}(\Delta p)\right)(\Delta p + f^* + \lambda(N)) \right\} + \max_{\Delta p \leq 0} \left\{ \frac{1}{2}\left(1 - H^{(N)}(\Delta p)\right)(\Delta p + \lambda(N) + b) \right\}$$

$$= \max_{\Delta p} \left\{ \frac{d_j[\underline{\sigma}]}{2} \left(1 - H^{(N)}(\Delta p)\right)(\Delta p + f^* + \lambda(N)) \right\} + \frac{\lambda(N)+b}{2N}.$$

Therefore,

$$\frac{\lambda(N)+b}{N} < \max_{\Delta p} \left\{ d_j[\underline{\sigma}] \left(1 - H^{(N)}(\Delta p)\right)(\Delta p + f^* + \lambda(N)) \right\}, \tag{10}$$

which contradicts the definition of f^*.

We now argue, if a platform deviates and offers a firm some fee $\hat{f} > f^*$, then the firm will delist from both platforms. To see why, note that

$$\frac{\lambda(N)+b}{N} > \max_{p_1} \left\{ \frac{d_j[\underline{\sigma}]}{2} \left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1) + \frac{1}{2}\left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1 - \hat{f} + b) \right\},$$

which follows from (9) and the envelope theorem. As a result, no platform can individually set a fee above f^*. Setting a fee below this level is obviously suboptimal. Therefore, both platform have no profitable deviation from the putative equilibrium.

Consider now the case of narrow price parity. If a firm joins any platform, this weaker form of price parity prevents the firm from setting the direct-sales price smaller than any platform price. This implies

$$\max_{p_1} \left\{ \frac{d_j[\underline{\sigma}]}{2} \left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1) + \frac{1}{2}\left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1 - f^* + b) \right\}$$

is the maximum profit a firm can obtain by patronizing a single platform in the putative equilibrium. The same arguments above then lead to (9), which implies that, as under broad price parity, firms strictly prefer joining neither platform than joining one but not the other in the putative equilibrium. It then follows that if a platform offers a firm some $\hat{f} > f^*$, the firm will optimally decide to delist from both platforms. This implies that platforms can do no better than offering f^*, concluding the proof. Q.E.D.

References

[1] Anderson, S., A. de Palma, and J.-F. Thisse, 1992. Discrete Choice Theory of Product Differentiation, MIT Press.

[2] Besanko, D., and T. P. Lyon, 1993. "Equilibrium incentives for most-favored customer clauses in an oligopolistic industry." International Journal of Industrial Organization, 11(3), 347-367.

[3] Boik, A., and K. S. Corts, 2016. "The effects of platform most-favored-nation clauses on competition and entry." Journal of Law and Economics, 59(1), 105-134.

[4] Cazaubiel, A., Cure, M., Johansen, B. O., and T. Vergé, T., 2018. "Substitution between online distribution channels: An empirical evaluation on the scandinavian hotel industry." Working Papers in Economics 08/18, University of Bergen.

[5] Edelman, B., and J. Wright, 2015. "Price coherence and excessive intermediation." Quarterly Journal of Economics, 130(3), 1283-1328.

[6] Gabaix X., D. Laibson, D. Li, H. Li, S. Resnick, and C. de Vries, 2016. "The Impact of Competition on Prices with Numerous Firms." Journal of Economic Theory, 165, 1-24.

[7] Hunold, M., Kesler, R., Laitenberger, U., and F. Schlütter, F., 2018. "Evaluation of best price clauses in hotel booking." International Journal of Industrial Organization, 61, 542-571.

[8] Johansen, B. O., T. and Vergé, 2017. "Platform price parity clauses with direct sales." Working Papers in Economics 01/17, University of Bergen.

[9] Johnson, J. P., 2017. "The agency model and MFN clauses." Review of Economic Studies, 84(3), 1151-1185.

[10] Mantovani, A., Piga, C., and C. Reggiani, 2019. "Much Ado About Nothing? Online Platform Price Parity Clauses and the EU Booking.com Case". Available at SSRN: https://ssrn.com/abstract=3381299

[11] Perloff, J. M., and S. C. Salop, 1985. "Equilibrium with Product Differentiation." Review of Economic Studies, 52(1), 107-120.

[12] Schnitzer, M., 1994. "Dynamic duopoly with best-price clauses." RAND Journal of Economics, 25(1), 186-196.

[13] Segal, I., 1999. "Contracting with Externalities." Quarterly Journal of Economics, 114(2), 337-388.

[14] Varian, H., 1980. "A Model of Sales." American Economic Review, 70, 651-659.

[15] Wang, C., and J. Wright, 2019. "Search platforms: showrooming and price parity clauses." RAND Journal of Economics, *forthcoming*.