Regulating Platform Fees under Price Parity*

Renato Gomes† **Andrea Mantovani**‡

First Version: 17/02/2020 This Version: 27/07/2021

Abstract

Online intermediaries greatly expand consumer information, but also raise sellers' marginal costs by charging high commissions. To prevent disintermediation, some platforms adopted price parity and anti-steering provisions, which restrict sellers' ability to use alternative sales channels. Whether to uphold, reform, or ban these provisions has been at the center of the policy debate, but, so far, little consensus has emerged. As an alternative, this paper studies how to cap platforms' commissions. The utilitarian cap reflects the Pigouvian precept according to which the platform should charge net fees no greater than the informational externality it exerts on other market participants.

JEL Classification: D83, L10, L41.

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

---

*For very helpful comments, we thank, without implicating, Emilio Calvano, Claire Chambolle, Nicolas Economides, Emanuele Giovannetti, Clement Gras, Doh-Shin Jeon, Leonardo Madio, Davide Raggi, Carlo Reggiani, Patrick Rey, Benjamin Rosa, Greg Taylor, Tat-How Teh, Jean Tirole, Thibaud Vergé, Chensgi Wang, Julian Wright and Jidong Zhou, as well as participants at the TSE Competition Policy Workshop, the Virtual Market Design Seminar (Bonn, Mannheim and Karlsruheand), the Online Economics of Platforms Seminar (TSE, BU and NUS), the Symplatform: Matching Scholars and Practitioners (Trinity College Dublin, Politecnico di Milano, and Audencia Business School), the 2021 Annual MaCCI Conference, the NYU Stern IO Seminar, the 12th Paris Digital Economics Conference, the 19th IIOC, PSE, the University of Padua, the 21st CEPR/JIE Applied Industrial Organization Conference, and the 2021 Platform Strategy Research Symposium (Boston University). We gratefully acknowledge financial help from the NET Institute (www.NETinst.org). 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), and funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir programme).

†Corresponding Author: Toulouse School of Economics, CNRS, University of Toulouse Capitole: renato.gomes@tse-fr.eu.

‡TBS Business School: a.mantovani@tbs-education.fr.

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 (OTAs), 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, 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 15% per sale, whereas in Booking.com the average fee ranges between 15% and 25%, depending on property type or location. Apple charges 30% in its App Store (for music, apps, and e-books).

A crucial challenge for this business model is the possibility of *show-rooming*, whereby consumers use the marketplace 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. In other cases, platforms imposed "anti-steering provisions," limiting the sellers' ability to inform buyers of alternative sales channels, or outright suppressing direct selling by firms (e.g., Apple's App Store).

Platforms claim that price parity (or anti-steering) is essential to their business, as it curbs opportunistic behavior.[2] By contrast, competition authorities and consumer associations often regard these practices as the source (or, at best, a reinforcer) of platforms' market power, enabling them to levy high commissions.[3] 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 OTAs (HRS and Booking.com in Germany, Booking.com in Sweden). In the US, Amazon recently decided to remove the price parity clause from its contracts with third-

[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 their direct channel more often in Germany, a country that prohibited the most prominent OTAs from applying price parity clauses, than in countries that did not abolish these clauses. Gu and Zhu (2021) document the importance of disintermediation, whereby buyers and sellers find each other on a platform but then connect directly.

[3] This presumption is at the heart of the legal disputes between Apple and Epic Games in the EU and the US.

party sellers, following political pressure.[4] This clause was also dropped by Amazon in the EU in 2013 after regulators in UK and Germany voiced concerns about its anticompetitive effects. The European Commission's proposal for the Digital Markets Act insists in that platforms should allow consumers to directly trade with third-party providers and/or in alternative marketplaces under no pricing constraints.

Yet, it is not entirely clear that a ban, or voluntary removal, of price parity or anti-steering clauses actually produces tangible results.[5] For one, sellers might still prefer operating through the platform, fearing that it may penalize them otherwise.[6] 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 PPPs is a voluntary action, such programs are often a legal way to bypass the ban.[7] Finally, in some markets, where price parity or anti-steering are not contractually imposed, transacting directly with sellers is impractical. This may occur because of physical constraints (e.g., in the case of Uber Eats, pick-up at the restaurant is often costly) or due to inconveniences that fare high relative to the price of the good (e.g., in the case of Google Play, most apps are modestly priced).

All in all, whether one should uphold, reform, or ban price parity (or, more broadly, favor direct selling by firms) has been at the center of the policy debate, but so far little consensus has emerged. In this paper, we investigate a natural alternative; namely, we study how to optimally cap platforms' commissions.[8] Regulation of this kind has been recently enacted for delivery apps during the Covid-19 crisis, with New York City and other jurisdictions introducing caps of 15% per delivery since May 2020.[9] A theoretical framework guiding cap regulation, and assessing the welfare impact of platform's commissions, has however been missing.

[4]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.

[5]Mantovani et al. (2021) collected listed prices on Booking.com in the period 2014-2017 for touristic areas in France and Italy. They compare prices before and after price parity was lifted in France, using as control similar tourist destinations in Italy. Their study finds limited response of hotel prices on Booking.com, both in the short and medium run, in line with evidence by the European Competition Network (2017).

[6]Hunold et al. (2020) show that OTAs down-list hotels that charge lower prices elsewhere with worse rankings, a practice euphemistically called 'dimming' . Amazon is reportedly removing the Buy Box feature for those merchants setting lower prices on different online channels.

[7]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. (2021).

[8]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 2015/2366.

[9]See Li and Wang (2021) for a description of these initiatives.

Model and Results

In our baseline model, we consider a monopolist platform that charges sellers a fee per sale. Our starting observation is that the platform greatly expands 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). The platform also adds convenience to transactions, what makes it the preferred sales channels by consumers (vis-à-vis direct purchasing from sellers).

Under price parity, 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.

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, in which the platform expands the consideration set 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;[10] the main difference being that the contribution to welfare generated by the platform is essentially informational in our setting, as it enables consumers

[10]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).

to realize purchases of much higher (match) value.

Assuming consumers' demand takes the logit form, we illustrate our findings in an application to OTAs. Namely, we evaluate by how much a dominant OTA needs to expand consumers' consideration sets (relative to alternative ways of gathering information online) so that the *current* average commission meets the utilitarian cap. In a plausible scenario,[11] we find that a 25% commission fee does not exceed the utilitarian cap only if the platform at least triples hotels' potential demands. If the welfare measure does not include platforms' profits, but only producer and consumer surplus, the optimal cap is tighter, as the platform would need to multiply by six the hotels' potential demands. 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.

We then study expanding markets, 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 may decrease as we move from mature to expanding markets. The reason is that, for the same potential demand, the expansion of consumers' consideration sets (produced by the platform) is lower in expanding markets, which explains the tighter cap. For instance, considering again the logit setting, the platform should now quadruple hotels' potential demand for a 25% commission fee to meet the utilitarian cap.[12]

The adoption of a specific demand model allows us to obtain quantitative insights about the optimal cap, at the peril of perhaps assuming too much about the distribution of consumers' tastes. In alternative to a logit specification, we can apply techniques from extreme-value theory (first developed by Gabaix et al. 2016) to isolate the key features of the demand model that matter for policy (therefore weakening the informational requirements of cap regulation). Namely, in the context of random utility models, we show that the consumer's informational gain from considering more firms is asymptotically equivalent to the firms' profit margin multiplied by a factor which is proportional to the relative expansion of consumers' consideration sets. The formula that we obtain is remarkably simple to use and leads to conclusions similar to those obtained under logit demands.

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

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

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

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 expanding, 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.[13] 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.[14]

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 expanding 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,

[13] 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.

[14] 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, or brought significant changes in the commission fees charged by OTAs.

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 platform 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. Besides 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.[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] Ronayne and Taylor (2020) analyze a market where two sellers produce a homogenous good which 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 analyze the effect of price parity clauses, which raise commission fees and induce some firms to delist, thus harming consumers.[17] Our paper differs from the above contributions on the way it models the informational role of platforms, which is instrumental to its distinctive focus on regulation.

A discording view is offered by Johansen and Vergé (2017) and Liu, Niu and White (2021), who find that price parity is welfare-enhancing in some circumstances. The former paper proposes a model where consumers are aware of all firms in the market, multi-home across all platforms, and

[15]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.

[16]Relatedly, Bisceglia et al. (2021) 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. The parity restrictions studied by this paper are not directly comparable to the price parity clauses that are the object of our investigation, as direct prices remain unconstrained.

[17]Calzada et al. (2019) consider a model in which both platforms and sellers are exogenously differentiated, focusing on the effect of price parity on the sellers' decision to single/multi-home. They show that these price restrictions induce sellers to single-home, thereby segmenting the market, and to eventually cut their direct sales channel.

are also able to purchase directly from firms.[18] 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. There are instances where the participation constraint is so tight that commissions decrease (relative to unrestricted pricing), benefiting consumers and firms.[19] The latter paper considers instead an economy populated by a monopolist platform, a merchant, and two types of consumers, those who only buy directly and those who may use the platform if advantageous. Regardless of the curvature of demand, both consumer surplus and total welfare increase under price parity. Yet, the platform gains with price parity if and only if the demand curvature is sufficiently small. In this case, the Pigouvian regulatory cap proposed here is slack, as the platform (practicing price parity) is beneficial to welfare.

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 13). By contrast, we study optimal regulation, which maintains price parity, but restricts the platform's ability to levy high commissions. The regulatory approach is arguably more flexible than its competition policy counterpart: As we show, lifting price parity is akin to capping the platform commission, 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 she observes the pair $(\hat{v}_j, p_j)$, where p_j is the price charged by firm j. A consumer can only transact with firms in her consideration set. Consumers are originally heterogeneous on their consideration sets. We describe this heterogeneity by means of a *consideration profile* $\sigma : 2^{\mathcal{N}} \rightarrow \mathcal{B}[0, 1]$, which maps each

[18] They also assume that commissions are secretly offered. Rey and Vergé (2020) 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.

[19] In turn, Schlütter (2020) revisits the set up of Johansen and Vergé (2017) to show that, under price parity, the platform may prefer sellers to collude on prices, therefore having less incentive to ensure competition.

subset of firms (contained in the power set $2^{\mathcal{N}}$) into the set of consumers (contained in the borelian sigma-algebra $\mathcal{B}[0,1]$) who consider exactly that set of firms.

We assume that consideration sets are independent of the consumer's profile of match values z. Because we normalized the mass of consumers to one, it follows that

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

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 (or not having access to) 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. The size of consideration sets is referred to as the *reach* of σ. Moreover, all firms are considered by the same number of consumers, and therefore have potential demands of the same size.[20] Letting n denote the reach of σ, d the size of potential demands, and d_0 the size of the latent demand, the feasibility condition (1) is then equivalent to

$$d = \frac{n}{N}\left(1 - d_0\right). \tag{2}$$

Intuitively, for symmetric consideration profiles, the size of potential demands equals the fraction of firms known by each consumer multiplied by the mass of potential consumers (i.e., those who possess some market information). While special, this class of consideration profiles offers a tractable way to study changes in consumers' information about market offers.

Before accessing 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 $\underline{n} \in \{2, \ldots, N-1\}$, potential demands of size $\underline{d} < 1$, and latent demand of size $\underline{d}_0 \geq 0$.

For simplicity, we assume that visiting the platform is costless to consumers, and that the platform (for instance, through its search tools) can seamlessly transmit market information. Accordingly, all firms *listed on the platform* are added to the consideration set of *all* consumers.

[20]Formally, this means 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.

In case all firms join the platform, consumer information (after the visit) is then described by the consideration profile $\bar{\sigma}$, which is symmetric with maximal reach $\bar{n} = N$. 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. As a result, the firms' potential demands under $\bar{\sigma}$ consist of the entire market: $\bar{d} = 1$.

Alternatively, suppose all firms join the platform, except for some firm j, which refuses to do it. 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 accessing 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

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

Accordingly, by not joining the platform, firm j keeps its original potential demand. However, those consumers in $D_j[\underline{\sigma}]$ now consider a much larger set of competitors. As the next sections reveal, the externality that firms' listing decisions impose on those who remain non-listed, captured by σ^{-j}, is a recurring theme of our analysis.

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,[21] 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. For example, Booking.com and Amazon Marketplace propose different (private) fees on the basis of firm characteristics such as size, location, logistics and distribution activities. 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.[22] As a result, if a firm joins the platform, all of its sales happen through the platform (due to the convenience benefit it generates). More broadly, our model captures the main features of markets where show-rooming is impractical, even when the platform does not contractually impose price parity. A case in point is delivery food apps, where physical constraints prevent consumers from directly purchasing from far-away restaurants.

The timing of the model is summarized below:

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

[22] We relax this assumption in Section 5.

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 purchases from some firm she is aware of (if any).[24]

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, a firm does not change its belief about the fees 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, consider the profiles (v, p) 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 the difference between the match values of 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

$$P_j\left(p_{-j}\left|\sigma, c_j\right.\right) \equiv \arg\max_{p_j} \quad D_j\left(p_j, p_{-j}; \sigma\right)(p_j - c_j), \tag{3}$$

while a pricing 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 the pricing 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.

[24] As we do here, the literature on oligopolistic competition often assumes that the market is covered (among consumers who know some firm). This will be the case if consumers do not have outside options, or if the vertical components of preferences are sufficiently large.

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 a pricing 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 pricing equilibrium.

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

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

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 $p_j^* = c_j + f - b + \lambda(N)$. We assume that firms follow the pricing equilibrium of Lemma 1 whenever consumers enjoy symmetric consideration profiles (or firms believe so).

An important special case of our model is that where the taste vector z is composed of N iid realizations of some univariate distribution G_1 with support contained on $\mathbb{R}$. This specification was first proposed by Perloff and Salop (1985), and has been widely applied thereafter (see Anderson, de Palma, and Thisse 1992). To sharpen some results, we shall specialize our analysis to this setting, which we refer to hereafter as the *random utility model.*[25] Away from random utility, another special case is the spokes model of Chen and Riordan (2007), which generalizes Hotelling.

3 Laissez-Faire

We now study the equilibrium outcome (of the full game) under a monopolistic platform that is able to impose price parity, and faces no regulation of any kind. The platform's profit from each firm j that joins is then the realized demand of firm j times f_j, which is the fee per sale contracted with firm j. The platform's total profit adds up revenues across all firms that join.

[25] In the context of a random utility model, Assumption 1 is implied by G_1 possessing a log-concave density. See, for instance, Caplin and Nalebuff (1991).

We focus on symmetric equilibria where all firms participate.[26] The unique such equilibrium is characterized in the next proposition.

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

$$\frac{\lambda(N)}{N} = \underline{d} \cdot \max_{\Delta p} \left\{ \left(1 - H^{(N)}(\Delta p)\right) (\Delta p + f^* + \lambda(N) - b) \right\}. \tag{4}$$

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 right-hand side of equation (4). The optimal price adjustment then solves

$$\Delta p - \left(\frac{1 - H^{(N)}(\Delta p)}{h^{(N)}(\Delta p)} \right) + f^* + \lambda(N) - b = 0. \tag{5}$$

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 (4) 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 this formally, note that the right-hand side of (4) is increasing in f^*, and that, at $f^* = b$, the optimal price adjustment is $\Delta p = 0$ (as implied by equation 5). Therefore, at $f^* = b$, the right-hand side of (4) equals

$$\underline{d} \cdot \left(\frac{\lambda(N)}{N} \right),$$

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

How high the platform sets its fee in equilibrium depends on the size of firms' pre-visit potential demands. The next corollary reveals that this quantity is a sufficient statistic for performing comparative statics with respect to the pre-visit information profile $\underline{\sigma}$.

[26] This equilibrium can be shown to exhibit the lowest equilibrium fee, and the highest welfare among all equilibria (if other exist). It therefore provides a conservative benchmark to assess the need of regulation (see Section 4).

Corollary 1 ***(comparative statics: potential demands)*** *Consider two pre-visit consideration profiles, $\underline{\sigma}$ and $\underline{\sigma}^\circ$, with potential demands of sizes $\underline{d}$ and $\underline{d}^\circ$, and let f and f° be their respective equilibrium fees under a monopolist platform. Then $f \leq f^\circ$ if and only if $\underline{d} \geq \underline{d}^\circ$.*

According to Corollary 1, the pre-visit information profile affects the equilibrium fee *only* through the size of potential demands $\underline{d}$: as it increases, the equilibrium fee goes down. This result suggests an explanation for differences between fragmented markets, such as Germany, with many family-run and independent hotels, and markets with a limited number of large suppliers, usually chain hotels, as in Scandinavia. The former markets are characterized by a relatively low potential demand, thus justifying high commission fees, whereas the opposite holds in the latter market.

Corollary 1 also reveals that, as long as the size of potential demand remains constant, the equilibrium fee is invariant to changes in the reach $\underline{n}$ (which captures the intensity of competition across firms in the absence of the platform).[27] 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.

More broadly, what explains the platform's ability to levy high commissions while guaranteeing firms' participation? Its source of market power lies in a *contractual externality* (Segal 1999) that listed firms impose on non-listed ones. Namely, delisting from the platform sends the deviating firm to a world where all consumers who 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.[28]

All in all, the equilibrium platform fee may be high enough to actually decrease both firms' profits and consumers' surplus relative to 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 (4), that f^* can be made arbitrarily large as the size of potential demands, $\underline{d}$, approaches zero.[29]

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, the platform sets a public fee and wishes to induce all firms to join. It is possible, however, that equilibria exist, under a public fee, where the platform induces a subset of*

[27] Of course, if the reach $\underline{n}$ increases, the size of the latent demand, $\underline{d}_0$, has to grow to keep potential demands at the same level (see equation 2). Therefore, provided potential demands have the same size, variations in the size of the latent demand also do not affect the equilibrium fee.

[28] 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.

[29] By contrast, the welfare gain generated by the platform is bounded if the cdf G has bounded support.

firms to join. In these equilibria, participating firms pay a fee larger than f^* and welfare (defined below) is lower than under full participation, which only strengthens the case for regulation.

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 (4) is only stronger were the platform to employ a two-part tariff.*[30] *The latter alleviates the double-marginalization problem, enabling the platform to extract rents without raising (as much) the perceived marginal costs of participating firms.*

4 Cap Regulation

The optimal regulation balances the gains from lower commissions and the potential losses from having no specialized 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).[31]

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 trying to book a hotel might directly visit this platform, rather than consulting information sources likely to be redundant ex-post (e.g., travel guides). Accordingly, the consideration profile $\underline{\sigma}$, which describes consumer information *before* visiting the platform, *but* at the doorstep of doing so, 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 $\underline{n}$ (as consumers may turn to information sources they would skip in the presence of the aggregator platform, such as general-purpose 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 alternative search technologies).

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 in a given market. We assume

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

[31] This is the risk mostly emphasized by industry participants on the verge of regulation. Uber's decision to quit Southeast Asia was triggered by the introduction of new regulations governing ride-hailing operations. Apple threatened to terminate ApplePay in Australia if regulation were to require the iPhones's NFC interface to be available to other payment services. Amazon was recently forced to remove a series of products from its website in India to comply with new rules protecting small retailers. Having previously committed to spend \$5.5 billion on e-commerce in India, Amazon may limit future investments in that country.

that the platform's operating cost is private information, being a draw from some distribution Φ, with density ϕ and support contained in $\mathbb{R}_+$. 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 the empty set.

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. Whenever the cap binds (which is the interesting case), the platform's revenue equals $\bar{f}$ (as all firms join under this fee).

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 maximum out of $n \leq N$ coordinates of the random vector z, which describes consumers' profiles of match values. 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. To guarantee quasi-concavity of the planner's objective, we assume that $f + (1-\alpha)\frac{\Phi(f)}{\phi(f)}$ is increasing in f.[32]

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 fee to satisfy $f \leq \bar{f}$. The*

[32] This condition is trivially satisfied if $\alpha = 1$ and boils down to the standard monotonicity of virtual values when $\alpha = 0$. When the reverse hazard rate, $\frac{\phi}{\Phi}$, is decreasing (which is satisfied by most distributions of interest), this requirement imposes an upper bound (strictly greater than one) on the value of α.

welfare-maximizing cap, written as a function of the weight α, implicitly solves

$$\bar{f}_{\alpha} = b + \mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] - (1-\alpha)\frac{\Phi(\bar{f}_{\alpha})}{\phi(\bar{f}_{\alpha})}. \tag{6}$$

Moreover, the welfare-maximizing cap $\bar{f}_{\alpha}$ is increasing in α.

Consider first the utilitarian case, where $\alpha = 1$. Here, the third term in the left-hand of side (6) 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). Accordingly, the platform operates if and only if it increases the surplus of consumers and firms.

More broadly, we interpret Proposition 2 as providing a welfarist foundation to a notion of fairness commonly employed in competition policy circles.[33] According to this notion, "gatekeepers" or "stacked platforms" should not be allowed to leverage on the externalities created by their presence to make other market participants worse off.[34] This precept, operationalized by a commission cap set at the surplus-neutral fee $\bar{f}_1$, is at the heart of recent policy proposals to regulate platforms (see, for instance, Cabral et al. 2021).

Away from the utilitarian case, the welfare-maximizing cap differs from the surplus-neutral fee. For α below 1, the planner gives more weight to consumers and firms, setting the cap below $\bar{f}_1$. 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, whenever the cap binds, the presence of the platform hurts consumers and firms.

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 $\alpha = 0$).*[35] *In this case, letting the quasi-fixed cost k be uniformly distributed on $[0, \bar{k}]$, the optimal cap is half of the surplus-neutral fee:*

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

[33]Lianos (2020) explores the possibility of a fairness-driven competition law. In his words, "adopting fairness as a guiding principle of competition law can be made either in utilitarian terms or on deontological terms" (p. 74).

[34]The externalities may be contractual, as in our model, or of other kind, such as network externalities. See, for instance, Lianos (2019) and Crémer et al. (2019).

[35]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.

The next corollary builds on Proposition 2 to evaluate how the optimal cap varies with the size of firm's potential demands under alternative search technologies.

Corollary 2 ***(comparative statics: potential demands)*** *Suppose the market is mature. In the context of a random utility model, consider two counterfactual consideration profiles $\hat{\sigma}$ and $\hat{\sigma}^{\circ}$, with potential demands of sizes $\hat{d}$ and $\hat{d}^{\circ}$, and let $\bar{f}_{\alpha}$ and $\bar{f}_{\alpha}^{\circ}$ be their α-welfare-maximizing caps. Then $\bar{f}_{\alpha} \leq \bar{f}_{\alpha}^{\circ}$ if and only if $\hat{d} \geq \hat{d}^{\circ}$.*

Intuitively, the cap is larger the higher is the informational benefit generated by the platform, which is inversely proportional to $\hat{d}$. If the platform is informationally redundant ($\hat{d} = 1$), the surplus-neutral fee equals the convenience benefit b (which captures the valued added by the booking and payment interface provided by the platform). 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. This expansion is maximal when firms' potential demands (under alternative search technologies) are small, in which case the cap reaches its highest level. It is worth noting that Corollary 2 can also be stated in terms of the counterfactual reach $\hat{n}$, as, in mature markets, $\hat{d} = \frac{\hat{n}}{N}$ (as revealed by equation 2).

Notice that whether the optimal cap binds or not depends on how the counterfactual consideration profile $\hat{\sigma}$ compares to its pre-visit counterpart, $\underline{\sigma}$. Even when the platform is informationally redundant ($\hat{d} = 1$), in that consumers could easily learn elsewhere the information it provides, the optimal cap binds provided the reach of the pre-visit profile $\underline{\sigma}$ is less than maximal ($\underline{n} < N$). More broadly, Corollaries 1 and 2 imply that the optimal cap binds provided the size of pre-visit potential demands $\underline{d}$ is sufficiently lower than its counterfactual counterpart $\hat{d}$.

The next example specializes Proposition 2 to a widely used discrete choice framework.

Example 1 ***(optimal cap under logit demand)*** *Suppose the market is mature and assume consumer match values are independent across firms and distributed according to a Gumbel distribution with scale parameter $\beta > 0$.*[36] *Then the optimal utilitarian cap is given by*

$$\bar{f}_1 = b - \beta \ln(\hat{d}) = b - \lambda(N) \left(\frac{N-1}{N} \right) \ln(\hat{d}), \tag{7}$$

where $\lambda(N) = \beta \left(\frac{N}{N-1} \right)$ is the logit markup.

Numerical Illustration. To get more quantitative insights about the optimal cap, it is useful to express the surplus-neutral fee in terms of firms' profit margins (which, being observable, allow us

[36]Namely, for each firm j, the match value z_j is an iid draw of the cdf $G_1(x) = \exp\{-\exp\{-x\beta^{-1}\}\}$ with support $\mathbb{R}$.

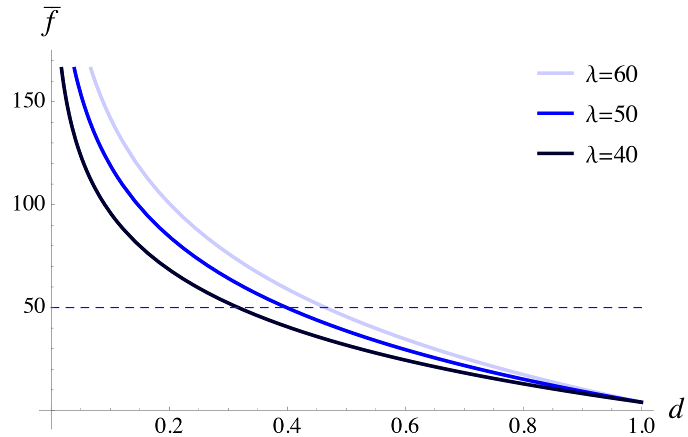

Figure 1: Utilitarian caps as a function of the size of counterfactual potential demands for different profit margins.

to calibrate the scale parameter β). This is the idea behind the second equality in (7). In turn, the size of counterfactual potential demands is the crucial quantity for assessing the informational contribution generated by the platform, and may be measured through surveys and experiments.[37]

We employ simple market data to numerically illustrate Example 1 in the context of OTAs. Specifically, we consider hotels in New York City in 2019. Over that year, the average rate for the cheapest available double room was around 200 dollars, and profit margins varied from 20% to 30% of the retail price.[38] 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). As these services charge around 2%, we set b to 4 dollars. We approximate the ratio $\frac{N-1}{N}$ by one, given the large number of hotels in the market under consideration. 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% (40 dollars per room per day). The utilitarian cap then equals 50 dollars provided the size of counterfactual potential demands is $\hat{d} = 0.31$. Fifty dollars is also the average commission of Booking.com in New York City (roughly 25% of retail price). Therefore, the current average commission does not exceed the utilitarian cap

[37]Our methodology therefore exhibits implementation challenges similar to those of standard market definition exercises, such as the SSNIP test (where survey/experimental evidence is needed). Two early contributions employing micro data on consumers' 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.

[38]Relevant data can be extrapolated from Hotstats (https://www.hotstats.com/), Cheaphotels.org, and Hospitalitynet.org. See also HOTREC's annual reports, available at https://www.hotrec.eu/.

if and only if the platform more than triples the hotel's potential demands under alternative search technologies. If the planner disregards platform profits ($\alpha = 0$), as in Remark 3, the platform would need to multiply by six the hotels' potential demands to meet the welfare-maximizing cap.

The optimal cap is less likely to bind the higher the hotels' profit margin is. For instance, if the profit margin is 30% of retail prices (60 dollars per room per day), current platform fees do not violate the optimal utilitarian cap if and only if the platform at least doubles the size of consumers' consideration sets (i.e. when the size of counterfactual potential demands is $\hat{d} = 0.46$).

Needless to say, the policy prescriptions derived above should not be taken at face value. We view this exercise primarily as an illustration of the results of this article. Empirical work, or experimental evidence, investigating by how much OTAs expand consumers' consideration sets (relative to other search tools) is needed in order to assess the optimal cap regulation in different markets.

4.2 Expanding Market

An important feature of mature markets is that, whenever profit margins decrease as competition intensifies,[39] 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 is expanding, as captured by the fact that the counterfactual latent demand $D_{\emptyset}[\hat{\sigma}]$ is non-empty. Letting $\hat{d}_0 > 0$ denote its size, the next corollary identifies a necessary and sufficient condition for producer surplus to increase when the platform operates.

Corollary 3 ***(industry profits)*** *Relative to the no-platform benchmark, firms gain with the presence of a monopolistic platform if and only if*

$$\hat{d}_0 > 1 - \frac{\lambda(N)}{\lambda(\hat{n})}.$$

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 expanding markets, 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.

[39] 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.

To investigate this point, let us assume for simplicity that the planner is utilitarian ($\alpha = 1$). In order to extend the welfare objective of the previous subsection to expanding markets, 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.[40] Under this assumption, whenever the cap binds, we obtain the following welfare objective:

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

The next proposition derives the optimal utilitarian cap for expanding markets.

Proposition 3 ***(optimal regulation: expanding market)*** *Suppose that the planner is utilitarian and consider regulation that mandates the platform's commission to satisfy $f \leq \bar{f}$. The utilitarian cap is then*

$$\tilde{f}_1 = b + \mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] + \lambda(N)\hat{d}_0. \tag{8}$$

It is useful to compare equations (6) and (8) for $\alpha = 1$. Relative to mature markets, the utilitarian cap in expanding markets exhibits an extra term that captures the profit gain that firms experience from the expansion in demand. As in the case of mature markets, this cap binds the platform's profit by the externality it imposes on firms and consumers. Its implementation is however more demanding, as it requires estimates of the latent demand $\hat{d}_0$ and the counterfactual reach $\hat{n}$ (or, equivalently, the size of potential demands $\hat{d}$, as determined by the feasibility condition 2).

The next example reconsiders the logit setting of Example 1.

Example 2 ***(optimal cap under logit demand)*** *Suppose the market is expanding and consider the logit specification of Example 1. Then the optimal utilitarian cap is given by*

$$\begin{aligned}
\tilde{f}_1 &= b - \beta \ln\left(\frac{\hat{d}}{1-d_0}\right) + \lambda(N)d_0 \\
&= \underbrace{b - \lambda(N)\left(\frac{N-1}{N}\right)\ln(\hat{d})}_{\textit{cap under mature market}} + \underbrace{\lambda(N)\left(\left(\frac{N-1}{N}\right)\ln\left(1-d_0\right) + d_0\right)}_{\textit{expanding market adjustment}},
\end{aligned}$$

where $\lambda(N) = \beta\left(\frac{N}{N-1}\right)$ is the logit markup.

[40] 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 other markets) 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 "discovering" the market. This option is likely to under-estimate the platform's contribution to consumer surplus.

For easier implementation, Example 2 expresses the utilitarian cap as a function of the sizes of potential and latent demands (under alternative search technologies). It decomposes the surplus-neutral fee $\tilde{f}_1$ as the sum of its mature-market counterpart and an adjustment term, capturing the market expansion effect.

Perhaps surprisingly, if there are sufficiently many firms (N high), the expanding market adjustment is negative, rendering the utilitarian cap in an expanding market lower than it would be were the market mature.[41] The reason is that, *fixing* the size of counterfactual potential demands, the reach of the informational profile $\hat{\sigma}$ is higher in expanding than in mature markets (by the feasibility condition 2). This renders the platform's informational contribution lower in expanding markets, which explains the tighter cap.

To numerically illustrate Proposition 3, consider a relatively non-obvious tourist destination such as Toulouse. In this market, in 2019, the average price for the cheapest available double room was around 80 euros.[42] If the platform increases latent demand by 50%,[43] and profit margins average 20% (16 euros), the current average commission by Booking.com (around 15% in a city like Toulouse, i.e. 12 euros) meets the cap if and only if potential demands are doubled by the platform (or quadrupled if the regulator disregards platform's profit).

4.3 Approximating the Optimal Cap

While being well-grounded in theory, the optimal caps derived in Propositions 2 and 3 may be hard to implement in practice, as they require knowledge of the distribution of consumer tastes across firms (needed to compute the informational benefit generated by the platform). One way to bypass this difficulty is to assume a specific demand model (logit, for instance, as often employed in empirical work, and in Examples 1 and 2 above). Another possibility, pursued here, is to explore properties of order statistics to construct an approximation that works for a large class of distributions. We do so by employing the techniques based on extreme-value theory introduced by the pioneering work of Gabaix et al. (2016). Due to its somewhat technical nature, readers more interested in the competition policy alternatives to cap regulation may directly skip to Section 5.

Applying extreme-value theory requires however that we specialize the discrete-choice setup adopted so far. Specifically, we adopt the random utility model of Perloff and Salop (1985) and assume that the distribution of consumer match values (iid across firms and consumers), denoted by G_1, satisfies the following.

[41] This follows from the fact that $\ln(1-d_0)+d_0<0$ for all $d_0\in(0,1)$, and that $\frac{N-1}{N}\approx 1$ for N high.

[42] See https://www.statista.com/statistics/752149/hotel-double-standard-room-costs-toulouse-city/

[43] This estimate is consistent with recent developments. In 2017, Booking.com was made unavailable in Turkey. The Small Hotels Association estimated that small hotels incurred a drop of $50-60\%$ in reservations (see https://businessturkeytoday.com/hotel-owners-are-waiting-for-a-solution-regarding-the-issue-about-booking-com.html).

Assumption 2 *In the context of a random utility model, assume that the match value distribution G_1 is twice continuously differentiable over a bounded support $(\underline{z}, \bar{z})$. Moreover, its density g_1 is non-vanishing and exhibits a bounded derivative.*[44]

For the class of distributions satisfying Assumption 2,[45] we construct an asymptotic approximation that holds constant the size of latent and potential demands. Namely, we fix $\hat{d}$ and $\hat{d}_0$ and take the limit as $\hat{n}, N \to \infty$ while satisfying the feasibility condition (2). This procedure weakens the informational requirements of cap regulation by isolating the key feature of the demand model that matters for policy (namely, the firms' profit margin). As in reality $\hat{n}$ is unlikely to be large, we see the approximation procedure as a useful mathematical tool, rather than reflecting the conditions of any specific market.

The main analytical result behind the approximation is the subject of the next proposition, which can also be seen as an independent contribution of this paper.

Proposition 4 ***(asymptotics)*** *In the context of a random utility model, assume the match value distribution G_1 satisfies Assumption 2. Then,*

$$\lim_{\hat{n}, N \to \infty} \left\{ \frac{\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right]}{\lambda(N)} \right\} = \frac{1 - \hat{d}_0}{\hat{d}} - 1, \tag{9}$$

where the limit above is taken as $\hat{n}$ and N grow large while satisfying (2).

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 (9) should co-move; they are indeed proportional to each other when the market is large. The proportionality factor is the relative expansion of consumers' consideration sets.[46]

We can employ Proposition 4 to obtain an "easier-to-use" formula for the utilitarian cap. Namely, we can express the informational benefit generated by the platform as a function of the firms' profit margin. The utilitarian cap of Proposition 3 is then approximated by

$$\tilde{f}_1 \approx b + \left(\frac{1 - \hat{d}_0}{\hat{d}} - 1 \right) \lambda(N) + \hat{d}_0 \lambda(N) = b + \left(1 - \hat{d}_0\right) \left(\frac{1 - \hat{d}}{\hat{d}} \right) \lambda(N). \tag{10}$$

[44] Namely, there exist $\varepsilon > 0$ and $K < \infty$ such that $g_1(x) > \varepsilon$ and $|g_1'(x)| < K$ for all $x \in (\underline{z}, \bar{z})$.

[45] Examples include the uniform, and censored versions of most distributions of interest, such as the normal, Gumbel, lognormal, exponential, among others.

[46] Indeed, by the feasibility condition (2), $\frac{N}{\hat{n}} = \frac{1 - \hat{d}_0}{\hat{d}}$, so the relative expansion of consumers' consideration sets is the first term in the right-hand side of (9).

This approximation works well in small samples, as discussed in detail in Appendix B.[47] Formula (10) is also remarkably simple to use. For instance, if the market is mature and the platform doubles potential demands, its fee should not exceed the convenience benefit added to the firms' profit margin.

Remark 4 ***(unbounded support)*** *Proposition 4 holds for distributions with bounded support, as required by Assumption 2. While this property makes sense in discrete-choice frameworks (as consumer valuations are bounded), most distributions employed in applied work have unbounded support. One can extend Proposition 4 to such distributions provided their rights tails are thin (in a precise mathematical sense). Away from this case, the limit relation (9) misses a non-vanishing second-order term, which explains the difference between (10) and the exact logit caps of Examples 1 and 2. Yet, formula (10) exhibits similar qualitative properties (e.g., decreasing in $\hat{d}$ and $\hat{d}_0$) and almost identical quantitative implications as the exact cap for most distributions of interest.*[48]

5 Other Remedies

We now investigate, in the context of our model, two potential alternatives to regulation capping commissions. One is an outright (*de facto*) ban of price parity, accompanied by unfettered access by consumers to firms' direct sales channels. These ideas are at the heart of the European Commission's proposal for the Digital Markets Act. The other is to relax price parity (moving from its wide to narrow version) in the presence of competing platforms. As discussed in more detail below, this possibility was actively pursued by some European countries in the case of OTAs. 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 $b < \bar{f}_1, \tilde{f}_1$), be the market mature or expanding.*

[47] When $\hat{G}$ is uniformly distributed, convergence is very fast if $\frac{\hat{d}}{1-\hat{d}_0}$ is above a half: a market size of $N = 20$ is enough to render negligible the difference between the exact value and the limit approximation. Convergence is slower if $\frac{\hat{d}}{1-\hat{d}_0}$ is close to zero. Even then, the approximation error is negligible for $N = 150$.

[48] This is described in the working paper version of this article (CEPR Discussion Paper No. DP15048, available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3661391).

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) welfare. In this regard, regulating commissions seems an attractive alternative 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 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,[49]
3. Each consumer decides which platform(s) to visit, and makes a purchasing decision among the firms in her consideration set.

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

Consumers choose a platform without observing neither commissions, the set of listed firms, or prices (but correctly anticipate these outcomes in equilibrium). 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 pre-visit 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, but preventing lower prices on the direct sales channel). In either case, we look for a perfect Bayesian equilibrium.[50] 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 the unique equilibrium where all firms participate in each platform.

Proposition 6 ***(platform competition)*** *Under either narrow or wide price parity, there exists an equilibrium where platforms offer all firms the monopolistic fee f^* of equation (4), half of all consumers join each platform, and firms join both platforms.*

To understand the result above, let us analyze first the case of a wide price parity. Note that condition (4) 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 a firm a fee higher than the equilibrium one. The firm has the option to stay on both platforms and to 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,[51] 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 one. The reason is that such a deviation does not result in more sales, as consumers have no way to detect the discounted fee.

[50] 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).

[51] Namely, the sales for clients of platform a who have the firm in their pre-visit consideration set, and prefer the firm to any other firm they met through platform a.

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 a 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.

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.[52]

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 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.

[52]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 OTAs for the service they provide to hotel operators. The German Federal Cartel Office appealed the Regional Court's judgment, and published in 2020 a report showing limited evidence of show-rooming. On May 2021, the Federal Court of Justice declared that the 'narrow' best price clauses imposed by Booking.com on hotels are incompatible with antitrust law.

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 a "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-visit) potential demands are. Moreover, provided potential demands remain constant, the equilibrium fee is invariant to the degree of competition among firms.

3) While in mature markets the platform always decreases producer surplus, firms are better-off with the platform in expanding markets if and only if the demand augmentation 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.

Policy Implications

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

4) 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.

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.

6) Perhaps counter-intuitively, holding constant the size of firms' potential demand in the absence of the platform, the optimal cap may decrease as we move from mature to expanding markets.

To ease the informational burden of regulation, we can apply techniques from extreme-value theory to approximate the optimal cap. This represents an alternative to assuming a specific demand model (e.g., logit).

7) The utilitarian cap can be approximated by a function of the firms' profit margin and the relative expansion of consumers' consideration sets.

We then investigate alternatives to regulation capping commissions.

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 expanding.

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

A large part of the current regulatory debate focuses on self-preferencing and algorithmic bias, whereby the platform's search tools favor specific sellers or (its own) brands.[53] While clearly important, addressing these issues is unlikely to align platforms' behavior to society's welfare objectives. Indeed, as our analysis reveals, even with a perfectly unbiased search tool (where consumers always find the best match), platforms have the ability to leverage on contractual externalities across firms to levy inefficiently high commissions (due to price parity or to search frictions outside of the platform). In this sense, measures curbing recommendation biases and capping commissions are likely to play a complementary role in the design of optimal regulation.

[53]See, among others, Aguiar and Waldfogel (2018) and Dinerstein et al. (2018) for empirical studies on platforms respectively using recommendation systems and search design, and de Corniere and Taylor (2019) and Teh and Wright (2020) for theoretical contributions exploring intermediation with biased recommendations.

Appendix A: 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,

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

or, equivalently,

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

Hence

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

By Assumption 1, the best response of firm 1 to p_{-1} uniquely 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. The 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. ■

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

$$\frac{1}{N} (p_1^* - c_1 - f^* + b) = \frac{\lambda(N)}{N}, \tag{11}$$

where we employ Lemma 1 to obtain the equality in (11). If firm 1 rejects and decides to go without the platform, it obtains

$$\max_{p_1} \quad \left\{ \underline{d} \cdot \Big(1 - H^{(N)} ((v_k - p_k^* - (v_1 - p_1)) \Big) (p_1 - c_1) \right\},$$

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

$$\begin{aligned}
&\max_{p_1} \quad \left\{ \underline{d} \cdot \cdot \Big(1 - H^{(N)} ((v_k - c_k - f^* - \lambda(N) + b) - (v_1 - p_1)) \Big) (p_1 - c_1) \right\} \\
&= \max_{p_1} \quad \left\{ \underline{d} \cdot \Big(1 - H^{(N)} (-f^* - \lambda(N) + b - c_1 + p_1)) \Big) (p_1 - c_1) \right\} \\
&= \max_{p_1} \quad \left\{ \underline{d} \cdot \Big(1 - H^{(N)} (p_1 - p_1^*) \Big) (p_1 - c_1) \right\}
\end{aligned}$$

$$= \underline{d} \cdot \max_{\Delta p} \left\{ \left(1 - H^{(N)}(\Delta p)\right) (\Delta p + f^* + \lambda(N) - b) \right\} \tag{12}$$

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 (11) is weakly greater than (12).

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 (11) remains greater than (12). Because, by the envelope theorem, (12) is increasing in f^*, and (11) is invariant to f^*, we conclude that (11) has to equal (12) in equilibrium. This is the unique symmetric equilibrium with full participation because the equality (4) admits a unique solution. ■

Proof of Corollary 1. Note that the right-hand side of (4) 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 $\underline{d}$ increases, the equilibrium fee f^* goes down, which establishes the claim. ■

Proof of Proposition 2. Equation (6) follows from differentiating $W(\bar{f})$, which is quasi-concave in $\bar{f}$. To see that $\bar{f}_\alpha$ is increasing in α, note that equation (6) can be rewritten as

$$\bar{f}_\alpha + (1-\alpha)\frac{\Phi(\bar{f}_\alpha)}{\phi(\bar{f}_\alpha)} = b + \mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right]. \tag{13}$$

The left-hand side is increasing in $\bar{f}_\alpha$ (by assumption) and decreasing in α, while the right-hand side is invariant to α and $\bar{f}_\alpha$. Because this equation admits a unique solution, we conclude that $\bar{f}_\alpha$ is increasing in α, as claimed. ■

Proof of Corollary 2. In the context of a random utility model, we will now prove that $\bar{f}_\alpha$ is decreasing in $\hat{d}$. To this end, consider a random variable X with cdf F and recall the following identity from probability theory:

$$\mathbb{E}[X] = \int_0^\infty [1 - F(x)]\, dx - \int_{-\infty}^0 F(x)dx.$$

Because the cdf of $Z^{1:n}$, denoted $\hat{G}_1^{1:n}$, is given by

$$G_1^{1:n}(x) = [G_1(x)]^n,$$

it follows that

$$\mathbb{E}\left[Z^{1:n}\right] = \int_0^\infty [1 - [G_1(x)]^n]\, dx - \int_{-\infty}^0 [G_1(x)]^n\, dx.$$

Hence,

$$\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] = \int_0^\infty \left[[G_1(x)]^{\hat{n}} - [G_1(x)]^N\right] dx - \int_{-\infty}^0 \left([G_1(x)]^N - [G_1(x)]^{\hat{n}}\right) dx.$$

By the feasibility condition (2), it follows that $\hat{d} = \frac{\hat{n}}{N}$. Simple rearrangements lead to

$$\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right] = \int_0^\infty \left[[G_1(x)]^{-N(1-\hat{d})} - 1\right] [G_1(x)]^N\, dx$$

$$-\int_{-\infty}^{0} [G_1(x)]^N \left(1 - [G_1(x)]^{-N(1-\hat{d})}\right) dx.$$

The derivative of $\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right]$ with respect to $\hat{d}$ is then

$$N \int_{\{x:G_1(x)>0\}} [G_1(x)]^{N\hat{d}} \ln [G_1(x)]\, dx < 0,$$

as $\ln(G_1(x)) < 0$ at all x in the support. This proves that $\mathbb{E}\left[Z^{1:N}\right] - \mathbb{E}\left[Z^{1:\hat{n}}\right]$ is decreasing in $\hat{d}$ (or, equivalently, in $\hat{n}$). The result then follows because the right-hand side of (13) is increasing in $\bar{f}_\alpha$ and the left-hand side is decreasing in $\hat{d}$. ■

Proof of Example 1. Let us first derive the cdf of $Z^{1:n}$, the maximum out of n iid draws of a Gumbel distribution. By standard arguments, this cdf is

$$G_1^{1:n}(x) = \left[\exp\{-\exp\{-x\beta^{-1}\}\}\right]^n = \exp\{-n\exp\{-x\beta^{-1}\}\}$$

$$= \exp\left\{-\exp\left\{-\left(\frac{x-\beta\ln n}{\beta}\right)\right\}\right\},$$

which is a Gumbel cdf with location parameter $\beta \ln N$ and scale parameter β. Its mean is therefore

$$\mathbb{E}\left[Z^{1:N}\right] = \beta\ln N + \beta\gamma,$$

where $\gamma \approx 0.577$ is the Euler-Mascheroni constant. From equation (6) evaluated at $\alpha = 1$ it follows that the utilitarian cap is

$$\bar{f}_1 = b + \beta\ln N - \beta\ln\hat{n} = b + \beta\ln\left(\frac{N}{\hat{n}}\right).$$

By the feasibility condition (2), $\hat{d} = \frac{\hat{n}}{N}$. The utilitarian cap is then $\bar{f}_1 = b - \beta\ln(\hat{d})$, as claimed. ■

Proof of Corollary 3. In the no-platform benchmark, the equilibrium profit of each firm is

$$\hat{d}\frac{\lambda(\hat{n})}{\hat{n}},$$

whereas, with a monopolistic platform, the equilibrium profit is

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

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

$$\frac{\lambda(N)}{N} > \hat{d}\frac{\lambda(\hat{n})}{\hat{n}} \quad\iff\quad \frac{\lambda(N)}{\lambda(\hat{n})} > \hat{d}\frac{N}{\hat{n}} \quad\iff\quad \hat{d}_0 > 1 - \frac{\lambda(N)}{\lambda(\hat{n})},$$

where the last step employed the feasibility condition (2). ■

Proof of Proposition 3. Equation (8) follows from differentiating $\tilde{W}(\bar{f})$, which is quasi-concave in $\bar{f}$. ■

Proof of Example 2. As shown in the proof of Example 1,

$$\mathbb{E}\left[Z^{1:N}\right] = \beta \ln N + \beta\gamma,$$

where $\gamma \approx 0.577$ is the Euler-Mascheroni constant. Plugging this into equation (8) leads to

$$\tilde{f}_1 = b + \beta \ln\left(\frac{N}{\hat{n}}\right) + \lambda(N)\hat{d}_0.$$

By the feasibility condition (2), $\frac{\hat{n}}{N} = \frac{\hat{d}}{1-\hat{d}_0}$. The utilitarian cap is then

$$\tilde{f}_1 = b + \beta \ln\left(\frac{1-\hat{d}_0}{\hat{d}}\right) + \lambda(N)\hat{d}_0.$$

Substituting for $\beta = \left(\frac{N-1}{N}\right)\lambda(N)$ and rearranging leads to the second equality in the statement. ■

Proof of Proposition 4. First, note that, under Assumption 2, the tail index of the cdf G_1 is

$$\gamma \equiv \lim_{x\to\bar{z}} \frac{d}{dx}\left(\frac{1-G_1(x)}{g_1(x)}\right) = -1 - \lim_{x\to\bar{z}}\left(\frac{g_1'(x)(1-G_1(x))}{g_1(x)^2}\right) = -1.$$

By the feasibility condition (2), $\hat{d} = \frac{n}{N}\left(1-\hat{d}_0\right)$, and therefore we can write that

$$N = k\hat{n}, \qquad \text{where} \qquad k \equiv \frac{1-\hat{d}_0}{\hat{d}}.$$

Note that $k > 1$, as $N > \hat{n}$. For natural numbers $\hat{n}$ and N, let us denote

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

where $\bar{G}_1(x) = 1 - G_1(x)$ is the survival function and $\left(\bar{G}_1\right)^{-1}$ is its inverse. We start with the following lemma.

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

$$\lim_{\hat{n},N\to\infty} \frac{a_k^{\hat{n}}}{b_k^{\hat{n}}} = 1,$$

which we denote by $a_k^{\hat{n}} \sim b_k^{\hat{n}}$.

Proof of Lemma 2. The cdf G_1 has bounded support and tail index $\gamma = -1$. By Theorem 3 of Gabaix et al. (2016), we know that

$$\lim_{n\to\infty}\left\{\frac{\bar{z} - \mathbb{E}\left[Z^{1:n}\right]}{\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{n}\right)}\right\} = 1. \tag{14}$$

Hence, with the understanding that $N = k\hat{n}$,

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

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

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

Again using the fact that G_1 has tail index -1 and bounded support, Pickands (1986) implies that

$$\frac{\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)}{\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)} = \frac{\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)}{\bar{z} - \left(\bar{G}_1\right)^{-1}\left(\frac{k}{N}\right)} \longrightarrow \frac{1}{k},$$

as $N \to \infty$ (or, equivalently, $\hat{n} \to \infty$). Together with (14), this implies that $c^{\hat{n}} \to 0$ and that $d^{\hat{n}} \to 1 - \frac{1}{k} > 0$ as $\hat{n} \to \infty$. Consequently, we obtain that $a_k^{\hat{n}} \sim b_k^{\hat{n}}$, as wished. ■

Letting $\mu(k, N) \equiv \left(\bar{G}_1\right)^{-1}\left(\frac{k}{N}\right)$, it follows that

$$b_k^{\hat{n}} = \left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right) - \left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right) = \mu(1, N) - \mu(k, N).$$

By the intermediate value theorem,

$$\mu(k, N) - \mu(1, N) = \frac{\partial \mu}{\partial k}(1, N)(k-1) + \frac{1}{2}\frac{\partial^2 \mu}{\partial k^2}(k^*, N)(k-1)^2,$$

where $k^* \in [1, k]$ depends on N. It is straightforward to compute that

$$\frac{\partial \mu}{\partial k}(k, N) = \frac{-\frac{1}{N}}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{k}{N}\right)\right)} \qquad \text{and} \qquad \frac{\partial^2 \mu}{\partial k^2}(k, N) = -\frac{1}{N^2}\left(\frac{g_1'\left(\left(\bar{G}_1\right)^{-1}\left(\frac{k}{N}\right)\right)}{\left[g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{k}{N}\right)\right)\right]^3}\right).$$

Therefore,

$$\frac{\mu(k, N) - \mu(1, N)}{\frac{\partial \mu}{\partial k}(1, N)(k-1)} = 1 + \frac{1}{2}\frac{\frac{\partial^2 \mu}{\partial k^2}(k^*, N)}{\frac{\partial \mu}{\partial k}(1, N)}(k-1)$$

$$= 1 + \frac{1}{2N}\left(\frac{g_1'\left(\left(\bar{G}_1\right)^{-1}\left(\frac{k^*}{N}\right)\right)}{\left[g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{k^*}{N}\right)\right)\right]^3}\right) g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)(k-1).$$

Because k^* is bounded, g_1' is bounded and g_1 is non-vanishing, the second term converges to zero as $N \to \infty$. Therefore,

$$\lim_{N\to\infty}\left\{\frac{\mu(k, N) - \mu(1, N)}{\frac{\partial \mu}{\partial k}(1, N)(k-1)}\right\} = 1.$$

Noting that

$$\frac{\mu(k,N)-\mu(1,N)}{\frac{\partial\mu}{\partial k}(1,N)(k-1)}=\frac{Nb_k^{\hat{n}}g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{k}{N}\right)\right)}{(k-1)}=\frac{b_k^{\hat{n}}g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{\hat{n}}\right)\right)}{\left(\frac{1}{\hat{n}}-\frac{1}{N}\right)},$$

we can conclude that

$$b_k^{\hat{n}}\sim\frac{\left(\frac{1}{\hat{n}}-\frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}. \tag{15}$$

By Theorem 1 of Gabaix et al. (2016), we also know that

$$\lambda(N)\sim\frac{1}{Ng_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)\Gamma(\gamma+2)}=\frac{1}{Ng_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}. \tag{16}$$

Combining (15) and (16) leads to

$$a_k^{\hat{n}}\sim b_k^{\hat{n}}\underset{\hat{n}}{\sim}\frac{\left(\frac{1}{\hat{n}}-\frac{1}{N}\right)}{g_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}=\frac{\left(\frac{N}{\hat{n}}-1\right)}{Ng_1\left(\left(\bar{G}_1\right)^{-1}\left(\frac{1}{N}\right)\right)}\sim(k-1)\,\lambda(N),$$

concluding the proof. ■

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 $f_j\leq b$. 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. ■

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

$$\frac{\lambda(N)}{N}>\max_{p_1}\left\{\frac{d}{2}\left(1-H^{(N)}(p_1-p_1^*)\right)(p_1-c_1)+\frac{1}{2}\left(1-H^{(N)}(p_1-p_1^*)\right)(p_1-c_1-f^*+b)\right\}, \tag{17}$$

where f^* is given by equation (4) 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 (17) then implies that, in the putative equilibrium, firms strictly prefer joining neither platform than joining one but not the other.

To prove (17), suppose, to obtain a contradiction, that

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

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

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

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

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

Therefore,

$$\frac{\lambda(N)}{N} < \max_{\Delta p} \left\{ \underline{d} \left(1 - H^{(N)}(\Delta p)\right) (\Delta p + f^* - b + \lambda(N)) \right\}, \tag{18}$$

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

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

which follows from (17) and the envelope theorem. As a result, no platform can individually set a fee above f^*. Setting a fee below this level is obviously sub-optimal. 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 that

$$\max_{p_1} \left\{ \frac{d}{2} \left(1 - H^{(N)}(p_1 - p_1^*)\right) (p_1 - c_1) + \frac{1}{2} \left(1 - H^{(N)}(p_1 - p_1^*)\right) (p_1 - c_1 - f^* + b) \right\}$$

is the maximum profit a firm can obtain by patronizing a single platform in the putative equilibrium. The same arguments above then lead to (17), 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. ■

Appendix B: Approximating the Optimal Cap

To illustrate Proposition 4, let G_1 be uniform in $[-1, 0]$: $G_1(x) = x + 1$, in which case $g_1(x) = 1$. The inverse of the hazard rate is then

$$\frac{1 - G_1(z)}{g_1(z)} = -z, \qquad \text{with} \qquad \lim_{x \to \bar{z}} \frac{1 - G_1(x)}{g_1(x)} = 0 \qquad \text{and} \qquad \gamma = \lim_{x \to \bar{z}} \frac{d}{dx} \left(\frac{1 - G_1(x)}{g_1(x)} \right) = -1.$$

It is straightforward to compute that, for $n \geq 2$,

$$\mathbb{E}\left[Z^{1:n}\right] = \int_{-1}^{0} x n (x+1)^{n-1} dx = -\left(\frac{1}{n-1}\right),$$

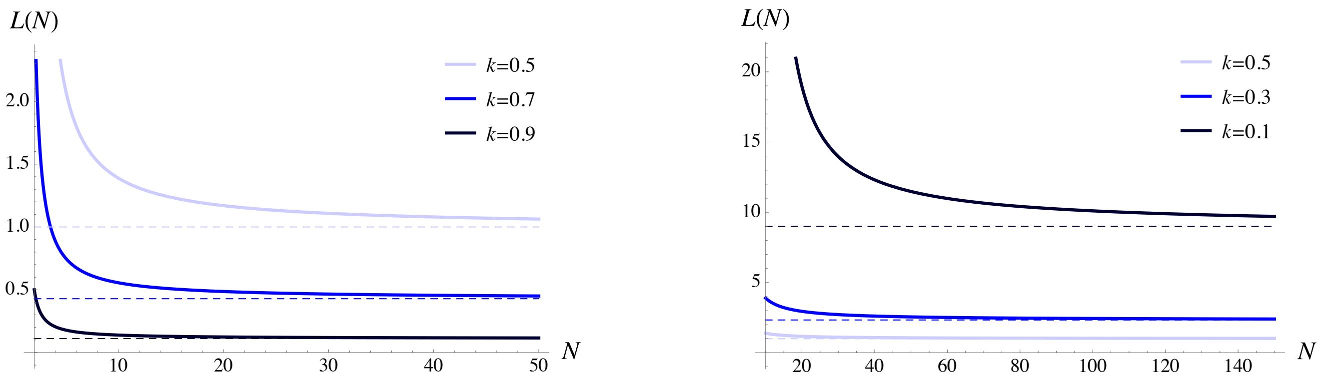

Figure 2: Convergence in the uniform example.

and so

$$\mathbb{E}\left[Z^{1:N}\right]-\mathbb{E}\left[Z^{1:\hat{n}}\right]=-\left(\frac{1}{N-1}\right)+\left(\frac{1}{\hat{n}-1}\right).$$

Moreover, as can be easily verified, the uniform markup is $\lambda(n)=\frac{1}{n}$.

Recall from the feasibility condition (2) that $\frac{\hat{n}}{N}=\frac{\hat{d}}{1-\hat{d}_0}$, in which case

$$L(N)\equiv\frac{\mathbb{E}\left[Z^{1:N}\right]-\mathbb{E}\left[Z^{1:\hat{n}}\right]}{\lambda(N)}=\frac{N^2\left(1-\frac{\hat{d}}{1-\hat{d}_0}\right)}{(N-1)\left(\frac{N\hat{d}}{1-\hat{d}_0}-1\right)}$$

$$=\left(\frac{N}{N-1}\right)\left(1-\frac{\hat{d}}{1-\hat{d}_0}\right)\frac{1}{\left(\frac{\hat{d}}{1-\hat{d}_0}-\frac{1}{N}\right)}.$$

Therefore,

$$\lim_{N\to\infty}L(N)=\left(\frac{1-\frac{\hat{d}}{1-\hat{d}_0}}{\frac{\hat{d}}{1-\hat{d}_0}}\right)=\left(\frac{1-\hat{d}_0}{\hat{d}}-1\right),$$

which is the statement of Proposition 4.

Because closed forms are available for the uniform case, we can easily assess the speed of convergence of the limit above. This is illustrated in Figure 2 for different values of $k\equiv\frac{\hat{d}}{1-\hat{d}_0}$. Inspection reveals that the convergence is very fast for k above a half (left-side panel), but slower for k close to zero (right-side panel). Even at $k=0.1$, a market size of $N=150$ is enough to render negligible the difference between the exact value $L(N)$ and the limit approximation.

References

[1] Aguiar, L., and J. Waldfogel, 2018. "Platforms, Promotion, and Product Discovery: Evidence from Spotify Playlists". NBER Working Paper No. 24713.

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

[3] 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.

[4] Bisceglia, M., J. Padilla, and S. Piccolo, 2021. "When prohibiting wholesale price-parity agreements may harm consumers." International Journal of Industrial Organization, 76, 102738.

[5] 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.

[6] Cabral, L., J. Haucap, G. Parker, G. Petropoulos, T. Valletti, and M. Van Alstyne, 2021. "The EU Digital Markets Act." Publications Office of the European Union.

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

[8] Caplin, A., and B. Nalebuff, 1991. "Aggregation and Imperfect Competition: On the Existence of Equilibrium." Econometrica, 59 (1), 25-59.

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

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

[11] 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.

[12] Crémer, J., Y.-A. de Montjoye, and H. Schweitzer, 2019. "Competition Policy for the Digital Era", Technical Report, European Commission.

[13] De Corniere, A., and G. Taylor, 2019. "A Model of Biased Intermediation." The RAND Journal of Economics. 50(4), 854-882

[14] Dinerstein, M., L. Einav, J. Levin, and N. Sundaresan, 2018. "Consumer Price Search and Platform Design in Internet Commerce." American Economic Review, 108(7), 1820-59.

[15] 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.

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

[17] 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.

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

[19] 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.

[20] Gu, G., and F. Zhu, 2021. "Trust and Disintermediation: Evidence from an Online Freelance Marketplace." Management Science, 67, 794–807.

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

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

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

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

[25] Li, Z., and G. Wang, 2021. "Regulating Powerful Platforms: Evidence from Commission Fee Caps in On-Demand Services." Working Paper.

[26] Lianos, I., 2019. "Competition Law for the Digital Era: A Complex Systems' Perspective", CLES Research Paper Series 6/2019.

[27] Lianos, I., 2020. "Competition Law as a Form of Social Regulation." The Antitrust Bulletin, 65(1), 3-86.

[28] Liu, C., F. Niu, and A. White, 2021. "Optional Intermediaries and Pricing Restraints." Working Paper.

[29] Mantovani, A., C. Piga, and C. Reggiani, 2021. "Online Platform Price Parity Clauses: Evidence from the EU Booking.com case." European Economic Review, 131, 103625.

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

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

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

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

[34] Rey, P., and T. Vergé, 2020. "Secret Contracting in Multilateral Relations." TSE Working Paper, n. 16-744, revised Dec. 2020.

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

[36] Ronayne, D., and G. Taylor, 2020. "Competing Sales Channels." Department of Economics, University of Oxford, Discussion Paper Number 843.

[37] Schlütter, F., 2020. "Managing Seller Conduct in Online Marketplaces and Platform Most-Favored Nation Clauses." Working Paper.

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

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

[40] 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.

[41] Teh, T.-H., and J. Wright, 2020. Intermediation and Steering: Competition in Prices and Commissions. American Economic Journal: Microeconomics, forthcoming.

[42] Tirole, J., 2011. "Payment Card Regulation and the Use of Economic Analysis in Antitrust", Competition Policy International, 7(1), 137–158.

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

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