Monetizing Steering*

Heski Bar-Isaac† Sandro Shelegia‡

September 3, 2024

Abstract

A monopoly marketplace chooses an ad-valorem fee associated with transactions. It can steer consumers toward products either by using an algorithm or an ad auction that raises further revenues. In a single competitive retail market, the marketplace can extract full monopoly profits with either scheme, although the optimal fee depends on the steering method. Generally, we show when the marketplace (and consumers) strictly prefer steering by algorithm or by ad auction. The marketplace prefers ad auctions when markets vary in demand/cost conditions, but algorithms when sellers on the marketplace have market power, or consumers vary in susceptibility to steering. We also consider policy implications associated with vertical integration and steering with both ad auctions and algorithms.

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 and the order in which to display them. These choices of 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 to them rather than to their

*We thank Sijie Lin, Luca Orsini, Krithika Ragupathi and Ramtin Salamant for excellent research assistance, Gary Biglaiser, Qihong Liu, Devesh Raval, Jen Redmond, Michelangelo Rossi, Andre Stenzel, and audiences at Cergy Paris Université, 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, the Competition and Markets Authority, the Competition Bureau, the Bank of Canada, Queen's, Frankfurt and the 2nd Antitrust Economics and Competition Policy Workshop, Rochester, for helpful suggestions. Shelegia acknowledges financial support from the BBVA Foundation, the Spanish Ministry of Economy and Competitiveness, and the 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

rivals. Commentators on policy have noticed these incentives. For example, Crémer et al. (2019) write that "a dominant platform could have incentives to sell 'monopoly positions' to their business users" (p.6); and Scott-Morton et al. (2019) state that "[p]latforms 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) More recently, the US Federal Trade Commission and 17 states issued a complaint alleging antitrust abuse by Amazon (Federal Trade Commission, 2023) amid reports of Amazon taking a large and increasing share of third-party revenues, with 15% coming from advertising (comparable to commission fee revenue); instead, ten years ago, there was negligible advertising revenue (Mitchell, 2021; Marketplacepulse 2023).

In this paper, we explore different methods that marketplaces can use for steering consumers towards retailers' offerings together with the fees that the marketplaces charge. 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, thus, overall welfare. Specifically, we contrast two typical approaches to steering: using an algorithm or using ad auctions 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.[1,2] This is how the marketplace earns in both regimes; if the marketplace steers through an ad auction it also earns the auction revenues. These differing monetization strategies are widely and variously employed: Google, of course, monetizes primarily through advertising; Airbnb through commission fees; Expedia through both, and Amazon on its marketplace began by monetizing through commission fees but as discussed above has seen ad revenue grow at a much faster rate over the last few years.[3]

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 many studies have shown to be the case.[4,5] This also highlights that some marketplaces do, in fact, feature privileged positions that are powerful consumer defaults.[6] Our approach allows us to highlight several key forces that

[1]That is the agency model explored in Johnson (2017), for example. This contractual form that fixes agency fees across all sellers and may involve additional auction revenues is therefore a little distinct from the literature on slotting allowances as described in Shaffer (1991), Lariviere and Padmanabhan (1997), and Desai (2000) where sellers compete in their contract offers.

[2]In our model we consider an ad-valorem fee—this corresponds to the commission fees of many marketplaces but may also reflect fees for service (such as fees for Fulfillment by Amazon) that some have alleged are, in effect, required in order to be featured.

[3]See `https://www.google.com/intl/en_ca/retail`, `https://www.airbnb.ie/e/commissione-semplificata-guida` and `https://www.statista.com/statistics/269382/expedias-revenue-distribution-by-product-category/` accessed on June 13, 2024. Figure 2 in Gutierrez (2021) also documents that third party fees as a fraction of third party revenue have remained fairly constant in the last ten years while the ad share has grown dramatically.

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

[5]As we discuss below, this consumer behavior can also arise naturally as equilibrium behavior in our setting

[6]See, for example Lee and Musolff (2021), who highlight the importance of the Amazon Buy Box, or RepricerExpress (2020) which states that "83 per cent of all Amazon sales happen through the Buy Box and even more on mobile."

would also be present in a richer model in which the marketplace manages several positions of varying visibility.

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, the marketplace will choose different fees when steering is a result of an auctioned ad or steering is by algorithm.

In our baseline model with a single competitive market, in which 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 regulations that focus on a single aspect may achieve little, as marketplaces can use 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 retailer 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.[7] Furthermore, in our benchmark model and several variations, at the optimal fee the algorithm can, in fact, 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.[8] Moreover, as it is easy to communicate this direct-towards-cheapest-offering algorithmic steering to retailers, it 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.

On the one hand, the ad auction prefers a fee of zero, and competition between retailers for the ex-post (undistorted) monopoly profits allows 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.

[7]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. Hunold et al. (2022) show that a seller may be removed from Buy Box for reasons outside of the marketplace.

[8]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.

On the other hand, 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 optimally "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 steering method. As our benchmark model makes clear, steering through ad auctions has the advantage of involving 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 such case 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 markets in which consumers observe all offerings regardless of ranking and so are not susceptible to steering (in this case, 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, in the absence of different demand conditions, the optimal fee for such markets is the same as in the baseline case and the pricing distortion does not arise.[9]

One clear reason that 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.[10] Thus, there may be considerably more markets than even the vast number of available products that feature in prominent digital marketplaces. If everything—that is, the choice of steering mechanism and fees—depends on the characteristics of each consumer/product pairing and the search term, this would correspond to our benchmark

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

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

model. In principle, "fine-tuning" to this level would appear to be technologically feasible and, clearly, is optimal.[11] In practice, it does not occur in most marketplaces.[12]

In this paper, we simply take it as given that the marketplace cannot fine-tune its fees and consider the consequences.[13] This seems both realistic and common. In practice, there may be many rationales for such limits: for example, marketplaces may lack the appropriate information themselves; 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—steering through ad auctions in some markets and algorithms in others. In many marketplaces, in fact, this does indeed seem to be the case.[14] What is optimal depends on the mix of different 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. However, in light of policy interest suggesting that a move from relying on algorithmic steering to advertising may harm consumers, in Section 6.2, we illustrate that allowing for a mixture might benefit consumers through the impact on endogenous fees.[15]

After discussing some related literature, we present our results and discuss key intuitions before highlighting further implications for policy. In particular, in light of considerable policy interest in understanding biases and distortions created by ranking algorithms (CMA 2020, 2022), we note that even when constrained to an algorithm that features whichever retailer maximizes consumer surplus, a marketplace can still earn monopoly profits. In circumstances where it cannot, constraining the marketplace to use this kind of "consumer-friendly" algorithm may lead to lower consumer surplus than an algorithm that maximizes marketplace profits.

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

[12]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" etc., 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 could imagine in principle many more combinations of feasible consumer characteristics and search terms.

[13]See also Tremblay (2022), who considers the role and limits of fee discrimination.

[14]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% are advertising fees (and 20% for 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).

[15]The FTC and 17 state attorney generals in a recent complaint against Amazon allege that it has engaged in illegal conduct that includes "Degrading the customer experience by replacing relevant, organic search results with paid advertisements" on `https://www.ftc.gov/news-events/news/press-releases/2023/09/ftc-sues-amazon-illegally-maintaining-monopoly-power`

1.1 Related Literature

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

First, there is a literature focused on prominence in consumer search initiated by Arbatskaya (2007) and Armstrong, Vickers, and Zhou (2009). Several 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 examines and compares different methods for allocating prominence, 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, that 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, on whether marketplaces should be allowed to offer their own products, and how this might bias steering. Theoretical papers in this literature, including 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), Zennyo (2022), and Choi, Kim and Mukherjee (2023). Long and Amaldoss (2022) while also studying a platform's role as a retailer, consider position auctions in addition to 'organic' listing of differentiated products. They assume that sellers' marginal costs are zero; as a result, effects arising from double marginalization are absent from their analysis, while they play a key role in ours. Further, when no auction is present, they assume that the top slot is allocated randomly, while we assume that the platform uses its commitment power to affect seller pricing in this mode too. Zou and Zhou (2023) study a related model to ours but, much like Long and Amaldoss (2022), study the zero marginal cost case and not allocation through the auction mechanism, however.

Guiterrez (2021), Lee and Musolff (2021), Lam (2021), Hunold et al (2022) 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 6.1, we show how this perspective has implications for steering and return to discussing 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 does), 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 with two identical retailers who can produce goods at a marginal cost c.[16] In order to access consumers with demand given by a well-behaved function $q(p)$,[17] these retailers require steering from a marketplace. Specifically, suppose that consumers consider only a single retailer and the marketplace can determine which retailer consumers observe (if any).[18] We call this the privileged position.

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

[16]Nothing of substance changes if we assume that there are $n > 2$ retailers.

[17]Specifically, as is standard, we suppose that $(p-c)q(p)$ is strictly quasi-concave for all c such that $q(c) > 0$.

[18]A similar assumption is made in Heidhues, Koster, and Koszegi (2022), for example. Although, for ease of exposition, we simply assume that consumers only consider a single retailer, in the settings that we consider this is easy to rationalize as equilibrium behavior in a standard consumer search problem in the style of Diamond (1971): Since consumers face identical or worse offers lower down the list they need not look beyond the first item if they incur even an infinitesimal search cost to do so (with the exception of the model variation considered in Section 3.1 where a non-negligible search cost is required). As discussed above, this also appears consistent with the behavior of most consumers with respect to the Amazon Buy Box.

pay the marketplace to appear on the marketplace—and its method to determine steering.[19] 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 in which retailers bid a lump sum for the privileged position.[20,21] To avoid uninteresting equilibrium multiplicity, we assume that sellers bid their profits in the auction.[22]

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 marketplace is effectively committing to prevent any retailer from matching with consumers with a positive probability, which entails a loss of fees. Second, even when $\alpha_1 + \alpha_2 = 1$, the individual allocation probabilities may not be revenue-maximizing for the marketplace, given retailer prices.

One natural algorithm to consider is one that we call the *(ex-post) "consumer-friendly" algorithm.* This algorithm allocates the privileged position to whichever retail offering affords the higher consumer surplus—if offerings are otherwise identical this would be the lower-priced one as long as it offers non-negative surplus. For many cases we consider, the platform can earn full monopoly rents through the use of this "consumer-friendly" algorithm, and, in Section 3, we will see that this algorithm can make both consumers and the marketplace worse off than they would be with other algorithms.

We suppose that the timing is such that the marketplace first chooses the fee and the allocation method (including the form of the algorithm). Retailers observe these choices and then simultaneously choose prices and bids (if relevant).[23] The privileged position is allocated. Consumers then make purchase decisions, and fees are collected.

[19] 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) to allow the marketplace to earn more.

[20] 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.

[21] We do not consider reserve prices. In cases where the auction leads to the marketplace earning monopoly profits, they play no role. In cases where the algorithm is unable to "fine-tune" the fee to vary market by market, then allowing the auction to fine-tune a reserve price in this way can naturally allow it to outperform algorithmic allocation.

[22] Without this assumption, everything we describe would remain an equilibrium but there will be other equilibria where a seller might bid more than their expected profits knowing that they will lose, akin to homogeneous Bertrand competition with asymmetric costs where in principle a higher cost firm can price below its cost (and above the cost of the rival) knowing that the rival will price just below this in equilibrium. As in that standard case, such equilibria are typically considered unreasonable and a vanishingly small amount of uncertainty can select the natural equilibrium uniquely.

[23] 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.

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.

If the privileged position is assigned 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 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.

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 the allocation of 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 over leaving 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 would be violated if the marketplace tried 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. [24] 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.

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

Proof. Trivially, $p^m\left(\frac{c}{1-f}\right) \geq \frac{c}{1-f}$ and $p^m\left(\frac{c}{1-f}\right) \geq p^m(0)$, with strict inequality in either case if at $c/(1-f)$ the demand is strictly positive. □

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 at marginal cost c.

Proposition 2. *The marketplace 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. Since profits are continuous, the auction is preferred if fees are low enough. 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 $p^m(c)$. Thus, 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. □

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

2.1.1 The consumer-friendly algorithm

We define the consumer-friendly algorithm as the one that allocates the privileged position to whichever retailer offers a higher consumer surplus. It is immediate that if the marketplace uses this algorithm, retailers compete à la Bertrand for the privileged position. As a consequence, the retailers will both set a price equal to $\frac{c}{1-f}$ in equilibrium and it is immediate that consumers (weakly) prefer the use of such an algorithm to any other that the marketplace might choose and prefer it to allocation by ad auction when fees are exogenous.

Note, also, following Proposition 1 that (for high-enough fees) this is the price that an otherwise unconstrained marketplace would want an algorithm to induce. It also follows by the proof of Proposition 2 that at $f = f^m(c)$, the marketplace extracts the full monopoly profit. This highlights that though the algorithm may be consumer friendly, the outcome also depends on the fee. We return to this theme below when we discuss endogenous fees. First we summarize this discussion as follows.

Proposition 3. *Consumers prefer the consumer-friendly algorithm both to any other that the marketplace would choose and to the use of an ad-auction. At $f = f^m(c)$ the marketplace earns the full monopoly profit with the consumer-friendly algorithm, and, more generally, for f high enough, it can attain its maximal profits among all algorithms with the consumer-friendly algorithm. In general, it might prefer to steer by consumer-friendly algorithm (when restricted to choosing only this algorithm) or by auction.*

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 fees are endogenous, 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. The marketplace can also extract full monopoly profits using the consumer-friendly algorithm.

Proposition 4. *When fees are endogenous, the marketplace can earn monopoly profits—that is, $\pi^m(c)$—when steering by auction or by algorithm. This is also the case when the marketplace must use the consumer-friendly algorithm.*

Proof. Following Proposition 1, when the marketplace sets $f = 0$, the auction earns $\pi^m(c)$. When the marketplace steers by algorithm, it dictates a price of $p^m(c)$ and sets a fee so that retailers earn no profits–that is, $f = f^m(c)$–the optimal algorithm earns $\pi^m(c)$. The marketplace can also earn $\pi^m(c)$ with the consumer-friendly algorithm, as described in Section 2.1.1. □

Note that this level of profit is the highest attainable. It is the most that a fully integrated retailer-marketplace can earn. Although the marketplace can achieve this level of earnings either with steering or with the algorithm, and the retail prices in both cases will be identical,

the two schemes involve different fees. The optimal auction attains it with a fee of 0, whereas a marketplace attains it with a fee of $\frac{p^m(c)-c}{p^m(c)}$.

We now move on from our benchmark model of a single competitive market in which 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. In particular, in Section 3 we argue that in a market with retail market power, the marketplace prefers steering by algorithm to ad auctions and that in this case, both the marketplace and consumers might suffer from the requirement of the consumer-friendly algorithm. Section 5 allows for some consumers who costlessly observe all offerings and shows how this leads the marketplace to prefer steering by algorithm. But in Section 4, where there is uncertainty or the marketplace chooses a common scheme for markets heterogeneous in demand or cost conditions, it prefers ad auctions. We then consider further issues around marketplace choices with broader notions of heterogeneous markets: notably the use of both auction-based and algorithmic steering simultaneously, and vertical integration, and we show how these may have pro-competitive effects.

3 Seller 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. 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 having a cost advantage.[25]

3.1 Match Probability Heterogeneity

We assume that with some probability η_i firm i's product is a bad match, in which case the 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 the realization of this uncertainty. The marketplace knows the match value probabilities, as do retailers. Without loss, let $\eta_2 \geq \eta_1$.[26] We will assume that the marketplace knows the match value probabilities, as do retailers.

We find that the algorithm dominates because it can extract Seller 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

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

[26]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 to total demand.

$f = f^m(c)$ and extract profits equal to

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

This is the maximum attainable profit that even an integrated monopolist can 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

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

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

Proposition 5. *When there is retailer market power through match probability heterogeneity, the marketplace can earn the full monopoly profit through algorithmic steering. Both the marketplace and consumers are better off with the algorithm.*

Proof. See the discussion in the text above. □

3.1.1 Consumer-friendly algorithm can be worse for consumers

Given Proposition 4, it is noteworthy that the result above does not hold when the marketplace is constrained to use the consumer-friendly algorithm. In this case, the Bertrand-like competition that the consumer-friendly algorithm induces leads the worse firm to charge a price equal to its fee-adjusted marginal cost, $\frac{c}{1-f}$. However, since the better Seller 1 can match the surplus that this induces at a strictly higher price and earn profits by doing so, the marketplace cannot extract the full monopoly profits at $f = f^m(c)$ or indeed at any fee.

Perhaps more surprisingly, it is easy to show that constraining a marketplace to use the consumer-friendly algorithm not only harms marketplace profits but also leads to higher retailer prices and lower consumer surplus compared to allowing the marketplace an unrestricted

choice among algorithms. This is a result of the marketplace being unable to appropriate all the profits that the better Seller 1 earns, and since the firm still retains some market power, through a standard double mark-up effect, the retail price that consumers face is higher than the monopoly price.

We illustrate this possibility in a simple example.

Assume there is one market with demand $q(p) = 1 - p$ with two sellers who have the same marginal cost c but differ in match value probability, where Seller 1 has match probability equal to 1 (that it is never a bad match for any consumer, $\eta_1 = 0$) while Seller 2 has a match probability $1 - \eta$, where we will assume $\eta > \frac{3}{4}$ for reasons that will become apparent momentarily.

We have shown that in such a situation the optimal algorithm allocates all the demand to Seller 1 and induces the seller to set a price equal to $p_1 = \frac{1+c}{2}$ and the optimal fee $f^* = \frac{1-c}{1+c}$, which would allow the marketplace to extract the monopoly profit.

Now assume that the platform must allocate the privileged position using the consumer-friendly algorithm. For a given price p, the consumer surplus associated with Seller 1 is given by $CS(p) = \frac{(1-p)^2}{2}$. It is easy to show that when $\eta > \frac{3}{4}$ then for any fee f, the Seller 1 generates a higher consumer surplus setting $p_1 = p^m(\frac{c}{1-f})$ than Seller 2 generates when setting $p_2 = \frac{c}{1-f}$ (its fee-adjusted marginal cost). Thus Seller 1 will choose to price at $p_1 = p^m(\frac{c}{1-f}) = \frac{1+\frac{c}{1-f}}{2}$. The marketplace maximizes its profits in this case by choosing the associated profit-maximizing fee, f_{CF} which is the solution to $(1-f)^3 = (1+f)c^2$.

It is easy to show that $f_{CF} < f^*$ due to double marginalization but also that $p^m(\frac{c}{1-f^*}) > p^m(c)$. Trivially, a fortiori, consumers are even worse off if the marketplace must move to the consumer-friendly algorithm but cannot adjust the fee when it does so.

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 the auction proceeds; that is, it would earn $(1-f)\pi^m\left(\frac{c_2}{1-f}\right)$.[27] 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

[27] 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. In our baseline analysis and all other extensions, allowing for a reservation price in the ad auction would not affect results.

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 perspective of the marketplace, which can rein in the price when using an algorithm.[28]

Proposition 6. *When there is retailer market power through cost heterogeneity, the marketplace earns the full monopoly profit through algorithmic steering. 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 converge to those when steering by algorithm.

Just as for match probability heterogeneity, requiring the marketplace to use the consumer-friendly outcome (which in this case of cost heterogeneity simply means privileging the lower-priced retailer) can lead to worse outcomes for consumers. The intuition is identical. The better retailer retains monopoly power through the consumer-friendly algorithm, and so a double mark-up effect arises that leads to lower consumer surplus than allowing for the flexibility in the algorithm that can constrain the retailer market power.

4 Markets Heterogeneous in Demand and Cost Conditions

Next, we return to the case in which retailers are equally competitive but we suppose that there are many markets. As discussed in the Introduction, we consider the case in which 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.

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

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, 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 or cost conditions, the marketplace can earn the full monopoly profit 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 and, specifically, the consumer-friendly algorithm. For consumers, ad-auction steering entails monopoly pricing in all markets. If the marketplace steers by algorithms it would prefer to charge different fees for different markets and so the constraint of a common fee will result in fees that are too high in some markets and too low in others, leading to retail prices above monopoly levels in some markets and below in others. Thus, some consumers would benefit and others suffer and overall effects are ambiguous.

5 Consumers Who Are Not Susceptible to Steering

Throughout, we have assumed that consumers in a market observe only the retailer in the privileged position—in that sense, all consumers are susceptible to steering. Although we have not modeled the consumer search behavior and, simply assumed that consumers only look at the first offering, such "susceptible" behavior arises naturally in a sequential search model

even with an infinitesimal cost for looking at a second offering.[29] This is a consequence of the fact that in much of our analysis, in equilibrium, the second offering is no better than the first. In this section, we suppose that there are consumers who observe *all* the retailers costlessly, akin to "shoppers" in Varian (1980) for example. We call these "non-susceptible" or "attentive". We allow for the possibility of both a mix of susceptible and non-susceptible consumers within the same market, denoting the fraction of susceptible consumers in market i by σ_i, in addition to heterogeneity in this fraction across markets.

For example, 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 plays 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 4.[30] 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, if 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. However, in an attentive market ($\sigma_i = 0$), Bertrand competition leads to 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 (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 susceptible 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. It can also do the same employing a consumer-friendly algorithm. Steering through an ad auction must lead to strictly lower revenues for the marketplace. The presence of susceptible consumers leads a retailer to price above (fee-adjusted) marginal cost, which is a double mark-up force; instead, a fee of 0 would result in

[29] The only exception is in the model variation of Section 3.1 where different retailers differ in their likelihood of being a suitable match for a consumer. In this case, similar results arise under the assumption of sufficiently high search costs, or that there is a perfect correlation in matches (so if the product less likely to be a good match does, in fact, turn out to be a good match, then the other product is also necessarily a good match).

[30] Moreover, just as in Proposition 4, the algorithm need not dictate price directly but would obtain an identical outcome using the consumer-friendly algorithm.

the retailers bidding no more than $\sigma\pi^m(c)$ (since the gains from winning the auction apply only to susceptible consumers) and $\pi^m(c)$ is an upper bound on the profits per consumer from winning the auction.[31] 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 also holds when steering with the consumer-friendly algorithm, but is generally not the case when steering through auctioned ads.*

Proof. Immediate following the discussion above. □

Consumers facing algorithmic steering will face monopoly prices in all markets. If instead the marketplace were constrained to employ ad auctions and a common fee, the presence of attentive consumers would lead it to charge a fee (leading to a retail price above the monopoly level in markets that feature only susceptible consumers). Some consumers would therefore face higher retail prices and others lower (say in markets that feature only attentive consumers) when the marketplace employs ad auctions and the impact on overall consumer surplus would be ambiguous.

It should be noted that another natural interpretation of the framework in this section 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 Multidimensional Market Heterogeneity

Thus far we have considered single markets or introduced a single form of heterogeneity at a time to highlight how the marketplace might have a strict preference for the use of steering algorithms or ad auctions. In this section, we allow for multiple forms of heterogeneity across markets. A complete analysis would be involved—both in terms of notation and in outlining a wide range of different cases to consider; however, a strength of our framework is that it is relatively tractable and it is feasible to consider many possible interesting variations. In this section, we highlight two such variations relevant to contemporary policy concerns: vertical integration and self-preferencing, and the extent to which allowing a marketplace to use ad auctions in addition to algorithmic steering might harm consumers.

[31] 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), 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 sellers 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 an optimal fee is, in general, not analytically tractable. Nevertheless, such a complete characterization is unnecessary for our results in this section.

6.1 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. This possibility is substantiated when there are reasons for strictly positive fees (such as attentive consumers or market power). As described above, steering by algorithm can be 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 in which a marketplace prefers to use an ad auction (for example, due to heterogeneous demand conditions) but also faces some markets with attentive consumers and others with susceptible consumers—which would otherwise call for steering by algorithm. If the marketplace were constrained to use only one steering method and chose an ad auction, it would want a fee of zero in markets with susceptible consumers but positive fees with attentive ones; 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. However, if the marketplace were steering by algorithm in every market (to account for a mix of attentive and susceptible ones), the heterogeneous demand conditions would imply that the marketplace would prefer different fees in different markets.[32] 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.

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 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. Rather, when it faces markets that are heterogeneous across different dimensions, it can gain. For example, consider the case in which 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 could choose a steering method and a fee that are 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

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

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 the case of algorithmic steering), will entail lower retail prices in at least one of the two markets.[33]

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 susceptible consumers. If the marketplace is constrained to employ the same steering method and fee across both markets, then, under vertical integration, the retail price will be lower in at least one of the two markets and indeed can be lower in both.*

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 (or the retail price reflects fee-adjusted marginal cost if all consumers are attentive), it can marginally increase the fees and the prices required to attain the privileged position in both markets and thereby strictly increase profits. This follows because under the algorithm, the price given to consumers is always optimally equal to the fee-adjusted marginal cost of that seller, and $f\frac{c}{1-f}q(\frac{c}{1-f})$ is always increasing in f for $\frac{c}{1-f} < p^m(c)$. Finally, to see that both retail prices may fall following vertical integration (that is both retail prices are above marginal costs), consider the following example. In Markets 1 and 2, demand is given by some well-behaved $q_1(p_1)$ and $q_2(p_2)$, respectively. All consumers are susceptible, and there are two sellers per market. Both sellers in Market 2 and one seller in Market 1 have marginal cost c and monopoly prices for that cost level in the two markets differ so that $f_1^m(c) \neq f_2^m(c)$. The remaining seller in Market 1 has marginal cost $c+\varepsilon$ where $\varepsilon > 0$. Given the demand heterogeneity, there clearly exists ε close enough to zero so that auction is the optimal mechanism. Furthermore, because $\varepsilon > 0$, $f^* > 0$ under the auction. It follows trivially that vertical integration into one of the markets reduces prices in both. □

The underlying mechanism contrasts somewhat with many discussions of self-steering that focus on effects on one particular retail market.[34] Instead, Proposition 9 highlights that

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

[34] See, for example Anderson and Bedre-Defolie (forthcoming), Etro (2023) 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.

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. The result that prices may fall in both markets contrasts with a standard multi-market monopoly model in which price discrimination always harms consumers in at least one market. The reason lies in the fact that, in our setting, vertical integration simultaneously allows for better price discrimination across markets and for relaxation of the double-marginalization problem.[35]

6.2 Steering With Both Algorithm and Ad Auctions

As highlighted in the Introduction, many marketplaces assign prominent positions using *both* ad auctions and steering algorithms. This is consistent with our model since each market can be understood as a particular individual entering a particular search term at a particular time. This practice, however, has come under scrutiny, especially in the context of some marketplaces that previously engaged in little advertising using it increasingly.[36] In this section, we show that restricting a marketplace (constrained to a common fee across markets) might lead to lower consumer surplus than would arise if it used both steering methods. We do so through an example that highlights that fees are endogenous and would be different in these two scenarios.

Consider a platform that manages two markets. In Market 1 there is only one seller with marginal cost $c < 1$ (and possibly another with marginal cost above 1) and the demand is given by $q_1(p) = 1 - p$. This market is characterized by (extreme) market power. In Market 2 there are two identical sellers, both with marginal cost c and unit demand (inelastic demand) with willingness to pay equal to 1. Market 2 is competitive, but the demand is less elastic than in the other market and if fees could be set separately, it would command a higher fee from the platform.[37] The relative sizes of Market 1 and Market 2 are given by q and $(1 - q)$, respectively.

The platform can use the same mechanism in both markets or mix-and-match between them, employing both. As we have learned in Section 3, the algorithm will do well in Market 1 (by taming the market power), but, following Section 4, the auction will have an edge due to demand heterogeneity across the two markets.

When employing both monetization strategies simultaneously, clearly it can never pay off to use ads in Market 1. Therefore, we will only consider the case where the algorithm is used in Market 1 and ads are used in Market 2. In either case, setting a fee above $1 - c$ can never

[35] de Cornière Alexandre, Montavani and Shekhar (2023) present a model where price discrimination that can lead to Pareto improvements but through a different mechanism that involves increased participation and network externalities.

[36] Our analysis does not speak to why this change has occurred. One might speculate that some marketplaces that in the past were focused on increasing participation of consumers and retailers have now acquired some "lock-in" and monopoly power and have shifted their strategies to exploit this monopoly power. This is, of course, far beyond what this paper can evaluate, and throughout our analysis, we assume the existence of monopoly power and a marketplace's objective to maximize its profits.

[37] The inelastic demand assumption helps to generate the stark result we are after—that this kind of "mix-and-match" mechanism benefits consumers. This is because in the market with inelastic demand, there is no increase in price when the platform switches from the algorithm to an auction.

be profitable (since it would mean that no seller is active), so from now on we will restrict our attention to $f \leq 1 - c$.

If the platform uses an algorithm in both markets, then it prefers to charge a fee equal to $\frac{1-c}{1+c}$ in Market 1 and (a higher) fee $1 - c$ in Market 2. Constrained to charge a common fee, the optimal uniform fee will maximize

$$q\left(1 - \frac{c}{1-f}\right)\frac{cf}{1-f} + (1-q)f,$$

where the first expression corresponds to revenue from the first market which is of size q and the second term in the summation corresponds to the revenue from the second market where unit demand ensures that any fee (below $1 - c$ as we have assumed) is collected (and scaled by the market size $1 - q$). It is immediate that the optimal fee will be too high for Market 1 and too low for Market 2.

If the platform uses only an auction instead, it will not extract anything from the bids in Market 1. The optimal fee in this case will maximize the platform profits given by[38]

$$q\left(1 - \frac{1+\frac{c}{1-f}}{2}\right)\frac{1+\frac{c}{1-f}}{2}f + (1-q)(1-c).$$

In general, whether the marketplace prefers steering by ad auction or algorithm depends on parameters, most crucially on q.

Now, let us consider what happens when the algorithm is used in Market 1 and an ad auction in Market 2.

In this case, platform profits are given by

$$q\left(1 - \frac{c}{1-f}\right)\frac{cf}{1-f} + (1-q)(1-c),$$

which must be higher than the profits under the previous two 'pure' regimes because by setting a fee $f^* = \frac{1-c}{1+c}$ the platform extracts integrated monopoly profits from both markets.

Consider now what happens if there is a ban on ads. In this case, when the platform switches to the algorithm, the fee must increase from $\frac{1-c}{1+c}$ because in Market 2 the optimal fee is higher. It is trivial that the platform will be harmed. However, consumers will also be harmed, because while in Market 2 they will pay the same (monopoly) price, in Market 1 the price will go up with the fee.

For sufficiently high q, the algorithm will always perform better than the auction because in Market 1 no bids are paid; thus, consumers will also be harmed relative to the mix-and-match regime if the platform is required to use the same mechanism in both markets (in which

[38] In Market 1 the only seller charges the monopoly price corresponding to the fee-inflated marginal cost. In Market 2 the platform extracts monopoly profits regardless of the fee because ex post the price is equal to 1. Thus, each seller is willing to bid $1 - f - c$, which alongside the fee revenues f adds up to $1 - c$, provided that the fee is no greater than $1 - c$, as is assumed.

case the platform would switch to the algorithm and increase the fee).

7 Conclusions and Policy Considerations

This paper presents 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. 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 to academics, policy makers, and practitioners interested in marketplaces and platforms, and we believe that the analysis in Section 6, which explores some aspects, is helpful in this regard.

First, we highlight that 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 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.

We also highlight that a natural "consumer-friendly" algorithm may, in fact, do little for consumers. In our baseline model with endogenous fees and many variations, consumers face monopoly prices when this algorithm is employed. Indeed, in Section 3.1.1 we illustrate that restricting the marketplace to use this algorithm can make consumers worse off. This is noteworthy in the context of calls to audit algorithms, such as CMA (2022), and highlights that there are equilibrium effects (in our case for the choice of commission fee and retail prices) that can make assessment involved.[39]

As explored throughout the paper, different steering 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, the extent to which heterogeneity across markets is then dominated by one or other of these types will govern the marketplace's preference for steering by auction or by algorithm.

These simple observations have immediate policy implications. Again, though perhaps trivial, these have not been discussed much. First, even in the absence of steering concerns

[39]See also Ichihashi and Smolin (2023), who make a similar point albeit in a different model.

(whether to steer 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.[40] 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. As discussed above, calls to ban marketplaces from producing their own goods and engaging in self-steering need not benefit consumers. Finally, while there has been concern surrounding some marketplaces beginning to employ ad auctions in addition to steering by algorithm (see Footnote 1 above), in Section 6.2 we illustrate that preventing a marketplace from doing so might be detrimental to consumers.

Part of the simplicity of our analysis arises from the monopoly position of the 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] Belleflamme, Paul and Martin Peitz (2021) *The Economics of Platforms*, Cambridge University Press.

[9] Bulow, Jeremy and Paul Klemperer (1996) " Auctions versus negotiations" *American Economic Review*, 86 (1), 180–194.

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

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

[11] Choi, Jay Pil, Kyungmin Kim and Arijit Mukherjee (2023) " "Sherlocking" and platform information policy" Available at SSRN 4332558.

[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 competition 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] CMA (2022) "Auditing Algorithms: The existing landscape, role of regulators and future outlook" at https://www.gov.uk/government/publications/findings-from-the-drcf-algorithmic-processing-workstream-spring-2022/auditing-algorithms-the-existing-landscape-role-of-regulators-and-future-outlook

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

[17] de Cornière Alexandre, Andrea Montovani and Shiva Shenkar (2023) "Third-degree price discrimination in two-sided markets." *mimeo.*

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

[19] Desai, Preyas S. (2000) "Multiple messages to retain retailers: Signaling new product demand." Marketing Science 19(4):381–389.

[20] Diamond, Peter (1971)"A model of price adjustment," textitJournal of Economic Theory, 3: 156–68.

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

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

[23] 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

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

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

[26] Furman, Jason, Diane Coyle, Amelia Fletcher, Derek McAuley and Philip Marsden (2019) *Unlocking Digital Competition: Report of the digital competition expert panel*. Report prepared for the Government of the United Kingdom

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

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

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

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

[31] 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.*

[32] Hervas-Drane, Andres and Sandro Shelegia (2022) " Retailer-led Marketplaces" mimeo.

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

[34] Hunold, Matthias, Ulrich Laitenberger, and Guillaume Thébaudin (2022) " Bye-Box: An analysis of non-promotion on the Amazon marketplace." mimeo.

[35] Ichihashi, Shota and Alex Smolin (2023) " Buyer-Optimal Algorithmic Consumption " mimeo.

[36] Inderst, Roman and Marco Ottaviani (2012)" Compeititon through commissions and kickback" *American Economic Review,* 102(2), 780–809.

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

[38] Kang, Zi Yang and Ellen V. Muir (2022) " Contracting and vertical control by a dominant platform" Stanford University, working paper

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

[40] Lam, H. Tai (2021) " Platform Search Design and Market Power. " mimeo

[41] Lariviere, Martin A., and V. Padmanabhan (1997) "Slotting allowances and new product introductions." Marketing Science 16(2):112–128.

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

[43] Long, Fei and Wilfred Amaldoss " Self-preferencing: Role of private labels and sponsored advertising in E-commerce marketplaces" *Marketing Science,* forthcoming

[44] Long, Fei, Kinshuk Jerath and Miklos Sarvary (2022) " Designing an online retail marketplace: Leveraging information from sponsored advertising" *Marketing Science*, 41(1):115–138.

[45] Madsen, Erik and Nikhil Vellodi (2021) " Insider imitation" Available at SSRN 3832712.

[46] 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).

[47] 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/.

[48] Raval, Devesh (2022) " Steering in One Click: Platform Self-Preferencing in the Amazon Buy Box."

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

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

[51] Shaffer, Greg (1991) "Slotting allowances and resale price maintenance: a comparison of facilitating practices." The RAND Journal of Economics, 120–135.

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

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

[54] Varian, Hal R (1980) "A model of sales." *American Economic Review*. 70(4): 651–659.

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

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

[57] Zou, Tianxin and Bo Zhou (2023) " Self-preferencing and search neutrality in online retail platforms, "Available at SSRN 3987361.