DISCUSSION PAPER SERIES

DP15048

REGULATING PLATFORM FEES UNDER PRICE PARITY

Renato Gomes and Andrea Mantovani

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

REGULATING PLATFORM FEES UNDER PRICE PARITY

Renato Gomes and Andrea Mantovani

Discussion Paper DP15048
Published 16 July 2020
Submitted 13 July 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

REGULATING PLATFORM FEES UNDER PRICE PARITY

Abstract

Online marketplaces, such as Amazon, 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. To prevent show-rooming, 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

Renato Gomes - renato.gomes@tse-fr.eu
Toulouse School of Economics and CEPR

Andrea Mantovani - a.mantovani@unibo.it
University of Bologna

Acknowledgements
For very helpful comments, we thank, without implicating, Patrick Rey, Julian Wright, Doh-Shin Jeon, Tat-How Teh, as well as participants at the TSE Competition Policy Workshop, the Virtual Market Design Seminar (Bonn, Mannheim and Karlsruheand), and the Online Economics of Platforms Seminar (TSE, BU and NUS). Renato Gomes gratefully acknowledges financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 759733 - PLATFORM).

Regulating Platform Fees under Price Parity*

Renato Gomes† **Andrea Mantovani**‡

First Version: 17/02/2020
This Version: 28/05/2020

Abstract

Online marketplaces, such as Amazon, 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. To prevent show-rooming, 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.

*For very helpful comments, we thank, without implicating, Patrick Rey, Julian Wright, Doh-Shin Jeon, Tat-How Teh, as well as participants at the TSE Competition Policy Workshop, the Virtual Market Design Seminar (Bonn, Mannheim and Karlsruheand), and the Online Economics of Platforms Seminar (TSE, BU and NUS). Renato Gomes gratefully acknowledges financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 759733 - PLATFORM).

†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

Platforms in the modern economy play an important role in providing consumers with information about potential trade opportunities. For instance, online travel agencies (OTA's), such as Booking.com or Expedia, are an essential tool for travelers to discover hotel offers. Marketplaces, such as Amazon, inform buyers about existing sellers, allowing for trades that would otherwise be elusive. A similar informational role is played by ride-sharing platforms (such as BlaBlaCar), food delivery apps (such as DoorDash or Uber Eats), and many other matching platforms (such as those for booking restaurants, hiring babysitters for the evening, or finding a caregiver).

Most of these 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. These commissions are often substantial: For instance, on Amazon Marketplace, professional sellers pay on average 13% per sale, whereas in Booking.com the average fee is 20%.

A crucial challenge for this business model is the possibility of *show-rooming*, whereby consumers use the platform to find their preferred seller, but then switch to the direct channel to obtain a discount (made possible by the fact that the seller then avoids the platform's commission). To prevent this practice, many platforms adopted price parity clauses, which restrict the sellers' ability to charge lower prices on alternative sales channels.[1] These covenants 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 their business, as it curbs opportunistic behavior.[2] By contrast, competition authorities and consumer associations often regard price parity as the source (or, at best, a reinforcer) of platforms' market power. In line with these concerns, European competition authorities have reached important decisions on price parity clauses over the past years. Currently, all types of price parities are forbidden in France, Italy, Belgium, and Austria, whereas in other countries they are prohibited only for certain OTA's (Amazon, HRS and Booking.com in Germany; Booking.com in Sweden; Amazon in the UK). In the US, Amazon recently decided to remove these clauses from its contracts with third-party marketplace sellers.[3]

Yet, it is not entirely clear that a ban, or voluntary removal, of price parity clauses actually produces tangible results.[4] For one, sellers might still practice price parity, fearing that the platform

[1] These clauses also prohibit sellers from offering more availability or better conditions on alternative sales channels.

[2] In line with this concern, Hunold et al. (2018) used metasearch data to show that hotels charged the lowest price on the direct channel more often in Germany than in countries that did not abolish price parity clauses.

[3] 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.

[4] Mantovani et al. (2020) collected listed prices on Booking.com in the period 2014-2017 for touristic areas in France, Italy, and Spain. They compare prices before and after price parity was lifted in France, using as control similar tourist destinations in Italy and Spain. Their study finds a limited response of hotel prices on Booking.com, both in the short

might "down-list" them otherwise.[5] 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.[6]

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 restricting price parity; namely, we study how to optimally cap platforms' commissions.[7] Regulation of this kind has been recently enacted for delivery apps during the Covid-19 crisis.[8] A theoretical framework guiding regulation has however been missing.

Model and Results

In our baseline model, we consider a monopolist platform that 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, augmenting each seller's potential demand (which is the set of consumers who consider the seller when making a purchasing decision). To capture this idea concisely, we assume that a firm listed in the platform becomes part of the consideration set of all consumers in the market. Conversely, if a firm does not join the platform, its potential demand consists of a much smaller set of consumers: those who know the firm 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 firms' 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 pay 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.

and medium run, in line with evidence by the European Competition Network (2017).

[5]Hunold et al. (2020) show that OTAs penalize hotels that charge lower prices elsewhere with worse rankings, a practice euphemistically called 'dimming' .

[6]Other impediments to *de facto* banning price parity are the fact that hotels have scarce propensity to price differentiate, and sometimes exhibit limited awareness of the policy changes. See Mantovani et al. (2020) for further discussion of these points.

[7]A parallel debate unfolded in the realm of payment cards, where the so-called "no-surcharge rule" prevents merchants from discriminating prices based on the payment method. Fearing this rule generates too much market power to payment networks, some countries lifted it (UK, Netherlands, New Zealand, Australia, etc), whereas others upheld it while regulating fees (US, the European Commission, Brazil, etc). These policies are seen as substitutes, as articulated by the European Directive (2015).

[8]Although these platforms do not explicitly impose price parity, it is impractical for restaurants to differentiate prices on whether the order is take-out or delivery. See for instance https://www.cnbc.com/2020/05/15/city-lawmakers-provide-restaurants-with-relief-from-delivery-fees.html

Crucially, the platform's market power stems from a contractual externality (Segal 1999) that listed firms impose on the non-listed ones. Namely, if a firm decides not to join the platform, it faces a world where all consumers who consider that firm *also* consider *all* competing firms listed on the platform. This makes the non-listed firm face a degree of competition (among its 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. We first consider mature markets, for which the platform expands the consideration sets of existing customers, but does not increase aggregate sales. If the regulator is utilitarian (assigning equal weights to the consumers' surplus, firms' profits, and the platform's profits), the optimal cap takes a familiar form: It equals the expected externality that the platform imposes on other market participants, i.e., the sum of convenience and informational benefits. This Pigouvian rule has similarities to the "tourist test" regulation adopted in the payment industry;[9] 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 knowledge of consumers' (distribution of) match values. To circumvent this difficulty, we apply techniques from extreme-value theory (first developed by Gabaix et al. 2016) 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' profit margin 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 an application to online travel agencies. In a plausible scenario,[10] we find that current average fees, which equal 20%, increase utilitarian welfare if and only if the platform at least doubles consumers' consideration sets (relative to alternative ways of gathering information online). In light of other sources of information easily available on the Internet (e.g., Google), this suggests that regulation capping commissions might bind in some markets, if optimally set. If the welfare measure does

[9] The "tourist" or "avoided-cost" test requires the merchant fee to not 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 Rochet and Tirole (2011).

[10] Namely, assuming firm's profit margins average 20% and that the convenience benefit of a transaction through the platform is 2%.

not include platforms' profits, but only producer and consumer surplus, the optimal cap is tighter. Namely, it is *half* of the utilitarian one (10% in the scenario alluded above) provided the platform's operating costs are uniformly distributed.

We then study markets with growth potential, for which the platform brings in new consumers, increasing aggregate sales. In this case, firms may gain from the presence of the platform, as increased competition for each consumer is accompanied by more sales for all firms. Moreover, the industries that are the most competitive in the absence of a platform are the most likely to gain once a platform enters the market, as demand expansion trumps the increase in competition. This is in contrast to mature markets, where the platform raises competition and marginal costs, but does not expand demand, rendering firms necessarily worse-off.

Perhaps counter-intuitively, holding constant the size of firms' potential demand in the absence of the platform, the optimal cap decreases as we move from mature to growing markets. The reason is that, for the same potential demand, the expansion of consumers' consideration sets (produced by the platform) is lower in growing markets, which explains the tighter cap. For instance, if the platform doubles aggregate sales and potential demands, it should be allowed to charge a commission no higher than 15% (under the utilitarian welfare criterium).[11]

Finally, we evaluate two alternatives to cap regulation: an outright ban of price parity, and a relaxation of the latter under competition among platforms. If consumers can seamlessly switch to the sales channel with the lowest price, banning price parity is outcome-equivalent to capping the platform fee at the convenience benefit of a transaction. This cap is inefficiently low, be the market mature or growing, as it prevents the platform from appropriating (any of) the informational (ex-ante) benefits it generates. If the platform enjoys some market power in the absence of price parity, the "equivalent" cap exceeds the convenience benefit, but in general differs from its welfare-maximizing level.

We also find that, as consumers single home in equilibrium, competition between platforms fails to reduce equilibrium fees. This holds true if price parity is practiced in either its wide or narrow forms.[12] The reason for this result is that, unless a firm delists from all platforms, price parity ties the firm's price at the direct-sales channel to the price charged at *some* platform where it is still listed. As a result, only two options are relevant for firms: delisting from all, or joining all platforms. The latter implies that platforms can sustain the monopolistic fee in equilibrium, rendering competition (even under narrow price parity) ineffective in curbing market power.[13]

[11] We assume the same scenario described in footnote 9.

[12] In the latter case, firms can charge a lower price in the competing platform, but not in the direct-sales channel. See subsection 5.2 for further discussion.

[13] Confirming this prediction, the European Competition Network (2017) documents that the change from wide to narrow price parity did not significantly increase price differentiation across hotels' sales channels, nor brought significant changes in the commission fees charged by OTAs.

Paper Outline

The next subsection reviews the pertinent literature. Section 2 describes the model. Section 3 studies the laissez-faire outcome, where the platform is not subject to regulation. Section 4 derives the optimal cap regulation in both mature and growing markets. Section 5 considers alternative remedies, such as (*de facto*) banning price parity, or relaxing it (to its narrow form) under platform competition. Section 6 concludes by collecting the empirical and policy implications of our analysis. All proofs are in the Appendix at the end of the document.

1.1 Related Literature

Price parity clauses, also known as "most-favored 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. 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, which may invest to provide a non-pecuniary benefit to consumers. They show that price parity raises the price of direct purchases, increasing demand for the intermediary's service. Financed by high commissions, the platform over-invests in the provision of the non-pecuniary benefit, which in turn renders firms more likely to join. This may lead to an increase in final prices, and a decrease in social welfare.

Boik and Corts (2016) and Johnson (2017) assume consumers must use one of two differentiated platforms and show that price parity restrictions typically increase platforms' fees, thereby raising the prices charged by sellers and ultimately damaging final consumers. In the former paper, demand is elastic and both the fees and prices are linear in quantities. Apart from raising fees and final prices, price parity prevents market entry by low-cost competitors. The latter paper compares instead the wholesale and the agency models in a framework with inelastic demand and revenue-sharing. With unrestricted prices, a shift from the wholesale to the agency model benefits platforms and consumers, but harms sellers, as retail prices tend to diminish. When price parities are imposed within the agency model, however, platform competition is softened, and fees tend to increase, driving up retail prices.

These articles do not explicitly model the role played by the platform in expanding consumers' consideration sets, nor study regulation. By contrast, the informational externalities produced by the platform are at the core of our theory of harm, and form the basis of optimal regulation.[14]

Ronayne and Taylor (2019) analyze a market where two sellers produce a homogenous good which

[14] For a more policy-oriented perspective, see Ezrachi (2015), who describes the benefits generated by price comparison websites (PCWs), but also warns about the anti-competitive consequences associated with the adoption of price parity.

can be sold both directly and through a competitive platform, which helps consumers to compare market prices, but charges sellers a fee per transaction being completed. They adopt a clearinghouse framework à la Varian (1980), and divide consumers into savvy and loyal shoppers, with the latter checking only one sales channel. They first compute equilibrium prices and fees and show that the platform size significantly affects market outcomes. In this context, they analyze the effect of price parity clauses, which raise commission fees and induce some firms to delist, thus harming consumers.[15] Wang and Wright (2020) consider a sequential search model in which 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.[16] Our paper differs on the way we model the informational role of platforms, which is instrumental to our focus on regulation.

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.[17] 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, firms become more prone to delisting, which reduces average costs.[18] 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 assess the welfare impact of price parity, in some cases lending support to a ban, or favoring a weakening of this clause (e.g., its narrow version - see footnote 12). By contrast, we study optimal regulation, which maintains price parity, but restricts the platform's ability to levy

[15]Calzada et al. (2019) consider a model in which both platforms and sellers are exogenously differentiated. They focus on the effect of price parity on the sellers' decision about whether to list on one or two OTAs. They obtain these price restrictions induce sellers to single-home, and to eventually cut their direct sales channel. This has an evident negative impact on consumer welfare, but not necessarily on the sellers' profits.

[16]Relatedly, Bisceglia et al. (2019) consider the possibility for two competing Global Distribution Systems (GDSs) to negotiate parity provisions with a monopolistic airline in order to obtain the same conditions than the rival. The two GDSs serve final consumers through travel agents. Consumers can also directly buy from the airline, thus avoiding the intermediation process. The authors compare the adoption of platform parity provisions in both the wholesale and the agency model, and find that the latter may result in higher benefit for both GDSs and consumers, whereas the former is preferred by travel agencies. However, these parity restrictions are not directly comparable to the price parity clauses that are the object of our investigation, as direct prices remain unconstrained.

[17]They also assume that commissions are secretly offered. Rey and Vergé (2019) extend the analysis of secret contracting within multilateral vertical relations, and find conditions under which price parity agreements do not necessarily raise retail prices at equilibrium.

[18]Cazaubiel *et al.* (2020) use an exhaustive database from one large chain of hotels in Oslo to evaluate the degree of substitution between online distribution channels, exploiting the chain's decision to delist from Expedia. Their results indicate that the chain's direct sales channel appears to be a credible alternative to the OTAs.

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, although typically away from the welfare-maximizing level.

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.[19] 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 conditions 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,

[19] 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 $\mathcal{B}[0, 1]$.

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 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 simplicity, 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.[20] The platform is profit-maximizing.

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.[21] As a result, if a firm joins the platform, all of its sales happen through the platform (as consumers prefer doing so).[22]

[20] For example, Booking.com proposes different fees on the basis of hotel characteristics such as size, location, logistics and distribution activities.

[21] We relax this assumption in Section 5.

[22] It is straight-forward to allow consumers to obtain a convenience benefit $\beta \geq 0$ for completing a transaction within the platform (rather in the direct-sales channel). For simplicity, we normalize $\beta = 0$ and interpret b, the firm's convenience benefit, as capturing all the convenience advantages produced by the platform.

The timing of the model is summarized below:

1. The platform privately offers the fee f_j for each firm $j \in \mathcal{N}$.

2. Firms simultaneously set prices and decide whether (or not) to join the platform,[23]

3. 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.

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.

[23] 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.

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

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

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. 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. To guarantee full market coverage, we assume that δ is sufficiently large.[24]

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 $j \in \mathcal{N}$,*

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

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 $p_j^* = c_j + f - b + \lambda(N)$.[25]

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 (2007), which generalizes Hotelling.

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.

[24] Namely, that $\delta \geq \max_n \lambda(n)$.

[25] Full market coverage requires that $\delta - (f - b) \geq \lambda(N)$, which we assume to hold in the relevant range.

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*

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

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 its 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

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

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 participating 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). An easy way to appreciate this point is to note, from (2), that f^* can be made arbitrarily large as the size of potential demands, $|d_j[\underline{\sigma}]|$, approaches zero.[26]

What explains the platform's ability to extract more rents from consumers and firms than it generates? Its source of market power lies in a *contractual externality* (Segal 1999) that listed

[26] To preserve full market coverage, we maintain the assumption that the stand-alone values v_j's are sufficiently large.

firms impose on non-listed ones. Namely, 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 obviously 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.[27]

As hinted above, the firms' pre-visit potential demand is a key determinant of the equilibrium fee. The next corollary indeed confirms that this quantity is a sufficient statistic for performing comparative statics with respect to the information profile σ.

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 the size of 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 Proposition 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).*

Remark 2 ***(two-part tariffs)*** *For simplicity, we assumed that the platform employs linear fees, but this is not essential for Proposition 1, nor for any of the results that follow. Indeed, the contractual externality captured by the equilibrium condition (2) is only stronger were the platform to employ a two-part tariff.*[28] *The latter alleviates the double-marginalization problem, enabling the platform to extract rents without raising (as much) the perceived marginal costs of participating firms.*

[27] Edelman and Wright (2015) identify an externality whereby consumers joining the platform make firms set higher prices, which hurts, under price coherence, those consumers who use the direct-sales channel. By contrast, the externality identified here pertains to listed firms rendering the potential demand of non-listed ones more elastic.

[28] Such fee structure is practiced, for instance, in Amazon Marketplace, where a flat fee is applied together with a sale percentage fee. The former depends on the type of seller (occasional sellers pay a flat $0.99 for each sales transaction, whereas pro merchants pay $39.99 per month), whereas the latter varies based on categories.

4 Cap Regulation

The optimal regulation balances the gains from lower commissions and the potential losses from having no platform to centralize market information. To perform this balancing act, one therefore has to assess the welfare level in case the platform refuses to operate (which occurs if its ability to extract rents from firms is too limited).[29] One main difficulty on this regard is that consumers' information acquisition behavior may change depending on the availability of an aggregator platform. Indeed, in a world where Expedia is available, 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 likely induces 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 operates).

To formally capture the platform's decision to operate or not in a regulated market, we introduce a 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 $\mathbb{R}_+$. Being profit-maximizing, the platform operates in a market if it expects profits to exceed the cost k.

4.1 Mature Market

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 inconsequential 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 revenue also equals f^r.

[29] This is the risk mostly emphasized by participants of industries on the verge of regulation. Uber's decision to quit Southeast Asia was partially motivated by the introduction of new regulations governing ride-hailing operations. Amazon was recently forced to remove a series of products from its website in India to comply with new rules protecting small retailers; as it previously committed to spending $5.5 billion on e-commerce in India, this may limit future investment opportunities by the platform in that country.

Our measure of social welfare combines the surplus derived by firms and consumers with that of the platform, which weight is $\alpha \geq 0$. To compute consumer surplus, we let $Z^{1:n}$ denote the first-order statistic out of $n \leq N$ coordinates of the random vector z, which describes consumers' profiles of match values (and is distributed according to the symmetric distribution G). Assuming the cap binds, the planner's objective is then

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

where the integral describes welfare when the platform operates (which occurs if its cost realization is low), while the second term describes welfare when the platform shuts down (which occurs if its cost realization is 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 informational frictions (that is, under reach N), in addition to the convenience benefit b, but discounted by the platform fee $\bar{f}$. In turn, the platform's profit is simply $\bar{f} - k$. When the platform stays inactive, the surplus of consumers and firms consists of the expected gains from trade under the counterfactual consideration profile $\hat{\sigma}$ (whose reach is $\hat{n} \leq N$), as described in the second term of $W(\bar{f})$.

The next proposition derives the welfare-maximizing commission cap.

Proposition 2 ***(optimal regulation: mature market)*** *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*

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

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

Consider first the utilitarian case, where $\alpha = 1$. Here, the third term in the left-hand of side (4) vanishes, and the cap equals the *surplus-neutral fee* $\bar{f}_1$, 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).

For α below 1, the planner gives more weight to consumers and firms, setting the cap below the surplus-neutral fee $\bar{f}_1$. This is optimal because having the platform inactive in some instances (in which it would operate were the planner utilitarian) is compensated by the increase in the surplus of consumers and firms that a tighter cap generates. Conversely, when α is above 1, the planner gives more weight to the platform vis-à-vis consumers and firms, setting the cap above the surplus-neutral fee $\bar{f}_1$. In this case, the presence of the platform may hurt the other market participants in some instances.

Remark 3 ***(welfare measure: excluding platform's profit)*** *If the platform spends its profit on activities of little social value (wasteful advertising), or if its dominant position is fortuitous (say, due to network externalities, not from technological innovation), the planner might be tempted to assign no weight to its profit (i.e., set $\alpha = 0$).*[30] *In this case, letting the quasi-fixed cost k be uniformly distributed on $[0, \bar{k}]$, the optimal cap is*

$$\bar{f}_0 = \frac{1}{2}\bar{f}_1.$$

Accordingly, cap regulation is more likely to bind when the planner cares exclusively about consumer surplus and firm profits.

4.2 Large Market

To get a better sense of how current platform fees compare to the welfare-maximizing cap, let us for simplicity focus on the utilitarian case, where the cap equals $\bar{f}_1$. 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 circumvent 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 approximation techniques introduced by Gabaix et al. (2016), based on extreme value theory.

Because the latent demand is null, it follows from the definition of symmetric information profiles that

$$\frac{\hat{n}}{N} = |d_j[\hat{\sigma}]|\,. \tag{5}$$

Accordingly, the size of potential demands equals the fraction of firms known by each consumer.[31] We construct the approximation by letting the market grow large (i.e., $\hat{n}, N \to \infty$) while holding $|d_j[\hat{\sigma}]|$ constant. This procedure allows us to express the cap as a function of firms' potential demands absent the platform, which is the crucial quantity determining the informational contribution generated by the platform. This size of potential demand is measurable through surveys,[32] or experiments.

[30] See Tirole (2011) for a discussion, in the context of the card payment application, of whether platform profits should be included into the measure of social welfare.

[31] Recall $\hat{\sigma}$ is a symmetric informational profile of reach $\hat{n}$. Denoting by $\hat{\sigma}_n$ the measure of consumers aware of each $\hat{n}$-sized subset of firms, it follows that $\hat{\sigma}_n = \frac{1}{C(N,\hat{n})}$, where $C(N, \hat{n})$ is the number of $\hat{n}$-combinations from a set of N elements. This implies that the potential demand of each firm is

$$|d_j[\hat{\sigma}]| = \hat{\sigma}_n C(N-1, \hat{n}-1) = \frac{C(N-1, \hat{n}-1)}{C(N, \hat{n})} = \frac{\hat{n}}{N}.$$

[32] Two early contributions employing micro data on consumer consideration sets are Clark et al (2009) and Draganska and Klapper (2011). The former paper measures brand awareness in 25 broad categories of goods, while the latter collects data on consumer information in the market for ground coffee. See also Moraga-Gonzalez et al. (2018) for cars, and Sengupta and Wiggins (2014) for airline tickets.

Applying extreme-value theory 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 ***(large market)*** *Assume that consumer match values are independent across firms, and let their cdf G_1 be well-behaved with tail index γ and support $(\underline{z}, \bar{z})$, where $0 < \underline{z} < \bar{z} \leq \infty$.*[33] *Letting $\Gamma(\cdot)$ be the gamma function, we obtain that*

$$\lim_{\hat{n},N\to\infty} \left\{ \frac{\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right]}{\lambda(N)} \right\} = \left(\frac{1 - |d_j[\hat{\sigma}]|}{|d_j[\hat{\sigma}]|} \right) \Gamma(\gamma + 2), \tag{6}$$

where the limit above is taken as $\hat{n}$ and N grow large while satisfying (5).

Equation (6) means that, in large markets, the consumers' benefit from enjoying larger consideration sets is proportional to the firms' profit margin. The proportionality factor is the relative expansion of consumers' consideration sets (or, equivalently, the expansion of firms' potential demands) multiplied by a constant reflecting the thickness the tail of the distribution of consumers' match values. Although instrumental for cap regulation (as described below), this result can be seen as an independent contribution of this paper.

The intuition is simple: the consumers' benefit from larger consideration sets is small when the highest match values across firms are close to each other. In this case, markups are small, as competition is intense for providing consumers with the highest surplus. As a result, the quantities in the numerator and denominator of the left-hand side of (6) should co-move; they are indeed proportional to each other when the market is large. Holding constant the consumers' benefit from enjoying larger consideration sets, the firms' profit margin decreases as $\frac{N}{\hat{n}} = |d_j[\hat{\sigma}]|^{-1}$ increases. The reason is that competition gets fiercer as the same consumers' benefit is associated with a larger expansion of consideration sets (which means that the highest realizations of match values are more likely to be close to each other).

Numerical Illustration. We can employ Proposition 3 to obtain an "easier-to-use" formula for the utilitarian cap. To do so, let us note that most distributions used in applied work have tail index $\gamma = 0$,[34] in which case $\Gamma(\gamma + 2) = 1$. We then adopt the approximation

$$\bar{f}_1 \approx b + \left(\frac{1 - |d_j[\hat{\sigma}]|}{|d_j[\hat{\sigma}]|} \right) \lambda(N). \tag{7}$$

[33] In extreme value theory, we say that the cdf G_1 is well behaved when its density is differentiable, bounded from above, $\lim_{x\to\infty} \frac{1-G_1(x)}{g_1(x)}$ exists in $\bar{\mathbb{R}}$, and $\gamma = \lim_{x\to\infty} \frac{d}{dx}\left(\frac{1-G_1(x)}{g_1(x)}\right)$ exists and is finite, which we call the *tail index.*

[34] Examples include the normal, log-normal, exponential, gamma, and Gumbel cdf's (used in logit demand models).

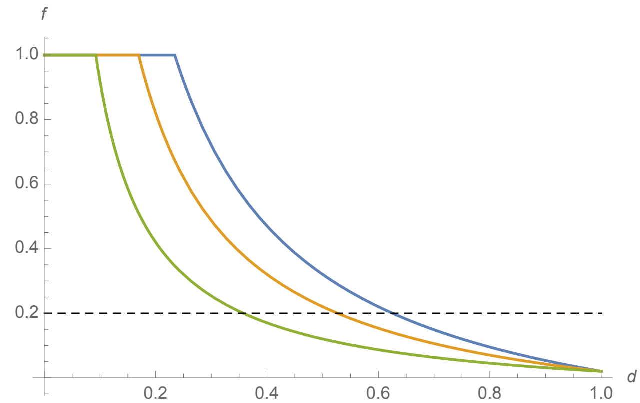

Figure 1: Utilitarian caps (as a percentage of the retail price) as a function of the relative size of hotels' potential demands in the absence of a platform. The blue curve sets the hotels' profit margin to 30%, the yellow to 20%, and the green to 10%.

Note that, because firms fully pass the platform's fee to consumers, their profit margin does not depend on the level of the fee, but only on the number of firms in the market, N.

We borrow from recent empirical work to numerically illustrate Proposition 3. To this end, we collected markup estimates in the hospitality industry, and found that in most markets it varies from 10% to 30% of the retail price.[35] The convenience benefit b (which captures the valued added by the booking and payment interface provided by the platform) should be commensurate to the market rates of online payment gateways (such as Paypal). So we set it to 2% of the retail price. Figure 1 illustrates the welfare-neutral fee as a function of the firms' potential demands.

Consider first the case where the profit margin of hotels is 20%. As illustrated by the yellow curve in the figure, the welfare-neutral fee is decreasing in the size of potential demands in the absence of a platform. The cap is irrelevant if potential demands are sufficiently small; namely, if $|d_j[\hat{\sigma}]| \leq 0.17$, the welfare-neutral cap equals 100% of the retail price. In other words: if the platform multiplies by at least 5 the number of consumers aware of each hotel, consumers are sufficiently better off to compensate for hotels' lost profits. By contrast, when the platform brings no informational benefit to consumers (that is, $|d_j[\hat{\sigma}]| = 1$), the welfare-neutral cap coincides with the convenience benefit: that is, it equals 2%. In this case, perhaps unsurprisingly, the optimal cap coincides with the tourist test regulation of payment cards, which advocates merchant fees should equal the convenience benefit of a card payment. In general, the regulation proposed here differs from the tourist test paradigm precisely for the expansion in the consumers' consideration sets produced by information platforms.

Turning to perhaps a more relevant range for regulation, the utilitarian cap is 20% (which is the average fee of Booking.com) when potential demands absent the platform are roughly half the market: $|d_j[\hat{\sigma}]| \approx 0.52$. Accordingly, a fee of 20% does not exceed the optimal cap if and only if

[35] See, for example, HOTREC's annual reports, available at https://www.hotrec.eu/.

the platform more than doubles the size of consumers' consideration sets. If the planner disregards platform profits ($\alpha = 0$), and operating costs are uniform, the optimal cap (under $|d_j[\hat{\sigma}]| \approx 0.52$) is 10%, which is below the fees practiced in most markets.

The optimal cap is more likely to bind the lower is the hotels' profit margin. For instance, if the profit margin is 10% of retail prices (instead of 20%), a fee of 20% does not violate the optimal utilitarian cap if and only if the platform at least triples the size of consumers' consideration sets (or, equivalently, $|d_j[\hat{\sigma}]| \approx 0.35$). This is illustrated by the green curve on Figure 1. By contrast, under a 30% profit margin, the 20% fee is welfare-neutral if and only if the platform increases the size of consumers' consideration sets by 60% (or, equivalently, $|d_j[\hat{\sigma}]| \approx 0.62$). This is illustrated by the blue curve on Figure 1.

Obviously, empirical work, or experimental evidence, investigating by how much OTA's expand consumers' consideration sets (relative to other online tools) is needed in order to assess the optimal cap regulation in different markets.[36] We conjecture that, in markets with relatively homogenous offers, general-purpose search engines (or other non-commission-based information sources) should give consumers as much information as an OTA. By contrast, non-specialized information sources should perform significantly worse in markets with greater potential for customization.

4.3 Growing Market

An important feature of mature markets is that, whenever profit margins decrease as competition intensifies,[37] the presence of the platform is bound to hurt firms' profits. The reason is that, by enlarging consumers' consideration sets, the platform intensifies competition without 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*

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

[36] One possibility is to use the equilibrium condition (2) to retrieve the potential demands of firms. The difficulty here is that the "pre-visit" potential demands implied by this condition belong to a world where consumers know they shall soon visit the platform to learn about the hotels available. By contrast, the potential demands relevant for regulation belong to a world where consumers already engaged in search through other means. So this approach is likely to over-estimate the optimal cap.

[37] Zhou (2017), in the context of a Perloff-Salop model, proves that $\lambda(n)$ is decreasing in n if the hazard rate of the match-value cdf is increasing - see also Quint (2014) and Gabaix et al (2016). Allowing match values to be correlated, as we do here, Chen and Riordan (2007, 2008) provide conditions under which $\lambda(n)$ is increasing. While unexceptional in theory, price-increasing competition is unlikely to occur in most markets pertinent to our analysis.

The existence of a platform expands sales (by attracting 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 knowing 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 platform, which is left unmodeled in our partial-equilibrium analysis. We take a conservative stance by assuming that the outside option of latent consumers coincides with the expected payoff of all other consumers were the platform absent.[38] Under this assumption, whenever the cap binds, we obtain the following welfare objective:

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

As in the case of mature markets, we want to express the optimal cap as a function of firms' potential demands and the market's latent demand, which are the key indicators of the platform's informational contribution. To this end, we employ the following identity, which generalizes equation (5) to symmetric information profiles which latent demand is non-empty:[39]

$$\frac{\hat{n}}{N} = \frac{|d_j[\hat{\sigma}]|}{1 - |d_\emptyset[\hat{\sigma}]|}. \tag{8}$$

We construct the approximation by letting the market grow large (i.e., $\hat{n}, N \to \infty$) while holding $|d_j[\hat{\sigma}]|$ and $|d_\emptyset[\hat{\sigma}]|$ constant, as described in the next proposition.

Proposition 4 ***(optimal regulation: growing market)*** *Suppose that the planner is utilitarian. 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 $0 < \underline{z} < \bar{z} \leq \infty$. The utilitarian cap is*

$$\tilde{f}_1 \;=\; b + \mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] + |d_\emptyset[\hat{\sigma}]|\,\lambda(N) \;\approx\; b + (1 - |d_\emptyset[\hat{\sigma}]|)\left(\frac{1 - |d_j[\hat{\sigma}]|}{|d_j[\hat{\sigma}]|}\right)\lambda(N), \tag{9}$$

where the approximation applies for $\hat{n}$ and N large satisfying (8).

[38] Two other possibilities stand out: One is to set the outside option of latent consumers to zero, which is unsatisfactory, as these consumers should find other beneficial purchases in the absence of the platform. Another possibility is to assume that latent consumers gain nothing from the platform's service, being therefore equally well-off before and after joining the market. This option is likely to under-estimate the platform's contribution to consumer surplus.

[39] This identity is established in the proof of Corollary 2. The reasoning is similar to that of footnote 31.

It is useful to compare equations (7) and (9). Perhaps surprisingly, holding constant the size of firms' potential demand, the utilitarian cap decreases as we move from mature to growing markets. The reason is that, fixing the size of potential demands absent the platform, the reach of the informational profile $\hat{\sigma}$ is higher in growing than in mature markets. This renders the platform's informational contribution lower in growing markets, which explains the tighter cap.

To illustrate Proposition 4, consider a market where hotels enjoy a 20% profit margin, and suppose the presence of the platform is credited for increasing aggregate sales by 50%. Accordingly, the market's latent demand comprised $\frac{0.5}{1.5} \approx 33\%$ of the market. Formula (9) then implies that the utilitarian cap equals 20% when the firm's potential demands in the absence of the platform have size $|d_j[\hat{\sigma}]| \approx 0.42$. Equivalently, a fee of 20% does not exceed the optimal cap if and only if the platform expands hotels' potential demands by at least 140%. By comparison, in a mature market, the utilitarian cap exceeds 20% only if the platform expands potential demands by 100% or more.

5 Other Remedies

We now investigate, in the context of our model, two potential alternatives to regulation capping commissions. One is an outright ban of price parity. The other is to, in the presence of competing platforms, relax price parity (moving from its wide to narrow version). We start with the former.

5.1 Banning Price Parity

Suppose firms are free to set two prices, one for transactions through the platform, and the other for direct sales. Also assume consumers can gather information through the platform and seamlessly complete the purchase in the sales channel offering the lowest price. Then the following is true.

Proposition 5 ***(banning price parity)*** *Banning price parity is outcome-equivalent to capping the platform fee according to $f \leq b$, which leads in equilibrium to $f^* = b$. If the planner is utilitarian, this cap is inefficiently low (as $b < \bar{f}_1$), be the market mature or growing.*

The intuition for the result above is simple: because consumers always choose the lowest price of any given product, firms effectively compete in the sales channel exhibiting the lowest *perceived* marginal cost. While this cost is simply c_j at the direct sales channel, it equals $c_j + f_j - b$ at the platform. Therefore, transactions occur inside the platform if and only if $f_j \leq b$. In equilibrium, the platform sets $f_j = b$ for all firms.

Crucially, banning price parity prevents the platform from appropriating any of the informational (ex-ante) benefits it generates. Only the (ex-post) convenience benefit b is recovered in equilibrium, which leads to an inefficiently low fee. This would push the platform out of the market in many instances where it contributes positively to welfare.

The latter conclusion crucially relies on an (arguably stark) feature of our model; namely, that the platform enjoys no market power in the absence of price parity. This is not the case if consumers exhibit "horizontal preferences" towards purchasing through the platform (as in Johansen and Vergé 2017). Or if firms face an inelastic demand of "direct-sales-only" consumers, which dampens the firms' incentive to arbitrage prices and induce show-rooming (as in Wang and Wright 2020). In these cases, the equilibrium fee in the absence of price parity may well exceed (but also, depending on parameters, fall short of) the welfare-maximizing level.

In light of these possibilities, there is no reason to expect that removing price parity will implement (or even raise) efficiency. In this regard, regulating commissions seems a more flexible and accurate instrument for reining in the platform's monopoly power.

5.2 Platform Competition

It is tempting to think that platform competition, coupled with a relaxation of price parity, might alleviate market distortions, rendering commissions caps 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,[40]
3. Each consumer decides which platform(s) to visit, 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. We assume that consumers incur no cost to visit a platform, but prefer to visit a single platform if they expect not to gain from visiting both.

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 clients of platform a, who consider that firm are those who already had the firm in their original consideration set.

We allow the price parity clause to be either *wide* (preventing sellers from posting a lower price on *any* alternative sales channel), or *narrow* (allowing sellers to differentiate prices across platforms,

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

but preventing lower prices on the direct sales channel). In either case, we look for a perfect bayesian equilibrium.[41] Because we are interested in the impact of competition on equilibrium outcomes, we focus on a symmetric equilibrium, therefore discarding all equilibria where the market tips (leading to a de facto monopolistic platform). The next proposition describes our main finding.

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 joins 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 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 some sales would now occur in the direct channel,[42] where the marginal cost is lower. 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).

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 the firm's profit from platform a necessarily goes down, which renders delisting from both strictly preferable to remaining on both.

[41] 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).

[42] Namely, those for clients of platform a who have the firm in their original consideration set, and prefer the firm to any other firm they met through platform a.

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 wide price parity clause. This again renders discounting direct sales too costly for the firm, which then prefers leaving both platforms.

All in all, moving from wide to narrow price parity fails to foster competition among platforms, and does not reduce equilibrium fees. Indeed, the arguments above echo the skeptic reaction by the German competition authority vis-à-vis Booking.com's proposal to revise its price parity clause.[43]

The conclusion above relies on the inability of platforms to convey their lower commissions to consumers, who single home in equilibrium. This is in contrast to Wang and Wright (2020), who assume that, having learned about a firm on a particular platform, consumers can observe the firm's identity and its prices on all channels. As a result, if a platform (say, a) decreases its commission to a firm, that firm is able is reduce its price and increase sales on platform a. Narrow price parity is then effective in engendering some degree of platform competition. This is, however, no guarantee that the equilibrium outcome is efficient, as, under competition, platforms remain unable to appropriate the informational (ex-ante) benefits they generate.

Of course, it is an empirical question to determine which set of assumptions better describes the market. Our analysis however indicates that there are good reasons to doubt that competition is bound to reduce fees or improve welfare, even under mild versions of price parity.

6 Predictions and policy implications

Our analysis unveils a rich set of predictions, as well as implications for policy. We summarize these conclusions below.

Positive Implications

1) The platform is able to levy high commissions by exploiting the contractual externality that listed firms impose on non-listed ones (by rendering their demand more elastic). As a result, the platform is "must-join," even when its fee exceeds the convenience and informational benefits generated to consumers and firms.

2) Firms accept higher fees the smaller their (pre-platform) potential demands are. Moreover, provided potential demands remain constant, the equilibrium fee is invariant to the degree of competition among firms.

The latter prediction, relating consumer information to platform fees, is testable with cross-sectional data on different markets.

[43] In a recent decision, the Higher Regional Court of Düsseldorf, although acknowledging their anti-competitive effect, declared that narrow clauses are not illegal as they represent the only way to compensate OTA's for the service they provide to hotel operators. The German Federal Cartel Office has appealed the court's judgment.

Policy Implications

We investigate regulation capping the platform's commission. We start with mature markets, for which the platform does not increase aggregate sales.

3) If the planner is utilitarian, the efficient cap is such that the platform operates if and only if it increases the joint surplus of consumers and firms.

4) In large markets, the utilitarian cap is approximately equal to the relative expansion of firms' potential demands multiplied by the firms' profit margin. In a plausible scenario, a fee of 20%*, which is the average across OTA's, does not exceed the optimal cap if and only if the platform more than doubles the size of consumers' consideration sets.*

5) If the welfare measure does not include platforms' profits, the optimal cap is tighter. Namely, it is half of the utilitarian one provided the platform's operating costs are uniformly distributed. In the scenario alluded above, unless the platform more than doubles the size of consumers' consideration sets, a 10% *cap would bind. Most current fees exceed this threshold.*

We then study markets with growth potential, in that the platform brings in new consumers, increasing aggregate sales. The following prediction is empirically testable.

6) While in mature markets the platform always decreases producer surplus, firms are better-off with the platform in growing markets if and only if the demand expansion is sufficiently large. Moreover, 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.

Assuming the planner is utilitarian, we derive the following implications for policy:

7) Perhaps counter-intuitively, holding constant the size of firms' potential demand in the absence of the platform, the optimal cap decreases as we move from mature to growing markets. For instance, if the platform doubles aggregate sales and potential demands, it should be allowed to charge a commission no higher than 12% *(in the plausible scenario from above).*

8) Banning price parity is outcome-equivalent to capping the platform fee at the convenience benefit of a transaction. This cap is inefficiently low, be the market mature or growing.

9) Competition between platforms fails to reduce equilibrium fees if price parity is maintained, either in its wide or narrow forms.

Appendix: Omitted Proofs

Proof of Lemma 1. Consider profiles v and p such that $v_k - p_k$ is invariant to $k \geq 2$, and let $\mu \equiv v_2 - p_2$. Pick $j = 1$ and note that, because G is symmetric,

$$D_1(p_1, p_{-1}; \sigma) = |d_1[\sigma]| \cdot \text{Prob}_G [z_2 + (v_2 - p_2) - (v_1 - p_1) \leq z_1 \,| z_2 \geq \max\{z_2, \ldots, z_n\}],$$

or, equivalently,

$$D_1(p_1, p_{-1}; \sigma) = |d_1[\sigma]| \cdot \Big(1 - H^{(n)}(\mu - (v_1 - p_1))\Big).$$

Hence

$$\frac{\partial D_1}{\partial p_1}(p_1, p_{-1}; \sigma) = -|d_1[\sigma]| \, h^{(n)}(\mu - (v_1 - p_1)).$$

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. Assumption 1 implies that at most one symmetric equilibrium exists. Existence is assured by the fact that

$$v_j - p_j^* = v_j - c_j - \lambda(n) = \delta - \lambda(n) > 0,$$

which is invariant in j. Q.E.D.

Proof of Proposition 1. Consider the putative equilibrium fee profile $(f^*, \ldots, f^*)$. 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

$$\frac{1}{N}(p_1^* - c_1 - f^* + b) = \frac{\lambda(N)}{N}, \tag{10}$$

where we employ Lemma 1 to obtain the equality in (10). If firm 1 rejects and decides to do without the platform, it obtains

$$\max_{p_1} \left\{ |d_1[\sigma]| \cdot \Big(1 - H^{(N)}((v_k - p_k^* - (v_1 - p_1))\Big)(p_1 - c_1) \right\},$$

where the quantity $v_k - p_k^*$ 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

$$\max_{p_1} \left\{ |d_1[\sigma]| \cdot \Big(1 - H^{(N)}((v_k - c_k - f^* - \lambda(N) + b) - (v_1 - p_1))\Big)(p_1 - c_1) \right\}$$

$$= \max_{p_1} \left\{ |d_1[\sigma]| \cdot \left(1 - H^{(N)}(-f^* - \lambda(N) + b - c_1 + p_1)\right)\right) (p_1 - c_1) \right\}$$

$$= \max_{p_1} \left\{ |d_1[\sigma]| \cdot \left(1 - H^{(N)}(p_1 - p_1^*)\right) (p_1 - c_1) \right\}$$

$$= |d_1[\sigma]| \cdot \max_{\Delta p} \left\{ \left(1 - H^{(N)}(\Delta p)\right) (\Delta p + f^* + \lambda(N) - b) \right\} \tag{11}$$

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 (10) is weakly greater than (11).

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 (10) remains greater than (11). Because, by the envelope theorem, (11) is increasing in f^*, and (10) is invariant to f^*, we conclude that (10) has to equal (11) 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 size of 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 $\bar{f}$. Q.E.D.

Proof of Proposition 3. For natural numbers $\hat{n}$ and N satisfying (5), let us denote

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

where $\kappa > 0$. We start with the following lemma.

Lemma 2 *As $\hat{n}$ and N grow large while satisfying (5), we obtain that*

$$\lim_{\hat{n}, N \to \infty} \frac{a_1^{\hat{n}}}{b_1^{\hat{n}}} = 1,$$

which we denote by $a_1^{\hat{n}} \sim b_1^{\hat{n}}$.

Proof of Lemma 2. There are four cases to consider.

Case 1: The cdf G_1 has unbounded support and tail index $\gamma \neq 0$.

By Theorem 3 of Gabaix et al (2018), we know that

$$\lim_{n \to \infty} \left\{ \frac{\mathbb{E}\left[Z^{1:n}\right]}{\left(\bar{G}_1\right)^{-1} \left(\frac{1}{n}\right)} \right\} = \Gamma(1 - \gamma). \tag{12}$$

Now set $\kappa = 1$ and note that

$$\left|\frac{a_1^{\hat{n}}}{b_1^{\hat{n}}} - 1\right| = \left|\frac{a_1^{\hat{n}} - b_1^{\hat{n}}}{b_1^{\hat{n}}}\right| = \left|\frac{\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] - \Gamma(1-\gamma)\left[\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right) - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)\right]}{\Gamma(1-\gamma)\left[\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right) - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)\right]}\right|$$

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

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

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 from Pickands (1986) that

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

Together with (12), this implies that $c^{\hat{n}} \to 0$ and $d^{\hat{n}} \to 1 - \left(d_j[\hat{\sigma}]\right)^{\gamma} \neq 0$ as $\hat{n} \to \infty$. Consequently, we obtain that $a_1^{\hat{n}} \sim b_1^{\hat{n}}$, as wished.

Case 2: The cdf G_1 has unbounded support and tail index $\gamma = 0$.

Set $\kappa \neq 1$ and note that

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

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

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

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

Because the tail index of G_1 is 0 and it has unbounded support, we know from Pickands (1986) that

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

Together with (12), this implies that $c^{\hat{n}} \to 0$ and $d^{\hat{n}} \to \kappa - 1 \neq 0$ as $\hat{n} \to \infty$. Consequently, we obtain that $a_\kappa^{\hat{n}} \sim b_\kappa^{\hat{n}}$.

Take κ_0 and κ_1 in a neighborhood of 1 such that $\kappa_0 < 1 < \kappa_1$. We know that, for each $\varepsilon > 0$, and $j \in \{0, 1\}$, there exists N_j such that $\hat{n} > N_j$ implies that

$$\left| \frac{a_{\kappa_j}^{\hat{n}}}{b_{\kappa_j}^{\hat{n}}} - 1 \right| < \varepsilon.$$

Because, fixing $\hat{n}$, the expression $\frac{a_\kappa^{\hat{n}}}{b_\kappa^{\hat{n}}}$ is monotonic in κ, we conclude that

$$\frac{a_1^{\hat{n}}}{b_1^{\hat{n}}} \in \left(\min_{j\in\{0,1\}} \left\{ \frac{a_{\kappa_j}^{\hat{n}}}{b_{\kappa_j}^{\hat{n}}} \right\}, \max_{j\in\{0,1\}} \left\{ \frac{a_{\kappa_j}^{\hat{n}}}{b_{\kappa_j}^{\hat{n}}} \right\} \right).$$

Therefore, for $\hat{n} > \max\{N_0, N_1\}$,

$$\left| \frac{a_1^{\hat{n}}}{b_1^{\hat{n}}} - 1 \right| < \varepsilon.$$

This shows that $a_1^{\hat{n}} \sim b_1^{\hat{n}}$, as we wanted to prove.

Case 3: The cdf G_1 has bounded support and tail index $\gamma \neq 0$.

By Theorem 3 of Gabaix et al (2018), we know that

$$\lim_{n\to\infty} \left\{ \frac{\bar{z} - \mathbb{E}\left[Z^{1:n}\right]}{\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{n}\right)} \right\} = \Gamma(1-\gamma). \tag{13}$$

Now set $\kappa = 1$, and note that

$$\left| \frac{a_1^{\hat{n}}}{b_1^{\hat{n}}} - 1 \right| = \left| \frac{\bar{z} - \mathbb{E}\left[Z^{1:\hat{n}}\right] - \left(\bar{z} - \mathbb{E}\left[Z^{1:N}\right]\right) - \Gamma(1-\gamma)\left[\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right) - \left(\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)\right]}{\Gamma(1-\gamma)\left[\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right) - \left(\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)\right]} \right|$$

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

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

Because the tail index of G_1 is γ and it has bounded support, we know from Pickands (1986) that

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

Together with (13), this implies that $c^{\hat{n}} \to 0$ and $d^{\hat{n}} \to 1 - \left(d_j[\hat{\sigma}]\right)^\gamma \neq 0$ as $\hat{n} \to \infty$. Consequently, we obtain that $a_1^{\hat{n}} \sim b_1^{\hat{n}}$, as wished.

Case 4: The cdf G_1 has bounded support and tail index $\gamma = 0$.

The proof of this case combines the arguments in cases 2 and 3, and is therefore omitted. This concludes the proof of the Lemma. Q.E.D.

Now note that

$$b_1^{\hat{n}} = \left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right) - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right) = \frac{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} + O\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)^2,$$

which implies that

$$\frac{b_1^{\hat{n}} g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)} = 1 + g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)\frac{O\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)^2}{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)} \longrightarrow 1 \quad \text{as} \quad \hat{n} \to \infty,$$

since $g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)$ is bounded from above and

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

Therefore,

$$b_1^{\hat{n}} \sim \frac{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}. \tag{14}$$

By Theorem 1 of Gabaix et al (2018), we also know that

$$\lambda(N) \sim \frac{1}{N g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)\Gamma(\gamma+2)}. \tag{15}$$

Combining (14) and (15) leads to

$$a_1^{\hat{n}} \sim b_1^{\hat{n}} \underset{\hat{n}}{\sim} \frac{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} = \frac{\left(\frac{1}{d_j[\hat{\sigma}]} - 1\right)}{N g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} \sim \left(\frac{1}{d_j[\hat{\sigma}]} - 1\right)\Gamma(\gamma+2)\lambda(N),$$

concluding the proof. Q.E.D.

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

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

whereas, with a monopolistic platform, the equilibrium profit is

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

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

$$\frac{\lambda(N)}{N} > |d_j[\hat{\sigma}]|\,\frac{\lambda(\hat{n})}{\hat{n}} \quad\Longleftrightarrow\quad \frac{\lambda(N)}{\lambda(\hat{n})} > |d_j[\hat{\sigma}]|\,\frac{N}{\hat{n}}.$$

Denoting by $\hat{\sigma}_n$ the measure of consumers aware of each $\hat{n}$-sized subset of firms, it follows that

$$|d_\emptyset[\hat{\sigma}]| = 1 - \hat{\sigma}_n C(N, \hat{n}) \quad \text{and} \quad |d_j[\hat{\sigma}]| = \hat{\sigma}_n C(N-1, \hat{n}-1),$$

where $C(N, \hat{n})$ is the number of $\hat{n}$-combinations from a set of N elements. Solving for $\hat{\sigma}_n$ using the first equation and plugging into the second leads to

$$\frac{\hat{n}}{N} = \frac{|d_j[\hat{\sigma}]|}{1 - |d_\emptyset[\hat{\sigma}]|}.$$

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

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

as stated. Q.E.D.

Proof of Proposition 4. After differentiating the welfare measure $\tilde{W}(\bar{f})$ with respect to $\bar{f}$ we obtain that the optimal utilitarian cap is

$$\tilde{f}_1 = b + \mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] + |d_\emptyset[\hat{\sigma}]|\, \lambda(N).$$

Following the same steps as in the proof of Proposition 3, but replacing (5) by (8), we obtain that

$$\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] \sim \frac{\left(\frac{1}{\hat{n}} - \frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} = \frac{\left(\frac{1-|d_\emptyset[\hat{\sigma}]|}{|d_j[\hat{\sigma}]|} - 1\right)}{N g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)} \sim \left(\frac{1 - |d_\emptyset[\hat{\sigma}]|}{|d_j[\hat{\sigma}]|} - 1\right)\Gamma(\gamma+2)\lambda(N).$$

After setting $\Gamma(\gamma + 2) = 1$, we obtain the approximation

$$\tilde{f}_1 \approx b + \left(\frac{1 - |d_\emptyset[\hat{\sigma}]|}{|d_j[\hat{\sigma}]|} - 1\right)\lambda(N) + |d_\emptyset[\hat{\sigma}]|\, \lambda(N) = b + (1 - |d_\emptyset[\hat{\sigma}]|)\left(\frac{1 - |d_j[\hat{\sigma}]|}{|d_j[\hat{\sigma}]|}\right)\lambda(N),$$

concluding the proof. 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

$$\frac{\lambda(N)}{N} > \max_{p_1} \left\{\frac{d_j[\underline{\sigma}]}{2}\left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1) + \frac{1}{2}\left(1 - H^{(N)}(p_1 - p_1^*)\right)(p_1 - c_1 - f^* + b)\right\}, \quad (16)$$

where f^* is given by equation (2) and p_1^* is the putative equilibrium price. Notice that, by construction of f^*, in the putative equilibrium, firms are indifferent between joining both platforms or neither platform. The inequality in (16) then implies that, in the putative equilibrium, firms strictly prefer joining neither platform than joining one but not the other.

To prove (16), suppose, to obtain a contradiction, that

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

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

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

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

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

Therefore,

$$\frac{\lambda(N)}{N} < \max_{\Delta p} \left\{ d_j[\underline{\sigma}] \left(1 - H^{(N)}(\Delta p)\right)(\Delta p + f^* - b + \lambda(N)) \right\}, \tag{17}$$

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

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

which follows from (16) 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 (16), 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] Bisceglia, M., J. Padilla and S. Piccolo, 2019. "When Prohibiting Platform Parity Agreements Harms Consumers." Working Paper.

[4] 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.

[5] Calzada, J., E. Manna, and A. Mantovani, 2019. "Platform Price Parity Clauses and Segmentation." UB Economics Working Papers 2019/387.

[6] Cazaubiel, A., M. Cure, B. O. Johansen, and T. Vergé, 2020. "Substitution Between Online Distribution Channels: Evidence from the Oslo Hotel Market." International Journal of Industrial Organization, 69, art. 102577.

[7] Chen, Y., and M. Riordan, 2007. "Price and Variety in the Spokes Model." The Economic Journal, 117(522), 897-921.

[8] Chen, Y., and M. Riordan, 2008. "Price-Increasing Competition." Rand Journal of Economics, 39(4), 1042-1058.

[9] Clark, C. R., U. Doraszelski, and M. Draganska, 2009. "The Effect of Advertising on Brand Awareness and Perceived Quality: An Empirical Investigation Using Panel Data." Quantitative Marketing and Economics, 7(2), 207-236.

[10] Draganska, M., and D. Klapper, 2011. "Choice Set Heterogeneity and the Role of Advertising: An Analysis with Micro and Macro Data." Journal of Marketing Research, 48(4), 653-669.

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

[12] European Competition Network. 2017. "Report on the monitoring exercise carried out in the online hotel booking sector by EU competition authorities in 2016." European Commission.

[13] Ezrachi, A., 2015. "The competitive effects of parity clauses on online commerce." European Competition Journal, 11(2-3), 488-519.

[14] 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.

[15] 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.

[16] Hunold, M., Kesler, R., Laitenberger, U., 2020. "Rankings of online travel agents, channel pricing, and consumer protection." Marketing Science. 39(1), 92-116.

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

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

[19] Mantovani, A., Piga, C., and C. Reggiani, 2020. "Online Platform Price Parity Clauses: Evidence from the EU Booking.com case.". Available at SSRN: https://ssrn.com/abstract=3381299

[20] Moraga-Gonzalez, J. L., Z. Sandor and M. Wildenbeest, 2018. "Consumer Search and Prices in the Automobile Market." Working Paper.

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

[22] Pickands, J., 1986. "The Continuous and Differentiable Domains of Attraction of the Extreme Value Distributions." Annals of Probability, 14(3), 996-1004.

[23] Quint, D., 2014. "Imperfect Competition With Complements and Substitutes." Journal of Economic Theory, 152, 266-290.

[24] Rey, P., nad T. Vergé, 2019. "Secret Contracting in Multilateral Relations." Working Paper.

[25] Rochet, J.-C., and J. Tirole, 2011. "Must-take Cards: Merchant Discounts and Avoided Costs." Journal of the European Economic Association, Vol. 9(3), 462-495.

[26] Ronayne, D., and G. Taylor, 2019. "Competing Sales Channels." Working Paper.

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

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

[29] Sengupta, A. and S. Wiggins, 2014. "Airline Pricing, Price Dispersion, and Ticket Characteristics on and off the Internet." American Economic Journal: Economic Policy, 6, 272-307.

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

[31] Wang, C., and J. Wright, 2020. "Search platforms: showrooming and price parity clauses." RAND Journal of Economics, 51(1), 32-58.

[32] Zhou, J., 2017. "Competitive Bundling." Econometrica, 85, 145-172.