Monetizing Steering*

Heski Bar-Isaac† Sandro Shelegia‡

November 4, 2022

Abstract

Better-placed products enjoy greater sales. Sellers on marketplace platforms prefer that the marketplace steers customers in their direction. A monopoly marketplace can earn profits through an ad-valorem fee. It determines steering through the design of an algorithm or through an auction that raises further revenues. The platform's preferred algorithm might select the cheapest offering. Still, prices may be high due to fees or auction payments. In a single competitive retail market, the platform can extract monopoly profits fully. It can do so by steering either through algorithm or fees or through auction when setting appropriate fees. More generally, there are trade-offs. The marketplace (and consumers) might prefer either scheme. Specifically, we consider these model variants: retail markets with market power; markets heterogeneous (or uncertain) in demand conditions; and markets heterogeneous in the extent to which consumers are susceptible to steering. In this way, the model can rationalize the use of auctions to determine steering in some markets and algorithms in others. The approach also speaks to discussion of self-steering and highlights that marketplaces enjoy several means of raising revenue. As a result, assessing impacts on consumers requires a holistic view.

Marketplaces and platforms, ranging from Alibaba 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 and the order with which to display them.[1] These choices over how to "steer" consumers have consequences for the sales and profits of the sellers of these products. In turn, marketplaces might seek to charge these sellers for steering consumers towards them rather than their rivals. Commentators on policy have noticed and commented on these incentives. For example, Crémer

*We thank Ramtin Salamant for excellent research assistance and audiences at Cergy Paris Universite, the University of Toronto, the Humboldt University and University of Toronto Theory Conference, and UIUC/USTC Marketing seminar for helpful suggestions.

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

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

[1]For example, even a fairly focussed search on Amazon.ca on June 6, 2022 for "wireless earbud headphones" produced over 60,000 results. On the same date searching for a Saturday night dining reservation on opentable.com in Toronto for a steakhouse yielded 100 results, and so on.

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)

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 platforms do in fact feature privileged positions that are powerful consumer defaults.[3]

While there has been considerable academic and public discussion on self-steering (that is marketplaces directing consumers to products that they produce), we primarily consider an environment where the marketplace does not produce its own goods. Instead, our focus is to explore different methods that marketplaces can use for steering consumers towards sellers' offerings. These choices not only affect the marketplace's profits and the efficiency of matching sellers and consumers, but also the prices that sellers charge and, in this way, overall welfare.

Specifically, we contrast two typical approaches towards steering: using an algorithm or auctioning privileged positions. In both cases we suppose that the marketplace also employs some form of 'revenue sharing' agreement, specifically an ad-valorem fee. This is how the platform earns revenues. In case that the marketplaces allocates steering through an auction it also earns the auction revenues.

It will not surprise the reader that we find that either approach—a steering algorithm designed to maximize the platform's profits or an auction—coupled with a fee that maximizes the marketplace's profits will incentivize sellers to increase prices higher than in the absence of platform fees or auctions. We find that for a given fee, the auction results in higher prices than the optimal algorithm. However, since the marketplace's optimal fee depends on its approach to steering, in our baseline model with a single competitive market where steering is necessary for access to consumers, retail prices would be identical. They would be set at the monopoly level even though there are competing sellers.

An algorithm could 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 relevant retailers low in the search order). Even though this requires commitment and may be costly for a platform, there is evidence that platforms might engage in this kind of behavior.[4] However, in our benchmark model and some variations, at the optimal fee the algorithm can implement a desired price by steering consumers towards whichever seller is cheapest. In this sense, the marketplace may appear to be working in consumers' best

[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."

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

interests through their algorithmic design given the available offers, even if including the fee and the implications for firm pricing, consumers face monopoly prices.[5]

When there is market power in the retail market then the platform is limited in its ability to earn profits from an auction and so relies on fee revenue to a greater extent. The retailer's market power leads to a double-marginalization problem that leads prices to be "too high". The optimal algorithm in this case may involve higher fees but be preferred by consumers as well as by the platform since the algorithm can counter-act the double mark-up by disciplining the retailer's pricing.

It is noteworthy that in the competitive benchmark model with endogenous fees, although an auction and algorithm lead to identical profits and consumer surplus, they do so through different choices of fees: the auction prefers a fee of zero and competition between retailers for the ex-post (undistorted) monopoly profits allow the platform to earn monopoly profits as consumers face monopoly prices. Instead, with the algorithm the platform charges a high fee such that retailers' effective costs (which will also be equal to the retail price) are equal to the (undistorted) monopoly price.

This observation highlights that there are trade-offs away from our benchmark and, in particular, when instead of facing a single consumer market, the platform faces several markets or uncertainty surrounding the market it faces and must choose a fee common across a number of markets or robust to various settings, as is the case in practice.[6] There are many reasons why a platform cannot fine tune fees market by market; for example, it may be costly to do over all products and over time; it may lack the relevant information to optimally do so; doing so may deter entry of small-scale retailers who might find it rather complex to consider different fees for each different product; doing so may raise regulatory scrutiny. These considerations are all the more pronounced on observing that the relevant notion of a "market" may be subtle and may often correspond to a single consumer-search outcome since auctions and algorithms may depend on specific consumer characteristics or histories.[7] In this paper, we simply take it as given that the platform cannot fine-tune its fees and consider the consequences. [8]

As our benchmark model makes clear, steering through auctions has the advantage that it involves the same optimal fee (0) across different markets with different demand conditions; whereas an algorithm would prefer different fees in markets with different demand considera-

[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 from "unfair ranking" (see CMA (2020) Section 2.1.4). However, the gatekeeper role of the platform leads consumers to suffer monopoly prices.

[6] For example, as of June 6, 2022, https://sellercentral.amazon.ca/gp/help/external/200336920 Amazon.ca listed 23 categories for referral fees that range from 6% to 45% with the model fee at 15%. These categories are broad for example "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.

[7] Certainly, there is scope for platforms to personalize the order in which products are displayed, and a 2018 European Commision study found that 61 percent of the 160 ecommerce sites personalised search results rankings.

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

tions. This consideration clearly advantages the use of an auction. However, there are markets where an auction is less effective. As discussed above, markets where retailers have market power are one such example. Another important case is where the retail market includes consumers who 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. However, in the case of an auction, the higher fees that the presence of such markets entails, imply more pricing distortions (through double marginalization). Instead, for the algorithm absent different demand conditions, the optimal fee for such markets is similar to that as in the baseline case and the pricing distortion does not arise.

Thus overall, with heterogeneous markets and common fees, either an auction or an algorithm would be preferred. Indeed, the platform can, trivially, do better with a mix— allocating some steering through auctions in some markets and algorithm in others. Indeed, in many marketplaces, this does indeed seem to be the case.[9] 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 (SSH: simple models?) that highlight one force at a time and, thereby, consider the implications for the platform and for consumers.

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

0.1 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) and Chen and He (2011) or more recently Anderson and Renault (2021), Teh and Wright (2022) and a broader literature considers the extent to which intermediaries might bias recommendations, for example, to earn higher commissions (Jullien and Wright 2011, Inderst and Ottaviani, 2012). These papers typically consider one form of allocating prominence (often a commission, or an auction) and often focus on different questions; for example, whether the right product is recommended when retailers are heterogeneous (a recent exam-

[9]For example, Gutierrez (2021) suggests that on average that of the price that a consumer pays 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.

ple where in addition consumers might form incorrect expectations is Heidues, Koster and Koszegi, 2022). Instead, this paper is focused on different methods for selling prominence and compares them, albeit in an environment where 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.

A more recent, exciting, and rapidly growing literature, both theoretical and empirical and in line with contemporary policy discussion, has focused on the role of platforms as sellers and, in particular, whether marketplaces should be allowed to offer their own products, and how this might bias steering. Theoretical papers in this literature include Anderson and Bedre-Defoile (2021,2022), de Cornière and Taylor (2014), Etro (2021a, 2021b), 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). Empirical evidence on such self-preferencing and its consequences include Guiterrez (2021), Lee and Musolff (2021), Lam (2021) and Raval (2022). We highlight that even in the absence of self-preferencing, platforms with market power have varied means of earning revenues and that these alternative means operate in somewhat different ways from each other. In Section 5 we show how this perspective has implications for steering.

Indeed, we abstract from many of the considerations of this literature. Retailers are ex-ante identical (so recommending the "wrong" retailer would only involve a high-priced one). We consider a rather extreme form of prominence whereby (some) consumers only consider 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—platforms' choices of *how* they monetize steering. In our reading, we have come across two recent papers that examine this question, albeit with slightly different lenses. Ciotti and Madio (2022) consider a single retail market where 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 platforms 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 retailers on Alibaba in comparison to Amazon suggests that outside options are more attractive and, in their model, this leads to the use of auctions (and lower platform revenues). Instead, our model can interpret this as reduced retailer market power leading to greater use of auctions (and need not imply lower platform revenues).

1 Benchmark Model and the Equivalence of Steering by Auction or by Algorithm for a Single Competitive Retail Market

We consider a single retail market. There are two (or more) identical retailers who can produce goods at a marginal cost c. In order to access consumers with demand given by $q(p)$, these retailers require steering from a platform. Specifically, suppose that consumers only consider a single retailer and the platform can determine which retailer consumers observe (if any).[10] We call this the privileged position.

The platform determines both its fee f—a proportion of retailer revenues that retailers pay the platform to appear on the platform—and its method to determine steering.[11] We consider two methods. First, the platform can auction the privileged position. 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.[12] Alternatively, the platform 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 platforms is effectively committing to withhold the access 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 platform given retailer prices.

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

1.1 Exogenous fees

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

In case the privileged position is allocated by auction, the platform has no decisions to make. In choosing their prices, retailers would treat bids as sunk costs and would choose prices

[10] A similar assumption is made in Heidhues, Koster and Koszegi (2022) for example.

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

[12] For example, if the auction was on a pay for impression basis or pay-per-click where all consumers clicked, the platform would obtain identical revenues and prices would be identical.

to maximize their profits. That is both firms 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 corresponds 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.

It is immediate that the identical retailers will bid up to their anticipated maximized profits—that is $(1-f)\pi^m\left(\frac{c}{1-f}\right)$ for the privileged position.[13] Consequently, when allocating by auction with a fee given by f, the platform will earn:

$$\left((1-f)p^m\left(\frac{c}{1-f}\right) - c\right) q\left(p^m\left(\frac{c}{1-f}\right)\right) + fp^m\left(\frac{c}{1-f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right), \tag{2}$$

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 platform 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 platform chooses, retailers will choose prices. It is immediate that the platform 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 platform 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 platform's problem becomes choosing p^* to maximize $fp^*q\left(p^*\right)$ such that

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

If the constraint does not bind then the platform 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 platform tries to impose the revenue-maximizing price $p^m(0)$. It is immediate that the platform will choose the lowest price possible that satisfies the retailers zero profit constraint. 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 auction leads to retail prices $p^m\left(\frac{c}{1-f}\right)$ and platform 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 platform revenues $fp^*q(p^*)$.*

[13] While in principle there may be other equilibria, this is both natural and can be selected as a unique equilibrium through a trembling-hand argument.

Proof. Follows from the discussion in the text. □

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

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 $p^m\left(\frac{c}{1-f}\right) > p^m(0)$, thus $p^* = \max\left\{p^m(0), \frac{c}{1-f}\right\} < p^m\left(\frac{c}{1-f}\right)$. □

Although consumers prefer allocation by algorithm with a fixed fee, the platform 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 $c/(1-f)$:

$$f^m(c) \equiv \frac{p^m(c) - c}{p^m(c)} \tag{3}$$

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

Proposition 2. *The platform 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 platform might prefer to steer by auction or by algorithm.*

Proof. If $f = 0$ then the platform 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 platform 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 platform recovers in part through fee revenue and in part as auction receipts) is strictly less than the full monopoly profit. For sufficiently close fee to $f^m(c)$, the algorithm has to do better than the auction. □

We should note here that in the case that the platform 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. 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.

Corollary 2. *Prices paid by consumers and profits earned by the platform are the same when using the algorithm regardless of whether it has or does not have commitment power.*

Proof. Without commitment, the platform allocates the slot to the firm whose price is closest to $p^m(0)$, thus with two or more firms, at least two in equilibrium will set $p = \max\{p^m(0), c/(1-f)\}$ while remaining firms may set higher prices. Therefore, prices paid by consumers are the same with or without commitment, and the platform earns the same fee revenues. □

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 .

1.2 Endogenous fees

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

Proposition 3. *When fees are endogenous the platform can earn monopoly profits—that is $\pi^m(c)$—when steering by auction or by platform.*

Proof. Following Proposition 1 when the platform sets $f = 0$ then the auction earns $\pi^m(c)$ When the platform 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-platform could earn. Although the platform 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 platform 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 platform can obtain the same profit setting $f = f^m(c)$ and allocating the privileged position to whichever retailer has a lower priced. This yields the following result.

Corollary 3. *The platform 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 Proposition 3 and Corollary 2. □

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 platform in the choice of auction or algorithm for steering.

2 Retail Market Power

If a retailer has market power within a market then allocating steering through an auction earns the platform less than doing so by algorithm. This can be easily understood and demonstrated by considering the case where the source of such market power is match value probabilities.[14]

We assume that with some probability η_i firm i's product is a bad match in which case resulting demand is zero. Bad match probabilities are assume to be independent across retailers, and neither the retailers nor the platform know realization of this uncertainty. The platform knows the match value probabilities, as do retailers. Without loss, let $\eta_2 \geq \eta_1$.

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 platform can induce retailer 1 to set $p^m(c)$, and allocate all consumers to it if retailer 1 complies. In this case the platform can set $f = f^m(c)$ and extract profits equal to

$$(1-\eta_1)\pi^m(c). \tag{4}$$

This is the maximum attainable profit that even an integrated monopolist could achieve since consumers observe only a single offering. If instead the platform steers via auction, then both retailers will set the monopoly price given the fee, retailer 1 will win, and will pay the bid of retailer 2 which would be equal to its profits if it won the auction. Thus the platform earns

$$(1-\eta_1) f p^m\left(\frac{c}{1-f}\right) q\left(p^m\left(\frac{c}{1-f}\right)\right) + (1-\eta_2)(1-f)\pi^m\left(\frac{c}{1-f}\right). \tag{5}$$

Eq (5) 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.[15]

The optimal fee for the auction is positive whenever $\eta_2 > \eta_1$ because in (5) the second

[14] There are, of course, many ways to consider market power. For example, Cotti and Madio (2022) examine the case where one firm is of higher quality than the other .

[15] If bidders were to pay per click, and clicks would only occur upon a good match, then to retain our formula one needs to introduce quality adjustment (by click-through rate) as in Varian (2007). In our setting the click-through rate is $1-\eta_i$, and the winner (retailer 1) would have to pay $\frac{1-\eta_2}{1-\eta_1}$ times the bid of retailer 2, equal to $(1-f)\pi^m\left(\frac{c}{1-f}\right)$, which when multiplied by the number of clicks for retailer 1, given by $1-\eta_1$, gives $(1-\eta_2)(1-f)\pi^m\left(\frac{c}{1-f}\right)$ as in the following analysis for pay per impression case.

term is maximized at $f = 0$ and the first term is maximized at $f > 0$, thus the optimal $f > 0$. 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 4. *Both the platform and consumers are better off with the algorithm.*

It is noteworthy that commitment is substantive for this result. If the platform cannot commit to the form of the algorithm then it is better off with the auction and there is no clear ranking for consumers.

3 Markets Heterogeneous in Demand 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 platform 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. As in our benchmark model, 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 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 platform 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 platform 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 platform employing an algorithm becomes choosing $p_1, ..., p_N$ and f in order to maximize

$$f \sum_{i=1}^{N} p_i q(p_i) \text{ such that } (1-f)p_i \geq c_i \text{ for all } i. \tag{6}$$

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

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

Proposition 5. *When retail markets are competitive and heterogeneous in demand conditions, the platform prefers steering by auction. Consumers may prefer steering by auction or steering by algorithm.*

Proof. The first statement is proven above. To prove the second, examples suffice. This is left as a simple exercise or the reader can contact the authors. □

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

4 Markets Heterogeneous in 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 markets where all consumers observe all the retailers.[16] In such "attentive" markets the algorithm and the auction play no role: consumers simply purchase from whichever retailer offers a lower price.

In these attentive markets, the platform 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 with an endogenous fee as in Proposition 3.[17] 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 platform treats these markets identically.

Instead, since the auction imposes no discipline on the market power of the retailer who wins the auction, in case the market is one with susceptible consumers, the retailer charges the monopoly price $p^m\left(\frac{c}{1-f}\right)$ given its effective marginal cost. Instead, in an attentive market Bertrand competition leads prices equal to the effective marginal cost $\frac{c}{1-f}$. Thus when steering by auction and facing attentive and susceptible markets, the platform would prefer different fees for these different markets (0 for the susceptible market and $f^m(c)$ for the attentive one).

Thus, while the algorithm can optimize by using the same fee $f^m(c)$ across attentive and susceptible markets (with an allocation that leads retailers to charge their effective marginal cost $p^m(c)$ in the susceptible market), the auction must compromise between the optimal fee for the attentive and susceptible markets and choose an intermediate fee in between 0 and $f^m(c)$. In particular, when choosing the auction, the platform would choose a strictly positive fee and as a result prices in susceptible markets would be higher than the platform prefers and higher than consumers would face under steering by algorithm. Instead for the attentive markets, since the fees are below $f^m(c)$, prices would be lower under steering by auction than steering by algorithm. For overall consumer surplus, either effect may dominate depending on the mix of attentive and inattentive and other parameters.

We summarize this discussion as follows.

Proposition 6. *When demand conditions are identical across all retail markets differing only*

[16] The assumption that all consumers in a retail market are either susceptible or they are not is clearly a simple analytical case that illustrates economic forces crisply.If each market constitutes a different individual, this assumption might be considered a directly relevant case.

[17] Moreover, just as in Corollary 3, 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.

in that some comprise only attentive consumers and others only susceptible consumers, the platform prefers steering by algorithm (with or without commitment). Consumers may prefer steering by auction or steering by algorithm.

Proof. The first statement is proven in the discussion above. The second statement can be proven by example. □

Of course, a natural interpretation of the same result is to suppose that the platform is uncertain about whether a market is susceptible or attentive. For example, rather than supposing that there are s markets that are susceptible and a markets that are attentive, a single market that the platform understood to be susceptible with probability $\frac{s}{s+a}$ and otherwise attentive would lead to qualitatively similar results. Another interpretation is that consumers are in the same market, but retailers are able to price discriminate between attentive and susceptible.

Perhaps a more interesting variation allows for a mixture of attentive and inattentive consumers within a single market. Following the discussion above it is immediate that steering by algorithm to a lowest-priced retailer and setting an optimal fee (in this case $f^m(c)$) allows an algorithm to earn monopoly profits. Instead, steering by auction does not perform as well for the platform. The characterization of this scenario is a little involved. In equilibrium, retailers would choose mixed price strategies whose characterization depends on the demand curve, costs, fees and the share of attentive consumers. For particular parameters, this can easily be done; however, a general characterization is not analytically tractable, though it is immediate that the platform would earn less in steering by auction than steering by algorithm and that consumers may prefer either scheme.

5 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 has been described, 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.

Of course, another traditional means to eliminate double marginalization is vertical integration. In our context, this would involve the platform operating as a retailer—perhaps by directly purchasing a retailer, or by developing its own production capabilities. As described in Section 1.2, with a single competitive retail market, the platform can earn monopoly profits through auction revenues alone and there is no reason to vertically integrate. Instead, the prospect of double marginalization can arise when fees are positive and the platform uses an auction. Following the analysis above this may be the case if some markets feature retail market power or attentive consumers, and other do not and there is heterogeneity in demand conditions. In this way, the opportunity to vertically integrate can substitute for the ability to fee discriminate across markets.

Vertical integration can also be useful for the platform even when it only steers by algorithm if the optimal fee differs across different retail market but the platform is constrained to charge a common fee. Following the analysis above, this would be the case when markets are heterogeneous in their demand conditions.

Rather than analyze all variations and to move away from familiar double mark-up concerns, we present a simple example that illustrates the possibility that allowing the platform to vertically integrate can raise or lower fees in the other market and can raise or lower consumer surplus.

Example 1. *Consider two markets $q_1(p) = 1 - p$ and $q_2 = 2 - p$ and say $c_1 = c_2 = 0.25$ then requiring a common fee would lead to a fee of $f = 0.69$ and consumer surplus of 0.52 Instead, suppose that the platform integrates in the first market then the fee for the second market would be higher $f - 0.78$ and consumer surplus would be lower 0.45 Now suppose, instead that the second market had demand $q = 1 - 2p$ then the common fee would be $f = 0.6$ and associated consumer surplus 0.07 but vertical integration in the first market would now lead to a lower fee $f = 0.33$ and higher consumer surplus 0.09.*

6 Conclusions and Policy Considerations

This paper has presented an almost laughingly simple model in which consumers only observe the single retailer that the platform privileges and we consider only a single model variation at a time.[18] 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 and policy makers interested in platforms. First, that platforms have several tools to monetize their gatekeeping power. We have shown this by making this power extreme (in supposing that susceptible consumers observe only a single retailer). In particular, we have shown that in our baseline setting, whether steering by auction or steering by algorithm, a platform can extract full monopoly rents. However, these different methods extract rents in different ways and so respond differently to different settings. As seem relevant in practice, we suppose that fees are common across markets and find that steering by algorithm is better-equiped 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

[18]Of course, in practice there are markets that vary in susceptibility, market power, and demand conditions all at the same time, and the platform may choose to use steering by auction 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.

fees views susceptible and attentive consumers similarly). Instead, steering by auction with less direct control over prices can suffer from double-marginilazation and in our benchmark model optimizes by setting minimal fees; consequently, it is well-equipped for heterogeneity in demand conditions.

These simple observations have immediate policy implications. Again, though perhaps trivial, these seem not much discussed. First, even absent steering concerns (whether to their own products or to the "wrong" ones when there are heterogeneous retailers) platforms have the ability to effect retail market outcomes and consumer welfare. Second, the platform 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. Similarly, regulation on the ability of the platform to integrate into retailing and to prevent self-steering, could affect fees and, therefore, outcomes in other retail markets that share common fees.

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

[2] Anderson, Simon and Özlem Bedre-Defolie (2022) Online trade platforms: hosting, selling, or both?

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

[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 G (2011) Position auctions with consumer search. Quart. J. Econom. 126(3):1213–1270.

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

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

[10] CMA (2017) "Online Search: Consumer and Firm Behaviour—A Review of Existing Literature"

[11] CMA (2020) 'Algorithms: How they can reduce compeition and harm consumers."

[12] de Cornière A, Taylor G (2014) Integration and search engine bias. *RAND Journal of Economics*, 45(3):576–597.

[13] Cremer J., de Montjoye, Y.-A., and Schweitzer, H. (2019) Competition Policy for the Digital Era, report for the European Commission (DG Competition).

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

[15] Federico Etro et al. (2021) "Hybrid marketplaces with free entry of sellers," Technical report, Universita'degli Studi di Firenze, Dipartimento di Scienze per l'Economia.

[16] 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

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

[18] Furman, J., Coyle, D., Fletcher, A., McAules, D., and Marsden, P. (2019). Unlocking digital competition: Report of the digital competition expert panel. Report prepared for the Government of the United Kingdom

[19] Gutiérrez, Germán (2021) "The Welfare Consequences of Regulating Amazon " mimeo

[20] Hagiu A, Jullien B (2011) "Why do intermediaries divert search?" *RAND Journal of Economics*, 42(2):337–362.

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

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

[23] Andrei Hagiu, Tat-How Teh, and Julian Wright (2021) "Should platforms be allowed to sell on their own marketplaces?" Forthcoming in the *RAND Journal of Economics.*

[24] Andres Hervas-Drane and Sandro Shelegia (2021) "Retailer-led marketplaces" mimeo.

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

[25] Roman Inderst and Marco Ottaviani (2012) “Compeititon through Commissions and Kick-back,” 102(2), 780-809.

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

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

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

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

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

[31] 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/.

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

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

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

[35] Teh, Tet-How and Julian Wright J (2022) Intermediation and steering: Competition in prices and commissions.Am. Econom. J.Microeconom. Forthcoming.

[36] Tremblay, Mark J. (2022) The Limits of Marketplace Fee Discrimination, working paper

[37] Yusuke Zennyo. (2020) “Platform encroachment and own-content bias” Available at SSRN 3683287.

A Per-unit fees

In this appendix we suppose that the platform can charge a per unit fee u in addition to raising revenues from a proportionate fee and possibly through auction revenues. First, we argue that when there is a single market or when the platform can collect full monopoly profits (heterogeneous markets with steering by auction, and heterogeneous susceptibility with steering by algorithm) then introducing this additional instrument does not change our conclusions. We briefly review these observations in the benchmark model and variations presented in the main paper.

A.1 Single Competitive Retail Market Benchmark

Here we review the analysis in Section 1.1. 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 platform earns

$$(1-f)\pi^m\left(\frac{c+u}{1-f}\right) + f\left(p^m\left(\frac{c+u}{1-f}\right) + \frac{u}{f}\right)q\left(p^m\left(\frac{c}{1-f}\right)\right). \tag{7}$$

As in Section 1.1, we consider the platform's problem given exogenous fees to be choosing p^* to maximize $f\left(p^* + \frac{u}{f}\right)q(p^*)$ such that $[(1-f)p^* - c - u]\, q(p^*) \geq 0$.

Again, it is immediate consumers prefer allocation by algorithm as does the platform 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 1.2, when fees are endogenous, the platform can optimize in many ways. As we have already seen when the platform sets $f = u = 0$ then the auction earns $\pi^m(c)$ and when the platform sets $f = \frac{p^m(c)-c}{p^m(c)}$ then the optimal algorithm earns $\pi^m(c)$. The platform 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 platform 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 3) it would choose to set $u = f = 0$. Instead, when it does so by algorithm (as in

Sections 2 and 4 which consider market power and susceptibility) then as above, the platform has flexibility in choosing combinations of per unit fee u and revenue shares f that can achieve this.

When the platform cannot extract full monopoly profits: for example, if there are different conditions across different markets and the platform is constrained to steer by algorithm, then the ability to use both a per unit fee u and a revenue share f would allow it to earn more than if only one of these instruments were available.

Appendix II: Retailer market power with no commitment

Here, we turn to the case where the platform cannot commit to an algorithm. Thus the platform, *ex post*, chooses the seller whose price is most profitable in terms of fees. To simplify, we assume that when indifferent the platform chooses Retailer 1.

The platform will choose Retailer 1 if

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

and will otherwise choose Retailer 2.

In equilibrium retailer 2 cannot set $p_2 > \max(p^m(0), c/(1-f))$ because retailer 1 would set such a price as to just win the spot, but then retailer 2 could reduce p_2 and win. So in equilibrium

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

and p_1 solves

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

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

Now we move on to find the optimal fee. Clearly if $c/(1-f) < p^m(0)$ then the platform can increase revenue by increasing f (this would increase fees without changing p_1), so $c/(1-f) \geq p^m(0)$ must hold, and therefore $p_2 = c/(1-f)$. Since the platform aims to maximize $f p_1 q(p_1)(1-\eta_1)$, then it has to maximize $\frac{(1-\eta_2)}{(1-\eta_1)} p_2 q(p_2)$, which we know is maximized at $f^* = f^m(c)$. Thus, whithout commitment the algorithm earns the platform revenue equal to

$$(1-\eta_2)\pi^m(c), \tag{8}$$

so the platform can extract the worse retailer's monopoly profits.[19] We therefore can state:

Proposition 7. *When the platform cannot commit to an algorithm, the platform is better off with the auction, consumers might prefer steering by auction or by algorithm.*

While the auction dominates the algorithm without commitment from the platform'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

[19]The revenue is even worse if η_2 is small so that Retailer 1 is an unconstrained monopolist.

cost which is inflated by the positive fee. In case of the algorithm, Retailer 2 charges the monopoly price for zero fee, but Retailer 1 can charge a higher price when $\eta_1 < \eta_2$. If firms are symmetric, then under the auction, the optimal fee is zero, while with the algorithm also charges $p^m(c)$, thus consumers and the platform are indifferent between both mechanisms.

$c_1 < c_2$, respectively. In this case an auction will raise the monopoly profits of the higher cost firm as auction proceeds; that is it would earn