Monetizing Steering*

Heski Bar-Isaac† Sandro Shelegia‡

September 15, 2023

Abstract

Better-placed products enjoy greater sales. Retailers on marketplace platforms prefer that the marketplace steers customers in their direction. A monopoly marketplace can earn profits through an ad-valorem fee: how should the marketplace set it? It also chooses whether to steer through the design of an algorithm or through an ad auction that raises further revenues. We examine which it prefers. In a single competitive retail market or perfectly discriminating among markets, the marketplace can extract monopoly profits fully through an algorithm or through an ad auction with appropriate fees. The optimal fee for the marketplace depends on the steering method. With an auction, the optimal fee is 0; positive fees induce double marginalization. Thus algorithms are preferred where fees are positive—for example, if there are limits on the ability to steer or to counteract retail market power. However, (natural) limits to fee discrimination can favor steering by auctioned ads. This rationalizes using ad auctions in some markets and algorithms and commissions in others. The approach also speaks to self-steering and highlights that assessing policy impacts requires a holistic view.

1 Introduction

Marketplaces and platforms, ranging from Alibaba and Amazon to Zillow, host numerous products. Indeed, their raison d'être can be understood as allowing consumers to navigate this overwhelming array of options. Even within relatively narrowly defined segments there may be many options available and the marketplace has to choose which products to display

*We thank Ramtin Salamant for excellent research assistance, Gary Biglaiser, Qihong Liu, Devesh Raval, Jen Redmond, Michelangelo Rossi and audiences at Cergy Paris Universite, the University of Toronto, the Humboldt University and University of Toronto Theory Conference, Recanati Business School (Tel Aviv University), Ben Gurion Univesity, 5th University of Bergamo IO Winter Symposium, Toulouse School of Economics 15th Digital Economics Conference, 14th Paris Conference on Digital Economics, IIOC 2023, UIUC/USTC Marketing seminar, Cornell Johnson Conference on Strategy and Economics of Digital Markets and the Competition and Markets Authority for helpful suggestions. Shelegia aknowledges financial support from BBVA foundation, Spanish Ministry of Economy and Competitiveness and ERC.

†University of Toronto, CEPR, and CRESSE. Email: heski.bar-isaac@rotman.utoronto.ca

‡Universitat Pompeu Fabra, BSE, and CEPR. Email: sandro.shelegia@upf.edu

and the order with which to display them. These choices over how to “steer” consumers have consequences for the sales and profits of the retailers of these products. In turn, marketplaces might seek to charge these retailers for steering consumers towards them rather than their rivals. Commentators on policy have noticed and commented on these incentives. For example, Crémer et al. (2019) write “a dominant platform could have incentives to sell ‘monopoly positions’ to their business users” (p.6) and Scott Morton et al. (2019) “Platforms often have a financial incentive to steer customers to particularly profitable products and can use the power of defaults and ordering to accomplish that effectively.”(p.51)

Our focus is to explore different methods that marketplaces can use for steering consumers towards retailers’ offering together with the fees that the marketplace charges. These choices affect not only the marketplace’s profits and the efficiency of matching retailers and consumers, but also the prices that retailers charge and, in this way, overall welfare. Specifically, we contrast two typical approaches to steering: using an algorithm or auctioning ads to allocate privileged positions (and earning auction revenues). In both cases, we suppose that the marketplace also employs some form of a ‘revenue sharing’ agreement, specifically an ad-valorem fee.[1] This is how the marketplace earns in both regimes; if the marketplace allocates steering through an ad auction it also earns the auction revenues.

We take a stylized approach and suppose that (some) consumers consider only a single offering—the one suggested by the platform. This is a simple way to reflect that the prioritized search results are more likely to be purchased, as has been shown to be the case in many studies.[2] This also highlights that some marketplaces do in fact feature privileged positions that are powerful consumer defaults.[3] Our approach allows us to highlight several key forces that would also be present in a richer model where several positions of various visibility are managed by the marketplace.

We find that for a given fee, the ad auction (which grants the winner some market power) results in higher prices than the optimal algorithm due to a familiar ”double mark-up” force. However, when fees are endogenous and set to maximize marketplace profits, then, the marketplace will choose different fees when steering as a result of an auctioned ad or steering by algorithm.

In our baseline model with a single competitive market where steering is necessary for access to consumers, under endogenous fees, retail prices would be identical under either steering method. Prices would be set at the monopoly level and the marketplace would earn monopoly profits. This is a simple but key result in this paper. From a managerial perspective, it highlights that the marketplace has many tools to extract surplus and that it can extract the full monopoly profit through different means. From a policy perspective, it highlights that policies or regulation that focus on a single aspect may achieve little as marketplaces can use

[1]That is the agency model explored in Johnson (2017), for example.

[2]See, for example, Finding 3 of CMA (2017).

[3]See, for example Lee and Musolff (2021) who highlight the importance of the Amazon buybox, or Repricer-Express (2020) which states that “83 per cent of all Amazon sales happen through the Buy Box and even more on mobile.”

a suite of different means to extract surplus.

In our model, an algorithm can achieve any "desired" retail price by committing not to steer any consumers towards any retailers within a market unless they choose a particular price (for example, by listing otherwise relevant retailers low in the search order). Even though this requires commitment and may be costly for a marketplace, there is evidence that marketplaces might engage in this kind of behavior.[4] Furthermore, in our benchmark model and several variations, at the optimal fee the algorithm can implement a desired price by steering consumers towards the lowest-priced retailer. In this sense, the marketplace may appear to be working in consumers' best interests through their algorithmic design when the fees and approach to steering induce monopoly prices.[5] Moreover, this direct-towards-cheapest-offering algorithmic steering is easy to communicate to retailers and so can easily be applied to many markets.

It is noteworthy that in our competitive benchmark model with endogenous fees, although an ad auction and steering by algorithm lead to identical profits and consumer surplus, they do so through different choices of fees.

The ad auction prefers a fee of zero and competition between retailers for the ex-post (undistorted) monopoly profits allow the marketplace to earn monopoly profits as auction revenue while consumers face monopoly prices. Indeed, any circumstance that leads to a strictly positive fee would mean that steering by auction would lead to fees that are too high—above the monopoly price in the retail markets due to a familiar double marginalization problem that the marketplace would rather avoid.

Instead, with algorithmic steering, the marketplace charges a high fee such that retailers' effective costs (which will also be equal to the retail price because of Bertrand-like competition induced by the algorithm) are equal to the (undistorted) monopoly price. This fee, therefore, is "fine-tuned" to demand and cost characteristics of the particular market.

We go beyond the benchmark in several natural extensions and find that a marketplace may prefer either method of steering. As our benchmark model makes clear, steering through ad auctions has the advantage that it involves the same optimal fee (0) across different markets with different demand conditions, whereas an algorithm may prefer different fees in different markets. This highlights the advantage of using an auction when fees are common across markets.

However, there are markets in which an ad auction is less effective. One instance relates to asymmetric retailers where the commitment power embedded in an algorithm might give the marketplace more flexibility in forcing a retail price and extracting an appropriate fee. Another important, and perhaps more interesting, case is when the marketplace features some

[4] Lee and Musolff (2021), for example, note that sometimes nothing is in Amazon Buy Box even though this practice is associated with a 26.31% lower daily sales probability.

[5] In particular, in our model retail offerings are identical (or if not, consumers are directed to the "right" product, so there is no harm associated with "unfair ranking" (see CMA (2020) Section 2.1.4). However, the gatekeeper role of the marketplace leads consumers to suffer monopoly prices even if an algorithmic audit would, presumably, raise no concerns.

markets where consumers are not susceptible to steering (and so there is limited gain to the retailer from securing a privileged position with respect to steering). For these markets, both the auction and the algorithmic approaches rely on fees (and the ad auction would raise no revenue at all, so again the auction and the algorithm would be identical). However, when the same fees are applied across similar markets, some with susceptible consumers and others without, the non-negative fees that the presence of non-susceptible consumers entails, imply pricing distortions (through double marginalization) under auction-based steering. Instead, for the algorithm, absent different demand conditions, the optimal fee for such markets is the same as in the baseline case, and the pricing distortion does not arise.[6]

One clear reason why the two mechanisms diverge in many of these extensions is that, by assumption, we do not allow the fee to be fine-tuned, or else the benchmark model result that the two mechanisms are equivalent would prevail. But the benchmark result is derived for a single competitive market. In practice, the notion of a market for which steering is determined is subtle since steering typically occurs in response to a specific consumer entering a specific search term.[7] Thus, there may be considerably more markets than even the vast number of available products. If everything—that is choice of steering mechanism and fees—depends on characteristics of each consumer/product pairing and the search term, this would correspond to our benchmark model. In principle, ”fine-tuning” to this level would appear to be technologically feasible and, clearly, is optimal.[8] In practice, it does not occur in most marketplaces.[9]

In this paper, we simply take it as given that the marketplace cannot fine-tune its fees and consider the consequences [10] This seems both realistic and a common case. In practice there may be many rationales for such limits: for example, marketplaces may themselves lack the appropriate information, asking small retailers to anticipate consumer-specific fees may be unrealistic, this level of fine-tuning may entail regulatory scrutiny, and so on.

With heterogeneous markets and common fees, either an ad auction or algorithmic steering may be preferred. Indeed, the marketplace can often, trivially, do better with a mix—allocating some steering through ad auctions in some markets and algorithm in others. In

[6]Indeed, that the marketplace faces many markets simultaneously and must charge common fees across them makes the environment quite different from settings where a retailer chooses a selling mechanism for a single good as Bulow and Klemperer (1996).

[7]Certainly, there is scope for marketplaces to personalize the order in which products are displayed, and a 2018 European Commission study found that 61 percent of the 160 e-commerce sites personalised search results rankings. That the outcome of an ad auction depends on consumers characteristics and the search term entered is, of course, familiar.

[8]For example, Uber generates a particular offer of a fee to a driver whenever a rider enters a destination request.

[9]For example, as of June 6, 2022, Amazon.ca on `https://sellercentral.amazon.ca/gp/help/external/200336920` listed 23 categories for referral fees that range from 6% to 45% with the modal fee at 15%. These are broad categories, such as “toys and games”, “home & garden (including pet supplies)”, “consumer electronics” and so on, although there are estimated to be more than 350 million distinct products available through Amazon and Amazon marketplace, as at `https://www.bigcommerce.com/blog/amazon-statistics/#amazon-everything-to-everybody`, and one might imagine in principle many more combinations of feasible consumer characteristics and search terms.

[10]See also Tremblay (2022) which considers the role and limits of fee discrimination.

many marketplaces, in fact, this does indeed seem to be the case. [11] What is optimal depends on the mix of different kinds of markets and their demand conditions. Providing a crisp characterization would rely on strong parametric assumptions and/or a numeric approach. Instead, our approach is to illustrate the forces discussed above through a series of examples that highlight one force at a time and, thereby, consider the implications for the marketplace and for consumers.

After discussing some related literature, we present our results and discuss key intuitions before highlighting some implications for policy.

1.1 Related Literature

There is by now a voluminous literature on platforms—enough for practitioner-oriented and textbook treatments (Evans, Hagiu and Schmalensee, 2006, and Belleflamme and Peitz, 2021). Our focus on how marketplace platforms earn revenues from steering brings us closer to literatures more focused on this aspect.

First, there is a literature focused on prominence in consumer search initiated by Arbatskaya (2007) and Armstrong, Vickers, and Zhou (2009). A number of papers have considered the allocation of prominence, including Armstrong and Zhou (2011), Athey and Ellison (2011), Chen and He (2011) and more recently Anderson and Renault (2021), Teh and Wright (2022), Ke, Lin and Lu (2022). A broader literature has also focused on the extent to which intermediaries might bias recommendations for example, to earn higher commissions (Hagiu and Jullien, 2011; Inderst and Ottaviani, 2012). These papers typically consider one form of allocating prominence (often directly through an algorithm when earning a commission, or instead by selling the position through an auction as is often the case for advertising) and often focus on different questions: for example, whether the right product is recommended when retailers are heterogeneous (a recent example where, in addition, consumers might form incorrect expectations is Heidues, Koster and Koszegi (2022)). Instead, our paper is focused on different methods for allocating prominence and compares them, albeit in an environment in which demand is specified in a simple way; and in much of our analysis different retailers are identical, so many of the interesting and relevant questions around bias do not arise. This paper is also related to a literature that focuses on fees and their regulation (Wang and Wright, 2022); however, this literature has focused on platforms competing with other platforms or other selling venues (such as direct sales), aspects that we do not consider.

A more recent, and rapidly growing literature, both theoretical and empirical, in line with contemporary policy discussion, has focused on the role of marketplaces as retailers and, in particular, whether marketplaces should be allowed to offer their own products, and how

[11] For example, Gutierrez (2021) suggests that, on average, of the prices that consumers pay on Amazon, approximately 15% are ad-valorem referral fees and approximately 5% advertising fees (and 20% on fulfillment an aspect that we do not explicitly cover in this paper). See, also, Mitchell (2021), who provides similar estimates. More recent evidence suggests that Amazon's revenue from advertising has been increasing in recent years and is now closer to 15% (Marketplacepulse (2023).

this might bias steering. Theoretical papers in this literature include Anderson and Bedre-Defoile (2021,2022), de Cornière and Taylor (2014), Etro (2021, 2023), Hagiu and Wright (2015a, 2015b), Hagiu, Teh and Wright (2022), Hervas-Drane and Shelegia (2022), Kang and Muir (2021), Madsen and Vellodi (2021), and Zennyo (2022). Guiterrez (2021), Lee and Musolff (2021), Lam (2021) and Raval (2022), among others, offer empirical evidence on such self-preferencing and its consequences. We highlight that even in the absence of self-preferencing, marketplaces with market power have varied means of earning revenues and that these alternative means operate in somewhat different ways from each other. In Section 6 we show how this perspective has implications for steering and return to discuss the relevant literature.

Indeed, we abstract from many of the considerations of this literature. In much of our analysis, retailers are ex-ante identical (so recommending the "wrong" retailer would involve only a high-priced one). We consider a rather extreme form of prominence whereby (some) consumers consider only a single option listed in a prominent way. These simplifications allow us to focus on an aspect that appears to be largely overlooked in this literature—marketplaces' choices of *how* they monetize steering. In our reading, we have come across two recent papers that examine this question, albeit through slightly different lenses. Ciotti and Madio (2023) consider a single retail market in which retailers are vertically differentiated (most closely related to our extension on market power) and show that allocation by auction is dominated. Long, Jerath, and Sarvary (2022) inspired by the variety of ways that marketplaces monetize (for example, they argue that Alibaba relies on ad revenue to a much greater extent than Amazon) view asymmetric information and learning as key forces. Our model provides a somewhat different explanation for some of the facts they describe—in particular, Long et al.(2022) argue that the relatively large number of sellers on Alibaba compared to Amazon suggests that outside options are more attractive and, in their model, this leads to the use of auctions (and lower marketplace revenues). Instead, our model can interpret this as reduced seller market power leading to greater use of auctions (and need not imply lower marketplace revenues).

2 Benchmark Model

We consider a single retail market. There are two identical retailers who can produce goods at a marginal cost c.[12] In order to access consumers with demand given by a well-behaved function $q(p)$,[13] these retailers require steering from a marketplace. Specifically, suppose that consumers only consider a single retailer and the marketplace can determine which retailer consumers observe (if any).[14] We call this the privileged position.

The marketplace determines both its fee f—a proportion of retailer revenues that retailers

[12] Nothing of substance changes if we assume that there are $n > 2$ retailers.
[13] Specifically, as is standard, we suppose that $(p-c)q(p)$ is strictly quasi-concave for all c such that $q(c) > 0$.
[14] A similar assumption is made in Heidhues, Koster and Koszegi (2022) for example.

pay the marketplace to appear on the marketplace—and its method to determine steering.[15] We consider two methods. First, the marketplace can auction the privileged position—this can be thought of as selling an ad. There are many equivalent auction formats. To be concrete we suppose that this is a second-price auction where retailers bid a lump sum for the privileged position.[16] To avoid uninteresting equilibrium multiplicity, we assume that sellers bid their profits in the auction.

Alternatively, the marketplace can allocate the privileged position on the basis of retailers' prices; in the most general case this can be understood as a pair of functions $\alpha_1(p_1, p_2)$ and $\alpha_2(p_1, p_2)$ from the price vector (p_1, p_2) to probabilities of allocating the privileged position to each of the firms so $\alpha_i \in [0, 1]$ and $\alpha_1 + \alpha_2 \leq 1$. The allocation mechanism entails two forms of commitment. First, when $\alpha_1 + \alpha_2 < 1$ the marketplaces is effectively committing to withhold to consumer from retailers with a positive probability, which entails loss in fees. Second, even when $\alpha_1 + \alpha_2 = 1$ the individual allocation probabilities may not be revenue maximizing for the marketplace given retailer prices. We consider both full and no commitment cases and show that for some settings this assumption makes no difference while for other settings it is pivotal for assessing whether the algorithm is better than the auction or not.

We suppose that the timing is such that the marketplace first chooses the fee and allocation method (including the form of the algorithm). Retailers observe these choices and Then simultaneously choose prices and bids (if relevant).[17] The privileged position is allocated. Consumers then make purchase decisions and fees are collected.

2.1 Exogenous fees

Suppose that fees are set exogenously at f and consider the outcomes associated with the ad auction and with an optimally-chosen algorithm.

In case the privileged position is allocated by auction, the marketplace has no decisions to make. Retailers choices of prices are only relevant contingent on winning and they can, in effect, treat their bids as sunk costs. Retailers would choose prices to maximize their profits. That is each retailer would choose p in order to maximize

$$(1 - f)pq(p) - cq(p) = (p(1 - f) - c)\, q(p). \tag{1}$$

We introduce the notation $p^m\left(\frac{c}{1-f}\right)$ to denote the solution to this problem, which corre-

[15] In several variations of the model, allowing the marketplace to charge a per-unit fee in addition does not change the outcomes or insights. This should be clear—since in many instances the marketplace can earn full monopoly profits with the available instruments, an additional one does not bring further benefits. Of course, when the marketplace cannot earn the full monopoly profits, there is scope for additional instruments (such as a unit fee) may help the marketplace earn more. We discuss this in Appendix A.

[16] For example, if the auction were on a pay-per-impression basis or pay-per-click where the number of consumers clicking was perfectly predictable, the marketplace would obtain identical revenues and prices would be identical.

[17] In principle, one could study sequential choice of prices and bids. If prices are chosen earlier, nothing of substance changes in our analysis. If the reverse is true, mixed bidding strategies may obtain in equilibrium and we purposefully avoid this possibility.

sponds to the solution to the monopoly problem when costs are $\frac{c}{1-f}$. We write $\pi^m\left(\frac{c}{1-f}\right) = \left(p^m\left(\frac{c}{1-f}\right) - \frac{c}{1-f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right)$ as the associated monopoly profits at this marginal cost.

Retailers bid up to their anticipated maximized profits—that is $(1-f)\pi^m\left(\frac{c}{1-f}\right)$ for the privileged position. Consequently, when allocating through an ad auction with a fee given by f, the marketplace will earn:

$$\left((1-f)p^m\left(\frac{c}{1-f}\right) - c\right) q\left(p^m\left(\frac{c}{1-f}\right)\right) + fp^m\left(\frac{c}{1-f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right),$$

where the first term reflects the receipts from the auction and the second term represents the fee income given that the winning retailer charges a price equal to $p^m\left(\frac{c}{1-f}\right)$ and so fee revenues are $fp^m\left(\frac{c}{1-f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right)$.

Next, consider allocating the privileged position by algorithm. The marketplace can choose a rule $\alpha_1(p_1, p_2)$ and $\alpha_2(p_1, p_2)$ to allocate the privileged position. Given any algorithm allocation rule that the marketplace chooses, retailers will choose prices. It is immediate that the marketplace can implement any desired retail price p^* as long as it delivers non-negative profits to a retailer by setting $\alpha_i(p_1, p_2) = 0$ for all $p_i \neq p^*$. It is also immediate that the marketplace prefers to allocate the privileged position than leave it empty (and earn no revenue) and so it chooses $\alpha_1(p^*, p^*) + \alpha_2(p^*, p^*) = 1$; it is natural though unimportant to suppose that $\alpha_i(p^*, p^*) = \frac{1}{2}$.

Then the marketplace's problem becomes choosing p^* to maximize $fp^*q(p^*)$ such that

$$[(1-f)p^* - c]\, q(p^*) \geq 0.$$

If the constraint does not bind then the marketplace would choose p^* to maximize revenues; that is, it would set $p^* = p^m(0)$. If $p^m(0) < \frac{c}{1-f}$ then the retailers' zero profit constraint is violated if the marketplace tries to impose the revenue-maximizing price $p^m(0)$. Under standard assumptions, it is immediate that the marketplace will choose the lowest price possible that satisfies the retailers zero profit constraint. [18] Overall, therefore it chooses $p^* = \max\left\{p^m(0), \frac{c}{1-f}\right\}$.

We can summarize this discussion in the following result.

Proposition 1. *With exogenous fees f, allocating by ad auction leads to retail prices $p^m\left(\frac{c}{1-f}\right)$ and marketplace receipts $\left((1-f)p^m\left(\frac{c}{1-f}\right) - c\right) q\left(p^m\left(\frac{c}{1-f}\right)\right) + fp^m\left(\frac{c}{1-f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right)$. Instead allocating by algorithm leads to retail prices equal to $p^* = \max\left\{p^m(0), \frac{c}{1-f}\right\}$ and marketplace revenues $fp^*q(p^*)$.*

Proof. Follows directly from the discussion in the text. □

It follows that when fees are exogenous, consumers prefer allocation by algorithm.

[18] In particular, $2q'(p) + pq"(p) < 0$ is sufficient.

Corollary 1. *With exogenous fees, consumers prefer allocation by algorithm.*

Proof. Trivially, $p^m\left(\frac{c}{1-f}\right) > \frac{c}{1-f}$ and $p^m\left(\frac{c}{1-f}\right) > p^m(0)$. □

Although consumers prefer allocation by algorithm with a fixed fee, the marketplace might prefer either. For the following proposition, it is convenient to introduce notation for the fee level that induces the monopoly price provided that firms sets price equal to $c/(1-f)$:

$$f^m(c) \equiv \frac{p^m(c) - c}{p^m(c)} \tag{2}$$

This is the Lerner index for the monopoly price for marginal cost c.

Proposition 2. *The marketplace always prefers to steer by auction if fees are low enough, but prefers to steer by algorithm for fees close enough to $f^m(c)$. Thus, in general, the marketplace might prefer to steer by auction or by algorithm.*

Proof. If $f = 0$ then the marketplace recovers the full monopoly profit by using the auction and earns nothing with the algorithm. Profits are continuous so if fees are low enough, the auction is preferred. To see that the marketplace might prefer the algorithm suppose that $f = f^m(c)$ then $\frac{c}{1-f} = p^m(c) > p^m(0)$ so $p^* = p^m(c)$ and the algorithm extracts the full monopoly profit. Instead, the auction extracts strictly less at this fee since the retailer who wins the auction will set a price strictly greater than $p^m(c)$ and so the surplus available (which the marketplace recovers in part through fee revenue and in part as auction receipts) is strictly less than the full monopoly profit. By continuity, for a fee sufficiently close to $f^m(c)$, the algorithm must do better than the auction. □

Note that in the case that the marketplace steers by algorithm, it does not need the commitment power we have assumed. In fact, even if the algorithm designates the winning retailer after prices have been set in the manner that maximizes revenue *ex post*, the same pricing equilibrium obtains. To see this note that non-commitment algorithm always promotes a price closest to $p^m(0)$ (the revenue maximizing price), therefore in the pricing stage firms will not set prices above $p^m(0)$ unless their fee-adjusted marginal cost is higher than this. In this latter case, competition leads them to charge a price equal to this fee-adjusted marginal cost. We therefore conclude that without commitment the algorithm achieves the same profits. This result relies on having two or more retailers. As we discuss in Section 3, with monopoly power, steering, by algorithm but with no commitment to the algorithm, will not be able to achieve full profit maximization in general.

Proposition 3. *When the marketplace cannot commit to the algorithm before retailers choose prices, the algorithm leads to retail prices equal to $p^* = \max\left\{p^m(0), \frac{c}{1-f}\right\}$ and marketplace revenues $fp^*q(p^*)$, that is to the same prices and revenues as with commitment.*

Proof. The marketplace will seek to maximize revenues that is to whichever price is closest to $p^m(0)$. Competition between retailers anticipating this will, therefore, bring prices down to this level unless it is below their fee-adjusted marginal cost. In this latter case, competition between retailers will lead them to set price at their fee-adjusted marginal cost. □

2.2 Endogenous Fees and an Equivalence Result

We have shown that for a fixed fee, retail prices are lower with the algorithm than with the ad auction. Here, we show that when the fees are endogenous then retail prices are identical and consumers and the marketplace are indifferent between the schemes. With either scheme, the marketplace can extract the full monopoly profits.

Proposition 4. *When fees are endogenous the marketplace can earn monopoly profits—that is $\pi^m(c)$—when steering by auction or by algorithm.*

Proof. Following Proposition 1 when the marketplace sets $f = 0$ then the auction earns $\pi^m(c)$. When the marketplace steers by algorithm, dictates a price of $p^m(c)$ and sets a fee so that retailers earn no profits—that is $f = f^m(c)$—then the optimal algorithm earns $\pi^m(c)$. □

Note, that this level of profit is the highest attainable. It is the most that a fully integrated retailer-marketplace could earn. Although the marketplace can achieve this level of earnings either with steering or with the algorithm, and the retail prices in both cases would be identical, the two schemes involve different fees. The optimal auction attains it with a fee of 0; instead, a marketplace attains it with a fee of $\frac{p^m(c)-c}{p^m(c)}$.

Following the same logic as the discussion following Proposition 2, the marketplace can obtain the same profit setting $f = f^m(c)$ and allocating the privileged position to whichever retailer has a lower price. This yields the following result.

Corollary 2. *The marketplace earns $\pi^m(c)$ when steering by algorithm with no commitment to the algorithm that it uses or when committed to allocating the privileged position to the lowest-price retailer.*

Proof. Follows directly from above. □

We now move on from our benchmark model of a single competitive market where consumers are susceptible to steering, and highlight how even with endogenous fees, there may be trade-offs for the marketplace in the choice of ad auction or algorithm for steering.

3 Retail Market Power

If a retailer has market power within a market then allocating steering through an auction earns the marketplace less than doing so by algorithm with full commitment. This can be easily understood and demonstrated by considering the case where the source of such market power is a higher probability of being a good match or a cost advantage. [19]

3.1 Match probability heterogeneity

We assume that with some probability η_i firm i's product is a bad match in which case resulting demand is zero and otherwise is $q(p)$ as in the baseline model. Bad match probabilities are assumed to be independent across retailers, and neither the retailers nor the marketplace know realization of this uncertainty. The marketplace knows the match value probabilities, as do retailers. Without loss, let $\eta_2 \geq \eta_1$.[20] We will assume that the marketplace knows the match value probabilities as do retailers.

We find that the algorithm dominates because it can extract firm 1's monopoly profits while the auction cannot.

Let us start with the algorithm. The marketplace can induce Retailer 1 to set $p^m(c)$, and allocate all consumers to it if Retailer 1 complies. In this case the marketplace can set $f = f^m(c)$ and extract profits equal to

$$(1-\eta_1)\pi^m(c).$$

This is the maximum attainable profit that even an integrated monopolist could achieve since consumers observe only a single offering. If instead the marketplace steers via an ad auction, then both retailers will set the monopoly price given the fee, Retailer 1 will win, and will pay the bid of Retailer 2 equal to its profits, thus the marketplace earns

$$(1-\eta_1) f p^m\left(\frac{c}{1-f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right) + (1-\eta_2)(1-f)\pi^m\left(\frac{c}{1-f}\right). \tag{3}$$

Eq. (3) can be rewritten as

$$\begin{aligned}(1-\eta_1)&\left(p^m\left(\frac{c}{1-f}\right)-c\right) q\left(p^m\left(\frac{c}{1-f}\right)\right) - (\eta_2-\eta_1)(1-f)\pi^m\left(\frac{c}{1-f}\right)\\ &\leq (1-\eta_1)\left(p^m\left(\frac{c}{1-f}\right)-c\right) q\left(p^m\left(\frac{c}{1-f}\right)\right) \leq (1-\eta_1)\pi^m(c),\end{aligned}$$

where the first inequality is strict when $\eta_2 > \eta_1$ and the second one is strict when $f > 0$. We conclude that if $\eta_2 > \eta_1$ for any fee (including the optimal one) the auction will do worse than the algorithm.

The optimal fee for the auction is positive whenever $\eta_2 > \eta_1$ because in (3) the second term is maximized at $f = 0$ and the first term is maximized at $f > 0$, thus the optimal $f > 0$.

[19] There are, of course, many ways to consider market power. For example, Cotti and Madio (2023) examine the case where one firm is of higher quality than the other.

[20] One can think of these probabilities to be drawn among a unit mass of consumers thus the lack of knowledge pertains to particular consumers not total demand.

This means that consumers are better off with the algorithm because there they pay $p^m(c)$ while with the auction they pay $p^m\left(\frac{c}{(1-f^*)}\right) > p^m(c)$ for some $f^* > 0$.

Proposition 5. *The marketplace earns the full monopoly profit. Both the marketplace and consumers are better off with the algorithm.*

Proof. See the discussion in the text above. □

It is noteworthy that commitment is substantive for this result. If the marketplace cannot commit to the form of the algorithm then it is better off with the auction and there is no clear ranking for consumers. We discuss this case in Appendix B.1.

3.2 Cost heterogeneity

In this subsection Retailer 1's market power will stem from its cost advantage. Instead of supposing that both retailers have the same marginal cost, we suppose that Retailer 1 has a lower marginal cost than Retailer 2 where these are given by $c_1 < c_2$, respectively. In this case an auction will raise the monopoly profits of the higher cost firm as auction proceeds; that is it would earn $(1-f)\pi^m\left(\frac{c_2}{1-f}\right)$. [21] Since the lower-cost firm would win the ad auction and set the monopoly price given its (distorted) marginal cost, the marketplace would earn fee revenue of $fp^m\left(\frac{c_1}{1-f}\right)q\left(p^m\left(\frac{c_1}{1-f}\right)\right)$ and so it would choose f to maximize the sum of this fee revenue and the auction revenue. In general, this would involve a positive fee $f > 0$ and so a price $p^m\left(\frac{c_1}{1-f}\right)$ above the monopoly level $p^m(c_1)$.

Instead, an algorithm that can threaten not to allocate the privileged position can earn the marketplace the full monopoly profits by setting

$$\alpha_1(p_1) = \begin{cases} 1 & \text{if } p_1 = p^m(c_1) \\ 0 & \text{otherwise} \end{cases}$$

and the fee at $f^* = f^m(c)$ in a similar way to the case with a competitive retail market and endogenous fees. Trivially, this must do at least as well for the marketplace, and in general, strictly better than allocation by auction.

In this case, as above, although they face monopoly prices when the marketplace uses an algorithm consumers face even higher prices when the marketplace allocates steering through an ad auction. This is because the monopoly retailer retains market power but in effect faces a higher cost ($\frac{c_1}{1-f}$ rather than c_1) this distorts prices to be higher—too high from the marketplace's perspective who can rein in the price when using an algorithm.[22]

[21]Instead, if the marketplace could commit to an observed reservation price for the auction, it would set a fee of 0 and a reservation price of $\pi^m(c_1)$ and would be able to earn the full monopoly profits at the level of the more efficient firm. Thus, this result relies on the assumption that the marketplace cannot set reservation prices in the auction.

[22]As with match probability heterogeneity, with cost heterogeneity the optimal fee is strictly positive.

Proposition 6. *When there is retailer market power through cost heterogeneity, the marketplace earns the full monopoly profit. The marketplace and consumers prefer steering by algorithm to steering by auction.*

It is easy to see that as $c_2 \to c_1$ the marketplace's profits (and consumer surplus) when steering by auction converges to those when steering by algorithm.

Note that as for match probability heterogeneity but in contrast to the analysis in Section 2.2, the ability of the marketplace to commit to an algorithm that might leave the privileged position unassigned is substantive in this case. If the marketplace fully lacks this power, then in fact the auction does strictly better. We illustrate this point formally in Appendix B.2

4 Markets Heterogeneous in Demand and Cost Conditions

Next, we return to the case where there are competitive retailers but we suppose that there are many markets. As discussed in the introduction, we consider the case where the marketplace must choose a common fee across all markets. We index the markets by $i \in \{1, ..., N\}$ and suppose that the demand in each market is $q_i(p)$ and that the associated marginal costs are given by c_i. Similar considerations apply if the marketplace is uncertain regarding demand conditions in a particular market which is an alternative interpretation for this section. As in our benchmark model in Section 2, we suppose that there are at least two identical most-efficient firms in each market.

Following Proposition 1, the optimal fee when allocating steering with an ad auction is 0. In particular, this optimal fee is insensitive to demand conditions. This suggests that by setting a fee of $f = 0$ and allocating steering through an auction, the marketplace can earn the monopoly rents in each market; that is, it earns $\sum_{i=1}^{N} \pi_i^m(c_i)$. In each market, consumers face the monopoly price $p_i^m(c_i)$.

Instead, the optimal fee when employing the algorithm depends on the demand conditions for that market and so requiring a common fee across all markets will not do as well. As above, the marketplace will design the algorithm to dictate prices and allocate the privileged position to a retailer who chooses the preferred price for that market. Consequently, the problem for a marketplace employing an algorithm becomes choosing $p_1, ..., p_N$ and f in order to maximize

$$f \sum_{i=1}^{N} p_i q_i(p_i) \text{ such that } (1-f)p_i \geq c_i \text{ for all } i.$$

It is immediate that this earns the marketplace less than the auction.

For consumers, the comparison is not immediate. Steering by auction entails monopoly prices in all markets. Steering by algorithm leads to higher prices in some markets and lower prices in others and either effect may dominate for overall consumer surplus.

Proposition 7. *When retail markets are competitive and heterogeneous in demand conditions, the marketplace can earn the full monopoly profits through steering with an auction when charging a common fee. This is generally not the case when steering by algorithm.*

Proof. Immediate following the discussion above. □

Of course, if the marketplace prefers the use of an auction to an unrestricted algorithm (that can commit not to assign the privileged position) then a fortiori, it prefers the auction to a restricted class of algorithms.

5 Susceptibility to Steering

Throughout we have assumed that consumers in a market only observe the retailer in the privileged position—that is all consumers are susceptible to steering. In this section we suppose that there are consumers who observe all the retailers. We allow for the possibility of both a mix of susceptible and non-susceptible consumers with a market, denoting the fraction of susceptible consumers in market i by σ_i, and heterogeneity in this fraction across markets. For examples, suppose that $\sigma_i = 0$ so that all consumers are "attentive" it is clear that in such a market, neither the algorithm nor the auction play a role: consumers simply purchase from whichever retailer offers a lower price.

In such fully attentive markets, the marketplace earns fee revenue and Bertrand competition will lead to prices in such markets equal to $\frac{c}{1-f}$. Note that this is identical to the price that would arise under an optimal algorithm for a susceptible market with an endogenous fee as in Proposition 3. [23] In this way, since the optimal algorithm with an optimal endogenous fee leads retailers to price at their effective marginal cost just as Bertrand competition in an attentive market, the marketplace treats these markets identically. Trivially, this is also the case when there is a mix of susceptible and attentive consumers in a market: regardless of σ_i and with an algorithm, for example, that prioritizes the lowest-priced seller then both retailers charge a price equal to $\frac{c}{1-f}$, and the marketplace can extract the full monopoly surplus with a fee equal to $f^m(c)$.

Instead, since the ad auction imposes no discipline on the market power of the retailer who wins the auction, in case the market is one with only susceptible consumers ($\sigma_i = 1$), the retailer charges the monopoly price $p^m\left(\frac{c}{1-f}\right)$ given its effective marginal cost. Instead, in an attentive market ($\sigma_i = 0$) Bertrand competition leads prices equal to the effective marginal cost $\frac{c}{1-f}$. Thus when steering through an ad auction and facing attentive and susceptible markets, the marketplace would prefer different fees for these different markets (0 for the susceptible market and $f^m(c)$ for the attentive one). Imposing a common fee would entail a compromise between these, leading to lower fees (and prices) in the susceptible market and higher fees

[23] Moreover, just as in Corollary 2, the algorithm need not dictate price directly but would be just as effective if it was committed to reward the lower-priced firm with the privileged position.

(and prices) in the attentive market; in general, the consequences for overall consumer surplus would be ambiguous.

Consider the case of a single market with a mix of attentive and inattentive consumers (that is $1 > \sigma > 0$). As above, the marketplace can fully extract the monopoly profits with a fee of $f^m(c)$ and steering through an algorithm. Steering through an ad auction must lead to strictly lower revenues for the marketplace. The presence of susceptible consumers is a force that leads a retailer to price above (fee-adjusted) marginal cost that is a double mark-up force; instead, a fee of 0 would result in the retailers bidding no more than $\sigma\pi^m(c)$ (since the gains from winning the auction only apply to susceptible consumers and $\pi^m(c)$ is an upper bound on the profits per consumer from winning the auction. [24] Thus, it is immediate that for the marketplace steering by algorithm is strictly more profitable than steering through an ad auction.

We summarize this discussion as follows.

Proposition 8. *When demand conditions are identical across all retail markets differing in susceptibility, the marketplace can earn full monopoly profits through steering by algorithm even when charging a common fee across markets. This is generally not the case when steering through auctioned ads.*

Proof. Immediate following the discussion above. □

Of course, a natural interpretation of the same result is to suppose that the marketplace is uncertain about the extent of susceptibility or attentiveness in a market (that is there is uncertainty regarding σ).

6 Heterogeneous Markets, Vertical Integration, and Steering

A theme throughout this paper is that since steering by auction allows ex-post retail monopoly power, it raises the possibility of double marginalization. A possibility that is substantiated when there are reasons for strictly positive fees (such as attentive consumers or market power). As described above, steering by algorithm is one means of mitigating this double mark-up concern since the algorithm can allocate the privileged position on the basis of retail prices.

However, there may be instances where a marketplace prefers to use an ad auction (for example, due to heterogeneous demand conditions) but also faces some markets with attentive

[24] This is in fact a strict upper bound. It cannot be an equilibrium for retailers to bid $\sigma\pi^m(c)$ and charge a price of $p^m(c)$ in equilibrium. A retailer anticipating no profit with such a strategy, would rather deviate, bid lower for the auction (and lose it) and charge a price just under $p^m(c)$ and earn profits of almost $(1-\sigma)\pi^m(c)$. The characterization of this scenario is a little involved: familiar Edgeworth-cycle-like reasoning leads to mixed price strategies from the retailers in equilibrium whose characterization depends on the demand curve, costs, fees and the share of attentive consumers. In a particular, parametrization this can easily be done but deriving expected marketplace revenue and hence optimal fee in general is not analytically tractable. Nevertheless, such a complete characterization is unnecessary for our results in this section.

consumers and others with inattentive consumers–which would otherwise call for steering by algorithm. In case the marketplace was constrained to use only one steering method and used an ad auction, it would want a fee of zero in markets where consumers were inattentive but positive fees where they were attentive; a common fee across all markets would in effect be a compromise between the optimal fees across different markets. It would optimally charge a strictly positive fee and for some markets, double marginalization would arise and prices would be higher than the monopoly level. Instead, if the marketplace were steering by algorithm in every market (to account for a mix of attentive and inattentive ones), the heterogeneous demand conditions imply that the marketplace would prefer different fees in different markets. Again, a common fee would be a compromise entailing fees and retail prices that are too high (from the marketplace perspective) in some markets and too low in others.

Ideally, the marketplace prefers to set fees and the allocation method for steering on a market-by-market basis, but if it cannot do so then it may fail to capture monopoly profits and find itself charging a fee that is too high in some markets and too low in other ones. [25] This occurs when there is sufficient heterogeneity among markets.

Of course, an integrated entity faces fewer constraints on internal transfer or allocation. Vertical integration (well-known as a means of eliminating double marginalization) can allow the marketplace greater flexibility. In effect, it can act as a way for the marketplace to discriminate between markets. In our context, vertical integration into a retail market would involve the marketplace taking over that retail market—perhaps by directly purchasing retailers, or by developing its own production capabilities—and steering consumers towards its own good.

If the marketplace can fully extract monopoly profits with a single allocation mechanism and a common fee (as in the Sections 2.2, 3, 4 and 5 where heterogeneity across markets is in a single dimension) the marketplace has nothing to gain from vertical integration. Instead, when it faces markets that are heterogeneous across different dimensions, it can. For example, consider the case where the marketplace faces two markets simultaneously characterized by retail market power (as in Section 3) and heterogeneous demand conditions (as in Section 4). In general, it would be unable to earn the full monopoly profit with a common fee and allocation mechanism across both markets. Instead, vertically integrating into one market would, trivially, allow the marketplace to earn monopoly profits in that market; then, left facing a single market, the marketplace can choose a steering method and fee that is optimal for that market and fully extract the monopoly surplus there. This logic extends beyond the above example and to any combination of heterogeneities we have discussed in this paper. From the perspective of consumers this vertical integration implies monopoly pricing in both markets. However, the elimination of double-marginalization (with the use of ad auction), or reduction of an inappropriately high fee in at least one market (in case of algorithmic steering), will entail lower retail prices in at least one of the two markets. [26]

[25] See Tremblay (2022) for an analysis of (limited) marketplace fee discrimination.

[26] In the proposition, we do not allow for a mixture of attentive and inattentive consumers within a market since this can lead to mixed strategies in pricing which make comparative statics involved.

Proposition 9. *Consider a marketplace that uses a common steering mechanism and common fee across two heterogeneous retail markets, which may differ in demand conditions and/or retail market power and each consists of either attentive or inattentive consumers. If the market is constrained to employ the same steering method and fee across both markets, then, under vertical integration retail price will be lower in at least one of the two markets. It is possible that retail prices would be strictly lower in both retail markets under vertical integration.*

Proof. Following arguments in the text, after vertical integration, the retail price in each market will be equal to the monopoly price for that market. Thus if the proposition is false with no integration, the retail price in each of the two markets must be below the monopoly price for that market. First, note that this requires that the marketplace engages in steering by algorithm (or that all consumers are attentive so that in effect steering by algorithm or through an ad auction are equivalent). This follows, since if there is steering through an ad auction then after the auction the winning retailer has monopoly power and charges a fee-adjusted monopoly price $p^m\left(\frac{c}{1-f}\right) \geq p^m(c)$. Next, note that if the marketplace is steering by algorithm in both markets and in both markets putting a retailer charging a price below the monopoly price in the privileged position in both markets (or the retail price reflects fee-adjusted marginal cost in case all consumers are attentive), it can marginally increase fees and the prices required to attain the privileged position in both markets and thereby strictly increase profits. Finally, to see that both retail prices may fall following vertical integration (that is both retail prices are above marginal costs), consider the case where demand conditions vary sufficiently between markets, both markets feature some but not too much retail market power (in the sense of Section 3), and all consumers are inattentive so that steering by algorithm with a common fee performs worse than ad auction. In order to capture rents from the better retailer in each market, the marketplace will charge a strictly positive fee even though the heterogeneous demand conditions would call for an auction. As a result of double marginalization, the retailer price in each market would be greater than the monopoly price. It follows that vertical integration would lead to a strict reduction in the retail price in both markets.

□

The underlying mechanism contrasts somewhat with many discussions of self-steering which focus on effects on one particular retail market.[27] Instead, Proposition 9 highlights that vertical integration in one retail market can have (ambiguous) consequences for other retail markets through its impact on a common fee and choice of steering mechanism.

[27] See for example Anderson and Bedre-Defolie (forthcoming), Etro (2022) and Hagiu et. al. (2022). Hervas-Drane and Shelegia (2023) consider independent markets and a ban on dual mode results in a simultaneous change in the marketplace fee and seller markups, similar to this paper.

7 Conclusions and Policy Considerations

This paper has presented a very simple model in which consumers observe only the single seller that the marketplace privileges. For the most part, we consider only a single model variation at a time.[28] Logic and bitter experience in the antecedents to this paper suggest that similar forces arise in richer, less analytically tractable settings that combine elements of many of the model variations outlined above. We believe that our approach makes some clear points that are relevant for academics, policy makers, and practitioners interested in marketplaces and platforms. First, platforms have several tools to monetize their gatekeeping power. We have shown this by making this power extreme (by supposing that susceptible consumers observe only a single seller). In particular, we have shown that in our baseline setting, whether steering through ad auctions or steering by algorithm, a marketplace can extract full monopoly rents; of course, vertical integration also allows a marketplace to extract monopoly profits. This observation is relevant both for the marketplaces themselves and for policy-makers seeking to achieve particular outcomes.

As explored throughout the paper, these different methods extract rents in different ways and so respond differently to different settings. As seems relevant in practice, we suppose that fees are common across markets and find that steering by algorithm is better-equipped to deal with retail market power and with heterogeneity in susceptibility (since it has a greater ability to dictate prices directly and since extracting through fees views susceptible and attentive consumers similarly). Instead, steering by ad auction with less direct control over prices can suffer from double-marginalization and, in our benchmark model, optimizes by setting minimal fees; consequently, it is well-equipped for heterogeneity in demand and cost conditions. Broadly, then the extent to which heterogeneity across markets is dominated by one or other of these types will govern the marketplace's preference for steering by auction or steering by algorithm.

These simple observations have immediate policy implications. Again, though perhaps trivial, these have not been discussed much. First, even absent steering concerns (whether to their own products or to the "wrong" ones when there are heterogeneous sellers) marketplaces have the ability to affect retail market outcomes and consumer welfare.[29] Second, and consistent with a growing concern with the "Whack-A-Mole Challenge" (Franck and Peitz, 2023), marketplace policies must be viewed holistically. For example, direct regulation to reduce fees might lead to more allocation by auction and, in turn, higher prices and lower welfare. Conversely, banning advertising could lead to higher fees and worse outcomes for consumers.

Part of the simplicity in our analysis, arises from the monopoly position of the marketplace

[28] Of course, in practice, there are markets that vary in susceptibility, market power, and demand conditions all at the same time, and the marketplace may choose to use steering through auctioned ads for some of these markets and steering by algorithm for others. It is not hard to write down a model that combines all these elements simultaneously, but it is notationally intensive and the same forces that we describe in this paper would also apply in such a "richer" setting.

[29] It is unsurprising that marketplaces may have a significant influence on retail prices, when fees and ad prices sum to 30 per cent or higher of revenue spent on a marketplace.

and the perfect information of all participants. We believe that analyzing such aspects would prove fruitful for a deeper understanding of these markets.

References

[1] Anderson, Simon and Özlem Bedre-Defolie (2021) "Hybrid platform model, RAND Journal of Economics, forthcoming.

[2] Anderson, Simon and Özlem Bedre-Defolie (2022) "Online trade platforms: hosting, selling, or both?" International Journal of Industrial Organization, 84.

[3] Anderson, Simon and Régis Renault (2021) "Search Direction: Position Externalities and Position Auction Bias " mimeo

[4] Arbatskaya, Maria (2007) "Ordered Search," *RAND Journal of Economics*, 38, 119-126.

[5] Armstrong, Mark and Jidong Zhou (2011) "Paying for Prominence, *Economic Journal*, 121(556), F368-395.

[6] Armstrong, Mark, John Vickers and Jidong Zhou, 2009 "Prominence and Consumer Search," *RAND Journal of Economics*, 40, 209-233.

[7] Athey, Susan and Glenn Ellison (2011) "Position auctions with consumer search." *Quarterly Journal of Economics* 126(3):1213–1270.

[8] Bar-Isaac, Heski and Sandro Shelegia (2022) "Monetizing Steering." CEPR Discussion Paper 17685.

[9] Belleflamme, Paul and Martin Peitz, 2021, *The Economics of Platforms*, Cambridge University Press.

[10] Bulow, Jeremy and Paul Klemperer, (1996) "Auctions Versus Negotiations," *American Economic Review*, 86 (1), 180-194.

[11] Chen, Yongmin and Chuan He C (2011) "Paid placement: Advertising and search on the internet," Economic Journal, 121(556):F309–F328

[12] Ciotti, Fabrizio, and Leonardo Madio (2023) "Competition for Prominence," mimeo

[13] CMA (2017) "Online Search: Consumer and Firm Behaviour—A Review of Existing Literature" at `https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/607077/online-search-literature-review-7-april-2017.pdf`

[14] CMA (2020) “Algorithms: How they can reduce compeition and harm consumers.” at https://www.gov.uk/government/publications/algorithms-how-they-can-reduce-competition-and-harm-consumers/algorithms-how-they-can-reduce-competition-and-harm-consumers

[15] de Cornière Alexandre and Greg Taylor (2014) “Integration and search engine bias.” *RAND Journal of Economics*, 45(3):576–597.

[16] Cremer Jacques, Yves-Alexandre de Montjoye and Heike Schweitzer (2019) textitCompetition Policy for the Digital Era, report for the European Commission (DG Competition).

[17] Etro, Federico (2021a) “Product selection in online marketplaces,” *Journal of Economics & Management Strategy*, 30(3):614–637.

[18] Etro, Federico (2023) “Hybrid marketplaces with free entry of sellers,” Review of Industrial Organization, 62: 119–148

[19] European Commission (2018) “Consumer market study on online market segmentation through personalised pricing/offers in the European Union.” at https://ec.europa.eu/info/publications/consumer-market-study-online-market-segmentation-through-personalised-pricing-offers-european-union-en

[20] Evans, David, Andrei Haigiu and Richard Schmalensee, 2006, *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries,* MIT Press.

[21] Franck, Jens-Uwe and Martin Petiz (2023), “The Digital Markets Act and the Whack-A-Mole Challenge,” mimeo

[22] Furman, Jason, Diane Coyle, Amelia Fletcher, Derek McAuley, D., and Philip Marsden, P. (2019). textitUnlocking digital competition: Report of the digital competition expert panel. Report prepared for the Government of the United Kingdom

[23] Gutiérrez, Germán (2021) “The Welfare Consequences of Regulating Amazon,” mimeo

[24] Hagiu, Andrei and Bruno Jullien (2011) “Why do intermediaries divert search?” *RAND Journal of Economics*, 42(2):337–362.

[25] Hagiu, Andrei and Julian Wright (2015a) “Marketplace or reseller?” *Management Science*, 61(1):184–203.

[26] Hagiu, Andrei and Julian Wright (2015b) “Multi-sided platforms,” *International Journal of Industrial Organization*, 43:162–174.

[27] Hagiu, Andrei, Tat-How Teh, and Julian Wright (2022) “Should platforms be allowed to sell on their own marketplaces?” *RAND Journal of Economics, 53(2): 297–327.*

[28] Hervas-Drane, Andres and Sandro Shelegia (2022) “Retailer-led marketplaces” mimeo.

[29] Heidhues, Paul, Mats Koster, and Botond Koszegi (2022) “Steering Fallible Consumers” mimeo.

[30] Inderst, Roman and Marco Ottaviani (2012)“Compeititon through Commissions and Kickback,” *American Economic Review,* 102(2), 780-809.

[31] Johnson, Justin P. (2017) “The Agency Model and MFN Clauses,” *Review of Economics Studies,* 84(3), 1151-1185.

[32] Kang, Zi Yang and Ellen V. Muir (2022) “Contracting and Vertical Control by a Dominant Platform,” Stanford University, working paper

[33] Ke, Tony, Song Lin and Michelle Y. Lu (2022) “Information Design of Online Platforms,” mimeo

[34] Lam, H. Tai (2021) “Platform Search Design and Market Power. ” mimeo

[35] Lee, Kwok Hao and Leon Musolff (2021) “Entry into two-sided markets shaped by platform-guided search,” Princeton University, working paper

[36] Long, Fei, Kinshuk Jerath and Miklos Sarvary (2022) “Designing an Online Retail Marketplace: Leveraging Information from Sponsored Advertising” *Marketing Science*, 41(1):115-138.

[37] Madsen, Erik and Nikhil Vellodi (2021) “Insider imitation,” Available at SSRN 3832712.

[38] Marketplacepulse (2023) “Amazon Takes a 50% Cut of Seller’s Revenue,” February 13, 2023 Available at https://www.marketplacepulse.com/articles/amazon-takes-a-50-cut-of-sellers-revenue (accessed May 1, 2023).

[39] Mitchell, Stacy (2021) “Amazon’s Toll Road: How the Tech Giant Funds Its Monopoly Empire by Exploiting Small Businesses,” Institute for Local Self-Reliance at https://ilsr.org/amazons-toll-road/.

[40] Raval, Devesh (2022) “Steering in One Click: Platform Self-Preferencing in the Amazon Buy Box,”

[41] RepricerExpress (2022) “How to Win the Amazon Buy Box 2022,” Available at https://www.repricerexpress.com/win-amazon-buy-box/ on June 20, 2022.

[42] Scott-Morton, Fiona (2019) textitReport from the Committee for the Study of Digital Platforms Market Structure and Antitrust Subcommittee, Stigler Center, University of Chicago.

[43] Teh, Tet-How and Julian Wright J (2022) "Intermediation and steering: Competition in prices and commissions." *American Economic Journal: Microeconomics*. 14(2): 281-321.

[44] Tremblay, Mark J. (2022) "The Limits of Marketplace Fee Discrimination," mimeo

[45] Wang, Chengsi and Julian Wright (2022) "Regulating Platform Fees," mimeo

[46] Zennyo, Yusuke (2020) "Platform encroachment and own-content bias" Available at SSRN 3683287.

A Per-unit fees

In this appendix we suppose that the marketplace can charge a per unit fee u in addition to raising revenues from a proportionate fee and possibly through auction revenues. We show that in variations presented in the main paper introducing this additional instrument does not change our conclusions. We also suggest when it may and dicuss implications.

A.1 Single Competitive Retail Market Benchmark

Here we review the analysis in Section 2.1. In case of steering by auction, analogously to (1), firms choose prices to maximize

$$(1-f)pq(p) - (c+u)q(p) = (1-f)\left[\left(p - \frac{c+u}{1-f}\right)q(p)\right],$$

where this expression differs from (1) only in so far as the effective costs for each firm incorporate the per unit fee u as well as marginal costs c and the revenue fee. It is immediate that firms set prices at $p^m\left(\frac{c+u}{1-f}\right)$ and the marketplace earns

$$(1-f)\pi^m\left(\frac{c+u}{1-f}\right) + f\left(p^m\left(\frac{c+u}{1-f}\right) + \frac{u}{f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right).$$

As in Section 2.1, we now consider the marketplace's problem of allocating by algorithm given exogenous fees to be choosing p^* to maximize $f\left(p^* + \frac{u}{f}\right) q\left(p^*\right)$ such that

$$\left[(1-f)p^* - c - u\right] q\left(p^*\right) \geq 0.$$

Again, it is immediate that consumers prefer allocation by algorithm as does the marketplace if all fees are low enough (consider the case $u = f = 0$) but in general might prefer to steer by auction or by algorithm (since this is already true in case $u = 0$).

Just as in Section 2.2, when fees are endogenous, the marketplace can optimize in many ways. As we have already seen when the marketplace sets $f = u = 0$ then the auction earns $\pi^m(c)$ and when the marketplace sets $f = \frac{p^m(c)-c}{p^m(c)}$ then the optimal algorithm earns $\pi^m(c)$. The marketplace can also earn $\pi^m(c)$ by setting $f = 0$ and $u = p^m(c) - c$ with steering by an algorithm that forces retailers to price at $p^m(c)$; or more generally with any steering by algorithm that forces retailers to price at $p^m(c)$ and combination of f and u that extracts all profits.

A.2 Model variations

It is clear that when the marketplace can extract full monopoly profits, either by steering by auction or by algorithm, when $u = 0$, it can also do so with the freedom to set u. It is clear that when it does so by auction (for example, where markets vary in demand conditions as in Section 4) it would choose to set $u = f = 0$. Instead, when it does so by algorithm (as in Sections 3 and 5 which consider market power and susceptibility) then as above, the marketplace has flexibility in choosing combinations of per unit fee u and revenue shares f that can achieve this.

There are cases where the marketplace cannot extract full monopoly profits and a per unit fee does not help it collect more. When there are different demand conditions across different markets but firms have the same cost and markets are competitive, the marketplace using the algorithm can use unit and ad valorem fees interchangeably.[30] However, if markets differ in cost conditions, then using both instruments will, in general, be more profitable. The same is true when retailers have market power, where two taxation tools generally will more readily alleviate the double markup problem when using an auction.

If we were to assume that markets are competitive, but retailers across (but not necessarily within) markets differ in marginal costs, then having two fee instruments may (and in general will) be beneficial. In this case different markets command different optimal price and using u allows more flexibility in achieving this.

B Further results on market power where the algorithm cannot commit

B.1 Match Probability Heterogeneity

Here we consider match probability heterogeneity but we now turn to the case where the marketplace cannot commit to an algorithm. Thus the marketplace, *ex post*, chooses the retailer whose price is most profitable in terms of fees. To simplify, in this and the next section we assume that when indifferent the marketplace chooses Retailer 1. Further, since in equilibrium Retailer 2 does not get any demand, we assume that it prices in such a ways as to maximize the revenues that the marketplace obtains if consumers are allocated to it, subject to a non-negative profit constraint.[31]

In the case of allocating by algorithm, the marketplace will choose Retailer 1 if

$$p_1 q(p_1) \geq \frac{(1-\eta_2)}{(1-\eta_1)} p_2 q(p_2),$$

[30] Since in competitive markets firms price at $p = (c+u)/(1-f)$ and the marketplace extracts all profits $pf + u = \frac{f(c+u)}{1-f} + u$, so for a given target price p, the marketplace earns $p - c$ and the particular way in which u and f are weighted in p is irrelevant.

[31] This assumption is analogous to the standard Econ 101 equilibrium selection in the Bertrand game of perfect substitutes with asymmetric costs where the high cost firm charges its own marginal cost. Retailer 2 in effect maximizes its chance of being selected without risking any losses. This latter assumption—that Retailer 2 would not incur losses if selected—can be justified through trembling-hand type arguments. Subject to Retailer 1 being constrained by Retailer 2's price (as explained in detail below), it must be that $p_2 = p^m(0)$ if $p^m(0) > c/(1-f)$ or else Retailer 2 could profitably deviate to a price that would garner demand and positive profits. When Retailer 2 does not constrain Retailer 1's pricing in equilibrium, this assumption is inconsequential.

and will otherwise choose Retailer 2.

In equilibrium, Retailer 2 can never win the spot because Retailer 1 can always choose the same price (and, in general, can choose a different more profitable price) and get steered toward. In accordance with the assumption above, Retailer 2 sets

$$p_2 = \max(p^m(0), c/(1-f))$$

so Retailer 2 charges the revenue-maximizing price unless it falls short of the fee-adjusted marginal cost, in which case Retailer 2 charges price equal to this cost.

Now we turn to p_1. In equilibrium Retailer 1 has to best respond to p_2 by choosing p_1 that solves

$$p_1 q(p_1) = \frac{(1-\eta_2)}{(1-\eta_1)} p_2 q(p_2) \tag{4}$$

if the highest solution (of the possible two) satisfies $p_1 < p^m(c/(1-f))$ or, else set $p_1 = p^m(c/(1-f))$.

The outcome is rather simple to understand. Retailer 1 always wins access to consumers. If η_2 is close to η_1 then Retailer 1 is constrained by Retailer 2's competitive pressure and gives the marketplace just enough revenue to be selected. If instead η_2 is close to 0, then Retailer 2 is a non-factor and Retailer 1 charges the monopoly price taking the fee into account.

Now we move on to find the optimal fee. Let the smallest optimal fee be denoted by f^*. That such fee exists is guaranteed by the fact the the maximization problem is on a compact set and the profit function is continuous.

At the optimal fee, depending on η_2/η_1, either p_1 is equal to $p^m(c/(1-f^*))$ or it solves (4) with $p_1 < p^m(c/(1-f^*))$. If at the optimal fee $p_1 < p^m(c/(1-f^*))$, then marketplace revenue has to be equal to

$$(1-\eta_2)\pi^m(c).$$

To see this note that in this case since f^* maximizes $fp_1q(p_1)(1-\eta_1)$ so it also has to maximize $\frac{(1-\eta_2)}{(1-\eta_1)}p_2q(p_2)$ (by 4) which we know is maximized at $f^* = f^m(c)$. We know from above, that at the optimal fee the maximand is $(1-\eta_2)\pi^m(c)$. We therefore conclude that the auction has to do better because the auction can attain this level of revenue by setting (a non-optimal) fee of $f = 0$.

In the second case at f^* we have $p_1 = p^m(c/(1-f^*))$. Again, an auction can earn more because at the same (potentially non-optimal) fee the auction generates the same fee revenue (Retailer 1 would set the same monopoly price) and bid revenue equal to the positive profit of Retailer 2 at this fee.

The above discussion leads to the following result:

Proposition 10. *In the model with match value heterogeneity, when the marketplace cannot commit to an algorithm, the marketplace is better off with the auction, consumers might prefer steering by auction or by algorithm.*

Proof. The first statement follows from the discussion in the text. The second is easily proven by an example. □

While the auction dominates the algorithm without commitment from the marketplace's standpoint, for consumers either may be better. This is because both entail double-marginalization. In the case of the auction, the winning Retailer 1 charges a monopoly markup over the marginal cost which is inflated by a positive fee. In case of the algorithm, the winning Retailer 1 may charge lower than monopoly price but the fee is higher so consumers may end up paying more.

B.2 Cost heterogeneity

Let us now turn to the second market power example where retailers are heterogeneous in marginal costs. Without algorithmic commitment, the marketplace always favors a price closest to $p^m(0)$, the revenue-maximizing price. Retailer 2, as assumed, charges

$$p_2 = \max(p^m(0), c_2/(1-f)).$$

As in the analysis with commitment, at the optimal fee $c_2/(1-f) \geq p^m(0)$ has to hold, thus Retailer 1 will serve all consumers at a price

$$p_1 = \min\left\{\frac{c_2}{1-f}, p^m\left(\frac{c_1}{1-f}\right)\right\}$$

Let the smallest optimal fee be denoted by f^*. That such fee exists is guaranteed by the fact the the maximization problem is on a compact set and the profit function is continuous. There are two cases to consider here: At this fee, either (i) $p^m\left(\frac{c_1}{1-f^*}\right) > \frac{c_2}{1-f^*}$ so that Retailer 1 is constrained by the marginal cost pricing by Retailer 2 and $f^* = f^m(c)$ or (ii) $p^m\left(\frac{c_1}{1-f^*}\right) \leq \frac{c_2}{1-f^*}$ so that Retailer 1 prices as if it is a monopolist. In the former case we know that the marketplace's fee revenues will be equal to $\pi^m(c_2)$, which the auction can at least match with a zero fee, but in general can do better. In the latter case with the auction revenue is (weakly) higher because choosing the same fee results in the same downstream price by Retailer 1 and so the same fee revenue and (potentially some additional) bid revenue from Retailer 2.[32]

We thus conclude:

Proposition 11. *When the marketplace cannot commit to an algorithm, the marketplace is better off with the auction, consumers might prefer steering by auction or by algorithm.*

Proof. The first part follows directly from the discussion in the text. The second is easily proven by an example. □

[32] If c_2 is very high compared to c_1 then Retailer 2 may not be able to earn positive profits at the fee set by the marketplace and thus bid zero.

As the proposition makes clear, it is *a priori* unclear which of the two schemes consumers prefer because while the algorithm for a given fee will result in lower price, the fee may be higher with the algorithm.

C Fee customization

We have maintained throughout our analysis an assumption that the marketplace has to charge a common fee across markets and retailers. We believe this is realistic for large marketplaces where the number of markets/retailers makes it very hard for the marketplace to communicate many fees, and may indeed trigger various counter responses from retailers (e.g. setting up multiple accounts to get lower fees). Indeed, interpreting different individual consumers as different "markets"it is hard to imagine any marketplace doing otherwise. Nevertheless, in this section we briefly consider what happens in the other extreme scenario whereby the marketplace can freely customize fees. Of course, it is only relevant in the case of heterogeneity (that is, in extensions beyond the baseline model). It turns out that in some of the settings we have considered this assumption makes no difference, while in others it changes the result drastically.

C.1 Retailer Market Power

Propositions 5 and 6 are not altered even if the marketplace can charge retailer specific fee f_i. To see this, consider Proposition 6 where market power relies on cost heterogeneity. Here, the algorithm is able to extract monopoly profits and this remains so with fee customization. While the auction format benefits from separate fees f_1 and f_2 for the two retailers, full extraction remains elusive. The reason is the following. In order to extract full profits it has to be Retailer 1 who sells the good at the price $p^m(c_1)$, but this price can only be achieved with $f_1 = 0$ given the ex post monopoly power. But given that the marketplace cannot earn fee revenues from Retailer 1, it cannot earn full monopoly profits because the most Retailer 2 will bid is $\pi^m(c_2)$. Optimal fees are such that $f_2 = 0$ so that Retailer 2's bid is maximized, however, f_1 is now constrained by non-negativity of Retailer 1's profits

$$(p^m(c_1/(1-f_1))(1-f_1)-c_1)q(p^m(c_1/(1-f_1)))-\pi^m(c_2)\geq 0$$

Since $f_1 > 0$ is optimal, there is no way to extract $\pi^m(c_1)$ from Retailer 1, thus the auction will do strictly worse than the algorithm.

Proposition 12. *Assume the marketplace can charge per retailer fees in the model with cost heterogeneity, then the algorithm dominates the auction both for the marketplace and consumers.*

Proof. As shown in the proof to Proposition 6, even with a single fee the marketplace can extract $\pi^m(c_1)$ from Retailer 1. Through arguments in the preceding text, the profit from the auction is strictly lower. That consumers will do worse follows from the optimality of $f_1 > 0$. □

A similar result can be derived for the case where market power stems from match value heterogeneity, thus we conclude that fee customization is not crucial for our results regarding market power.

C.2 Heterogeneous Market Conditions

In the case of heterogeneous demand conditions example fee customization will restore parity between auction and algorithm because for each market the fee can be adjusted to extract monopoly profits. However, in the interpretation of that model where the marketplace does not know demand conditions in a particular market, then fee customization will play no role (it cannot be based on demand), while the auction will still extract full integrated profits.

Results on susceptibility heterogeneity also change. In this case, the auction was strictly worse than the algorithm because it was unable to customize fees to susceptible and attentive markets, and with the ability to do so the full parity between the auction and the algorithm is restored.