Econometrica, Vol. 94, No. 6 (November, 2026), 2081–2115

TWO-SIDED MARKETS SHAPED BY PLATFORM-GUIDED SEARCH

KWOK HAO LEE
Department of Strategy and Policy, National University of Singapore (NUS) Business School

LEON MUSOLFF
Business Economics and Public Policy Department, The Wharton School, The University of Pennsylvania

This paper investigates concerns that vertically integrated platforms like Amazon steer demand toward their own offers via algorithmic prominence, potentially harming consumers. On Amazon, for each product, the Buybox prominence algorithm selects one seller to feature, influencing which offers consumers consider. Using novel Amazon sales and Buybox (prominence) data, we estimate a structural model capturing the effects of such algorithmic prominence on consumer choices, seller pricing, and entry. We find that the platform can indeed steer demand as 95% of consumers consider only the Buybox offer. The Buybox is highly price-elastic (-21), but skews toward Amazon's own offers, which are featured as frequently as observably similar offers priced 5% cheaper. Still, as consumers prefer these offers, this skew does not amount to self-preferencing in the sense of harming consumers: consumer surplus is roughly maximized at the estimated Amazon Buybox advantage, which balances higher prices against showing consumers their preferred offers.

KEYWORDS: Two-sided markets, e-commerce, consumer search, self-preferencing.

1. INTRODUCTION

ONLINE PLATFORMS increasingly act not merely as passive intermediaries, but actively shape consumer experiences. Through search, ranking, and prominence algorithms, platforms like Amazon direct consumer attention. These algorithms significantly impact market outcomes: by influencing which options consumers consider, platforms effectively steer demand, thereby shaping price competition, merchant entry, and ultimately consumer welfare. We call this ability to steer demand *intermediation power*.

In this paper, we study the implications of such intermediation power on Amazon Marketplace, where Amazon both operates the platform and competes as a seller. Vertically integrated platforms like Amazon may engage in self-preferencing—guiding consumers preferentially toward their own offers, potentially harming consumers (Stigler Committee on Digital Platforms (2019)). This concern has drawn scrutiny from antitrust agencies (Federal Trade Commission and State Attorneys General (2023)) and regulators (European Commission (2022), U.S. Congress (2021), U.S. Senate Judiciary Committee (2020)). Besides potential static harm from not showing consumers their most preferred offers, proposed theories of harm from self-preferencing include a possible stifling of merchant entry and associated reduction in options available to consumers (Klingler, Bokemeyer, Della Rocca, and Nunes Bezerra (2020)).

Kwok Hao Lee: kwokhao@nus.edu.sg
Leon Musolff: lmusolff@wharton.upenn.edu

We thank Jakub Kastl, Kate Ho, Nicholas Buchholz, Adam Kapor, and Bo Honoré for invaluable guidance. We would also like to thank Bruno Baránek, Juan Camilo Castillo, Giovanni Compiani, Chris Conlon, Michael Dinerstein, Glenn Ellison, Chiara Farronato, Qiang Fu, Germán Gutiérrez, Andreas Haupt, Jean-François Houde, Ian Kaneshiro, Uri Miron, Devesh Raval, Imke Reimers, Tiffany Tsai, Raluca Ursu, Matthijs Wildenbeest, and Chuan Yu.

© 2026 The Authors. Econometrica published by Wiley Periodicals LLC on behalf of The Econometric Society. Leon Musolff is the corresponding author on this paper. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Our primary research question is whether Amazon's Buybox prominence algorithm is self-preferencing in the sense of harming consumers.[1] To this end, we build and estimate a tractable structural model of intermediation power, that is, a platform's ability to influence market outcomes by steering consumers. We model demand, the platform's prominence assignments (which steer consumers), pricing, and entry. We combine this model with unique proprietary data linking sales and prominence assignments on Amazon, a setting of prime regulatory concern. Our counterfactuals speak to various theories of harm from self-preferencing and the platform's tradeoff between attracting merchants and fostering competition.

Platforms commonly set defaults for consumers through their search (Yu (2024)), ranking (Ursu (2018)), and prominence algorithms (Johnson, Rhodes, and Wildenbeest (2023)). Economists have increasingly become aware that these defaults can influence consumer choice. We thus model consumers as having one of two kinds of consideration sets (Manski (1977), Ben-Akiva and Boccara (1995)): either they consider all available options, or they only consider the prominent offer and the outside option. Through its influence on consideration sets, the platform's prominence algorithm affects seller pricing (as in, e.g., Dinerstein, Einav, Levin, and Sundaresan (2018)), as well as *how many* and *which* firms enter.

We estimate this model using high-frequency Amazon data via maximum likelihood (prominence algorithm and demand) and GMM (pricing and entry). Variation in offer characteristics (and corresponding prominence assignments) identifies how Amazon decides which offers to make prominent. Similarly, variation in choice sets and observed sales identify consumer demand parameters. The fraction of consumers who consider all offers (as opposed to only the prominent offer) is identified by the difference in sales between offers eligible for prominence and those not eligible. The estimated demand structure and variation in merchant prices identify marginal costs. Finally, the model of variable profits implied by the estimated demand system and variation in market size identify fixed costs.

We find that Amazon can steer demand and indeed has a large influence over it as 95% of consumers consider only the Buybox offer. The Buybox is highly price-elastic (-21), but skews toward Amazon's own offers, which are featured as frequently as observably similar offers priced 5% cheaper.

After recovering the prominence algorithm, demand, and cost parameters, we simulate market outcomes under alternative prominence algorithms, proceeding in two stages: first with fixed prices and entry (short run), then allowing price and entry adjustments (long run). We consider counterfactuals that examine (i) the value of search guidance, (ii) how the price sensitivity of prominence affects consumer surplus, and (iii) self-preferencing, that is, whether the frequency with which Amazon promotes its own offers is consumer surplus maximizing.

First, we evaluate the impact of search guidance by comparing the status quo to a benchmark in which unsophisticated consumers only consider the cheapest offer. While stylized, this counterfactual illustrates how search guidance shapes market outcomes. We find that the algorithm increases consumer surplus in the short run, mostly because of better matching of consumers to high-quality offers. Relative to our benchmark, the algorithm also increases first-party producer surplus in the short run, incentivizing merchant

[1]Specifically, we study self-preferencing by Amazon in its role as a *retailer*—procuring products from manufacturers and competing in the Buybox against third-party sellers offering the same products—not in its role as a producer of its own brands (e.g., Amazon Basics).

entry and lower prices in the long run. Although these effects are modest, together they reinforce the increase in consumer surplus.

Second, we further investigate the prominence algorithm's price sensitivity, addressing concerns that the platform, which makes revenue from ad valorem intermediation fees, may be insufficiently incentivized to foster competition. Contrary to such concerns, we find that increasing the algorithm's price sensitivity would lower consumer surplus. Indeed, our evaluation suggests that the platform's price sensitivity is approximately optimal from a consumer surplus perspective.

Third, we investigate the question at the core of our paper: is Amazon self-preferencing in the sense of harming consumers? To answer this question, we explore how the frequency at which Amazon currently promotes its own offers compares to that which would maximize consumer surplus. While more frequent promotion of Amazon offers leads it to raise prices, it also allows more consumers to consider Amazon's offers, which they strongly prefer. We find that at the estimated promotion frequency, these two effects approximately balance, so that Amazon is not self-preferencing in the sense of harming consumers: increasing the Amazon promotion frequency would lead to consumer surplus losses from increased prices, but lowering it would lead to losses from worse matching of consumers to their preferred offers. The current promotion frequency is thus approximately consumer surplus maximizing.

2. LITERATURE REVIEW

This paper speaks to three strands of literature in industrial organization and the economics of digitization.

First, we contribute to a large literature on price competition and search frictions in online markets. Despite theoretically small search costs, online markets still exhibit substantial price dispersion (Bailey (1998), Smith and Brynjolfsson (2001), Baye, Morgan, and Scholten (2004), Einav, Kuchler, Levin, and Sundaresan (2015)), possibly because consumers do not search efficiently (Stigler (1961), Malmendier and Lee (2011), De Los Santos, Hortaçsu, and Wildenbeest (2012), Schneider (2016)) and consider only a subset of available alternatives (Goeree (2008)). Given these frictions, platforms face a tradeoff between incentivizing sellers to compete on price and guiding consumers to their preferred products (Dinerstein, Einav, Levin, and Sundaresan (2018)). Efficient search technologies can lead to high offer-level elasticities, prompting retailers to obfuscate product characteristics to hinder comparison and raise profits (Ellison and Ellison (2009)). Finally, personalized prominence as in recommender systems can also affect pricing incentives (Calvano, Calzolari, Denicolò, and Pastorello (2025)). While addressing similar platform-design issues, this literature rarely models entry or a firm competing on its own platform.

Second, we contribute to the literature on the downstream conduct of vertically integrated firms. In traditional markets, concerns include foreclosure (Crawford, Lee, Whinston, and Yurukoglu (2018)), raising prices of rivals' products (Luco and Marshall (2020)), or altering products to disadvantage competitors (McManus, Nevo, Nolan, and Williams (2020)). In digital markets, the focus shifts to steering, with self-preferencing defined as "platform [prominence assignments that] deliver outcomes below [a] frontier that maximizes a weighted sum of seller and consumer surplus" (Reimers and Waldfogel (2023)). Such steering may negatively affect consumers through higher commissions and prices (e.g., Teh and Wright (2022), Hagiu, Teh, and Wright (2022)).

Third, spurred partly by Khan (2016), our empirical setting—the Amazon marketplace—has emerged as a central focus for studies of price competition and search on platform marketplaces (Ciotti and Madio (2023), Gutiérrez (2021), Lam (2021), Crawford, Courthoud, Seibel, and Zuzek (2022)), including research on steering (Hartzell and Haupt (2025), Raval (2022), Chen and Tsai (2024), Farronato, Fradkin, and MacKay (2025), Waldfogel (2024), Boldrini and Clavorà Braulin (2025)). Etro (2024) provides a survey. Our paper establishes that algorithmic steering exists and is of first-order relevance to consumer welfare.

We highlight two particularly relevant studies. Farronato, Fradkin, and MacKay (2025) experimentally alter Amazon brand visibility in search results, finding consumer surplus would fall if these brands were removed. We, however, investigate a distinct steering dimension (Buybox prominence), which has garnered separate regulatory scrutiny (U.S. Senate Judiciary Committee (2020)). Raval (2022) also investigates the Buybox, finding steering toward Amazon. Our work, using sales data, analyzes not just the presence but also the welfare effects of this steering.

3. WHO DOES THE AMAZON BUYBOX FEATURE?

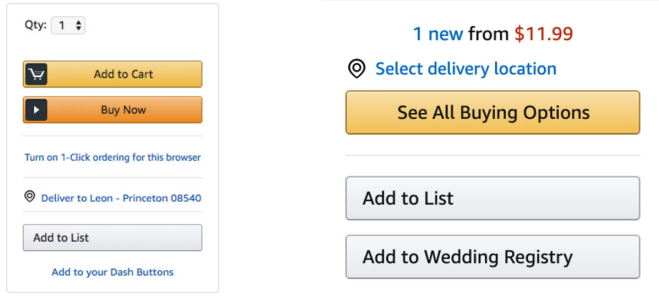
Our empirical setting is Amazon Marketplace, a platform permitting third parties to list offers alongside Amazon's. In 2020, about 1.9 million merchants utilized this opportunity (aboutamazon.com). Their listings were responsible for 60% of sales on this marketplace, yielding estimated merchant profits of $25 billion (aboutamazon.com). However, with 28% of purchases completed within three minutes, consumers use little time to explore their options (aboutamazon.com). This speed suggests that search, ranking, and prominence algorithms play a role in shaping consumers' choices.

E-commerce platforms typically employ multiple such algorithms. However, investigating algorithm behavior is challenging when products differ on unobserved dimensions. We exploit a unique feature of our setting to sidestep this issue: because Amazon requires sellers to list their offers on the correct product page, a specific product often has multiple offers.

When multiple offers are present, the platform automatically designates *at most one* offer as prominent, that is, featured in the "Buybox" (Figure 1). This offer does not vary across consumers.[2] Often, Amazon "rotates" prominence between sellers of sufficiently high quality to "avoid stock-outs" and encourage sellers to "delight their customers" and "get good feedback" (feedvisor.com). Prominence is critical for sellers, as the "vast majority of sales" go through the "Buybox" (europa.eu), with estimates from 80% (U.S. Senate Judiciary Committee (2020)) to 98% (Federal Trade Commission and State Attorneys General (2023, p. 29)).

Econometrically, this algorithm offers an ideal setting to study the impact of platform-guided consumer search. In particular, a market will be a specific product—for example, "Clarks Men's Bushacre 2 Chukka [Shoes], Dark Brown, Size 8.5." Once decided on this product, consumers must choose between various offers. For instance, the merchant "Zappos" offers these shoes for US$57.68 and will deliver them via Fulfillment by Amazon (FBA), a "fulfillment service that allows businesses to use Amazon to store, pick, pack, and ship customer orders" (amazon.com). Alternatively, "BHFO" charges only US$54.99 but does not offer FBA. Below, we examine the influence of prominence

[2]In rare situations, two offers are prominent: one for Prime customers, and one for non-Prime customers. We drop these rare cases.



(a) Occupied Buybox (b) Canceled Buybox

FIGURE 1.—The "Amazon Buybox." Amazon chooses which seller is assigned the sale if the buttons inside the rectangle are used. In (a), a seller has been assigned, while in (b), no seller has been featured.

on how consumers choose between such offers. Crucially, in this choice, all options share all product characteristics—obviating the need for complex demand models to explain substitution.

Amazon Marketplace is a key setting to evaluate antitrust concerns. Economists worry that "defaults can direct a consumer to the choice that is most profitable to the platform" (Stigler Committee on Digital Platforms (2019)). Indeed, sellers complain that competing with Amazon is challenging, as the platform frequently assigns the Buybox to itself. A Senate investigation concluded that it "can give itself favorable treatment [and] has done so through its control over the Buy Box" (U.S. Senate Judiciary Committee (2020)). The investigation sparked the introduction of a draft bill in Congress prohibiting self-preferencing (U.S. Congress (2021)).

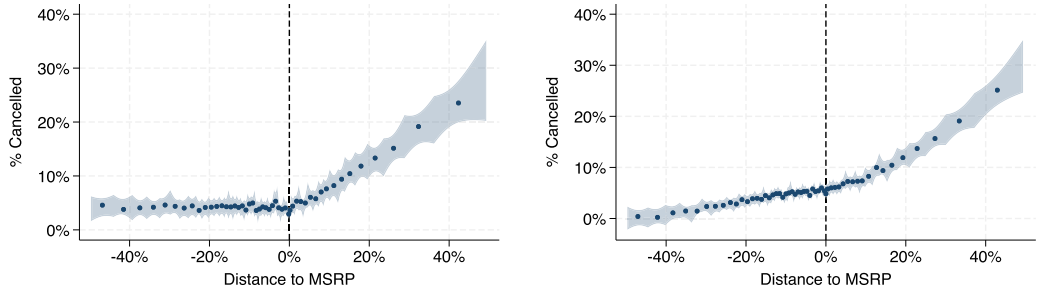
Recent allegations also claim Amazon modified its search algorithms to favor its own products (wsj.com), though evidence is mixed (Farronato, Fradkin, and MacKay (2023, 2025), Waldfogel (2024)). These claims suggest seller concerns about the Buybox may be justified. However, Amazon's general counsel counters that "the Buy box is aimed to predict what customers want to buy" and that "[the platform applies] the same criteria whether [the merchant is] a third-party seller or Amazon" (thehill.com).

3.1. *Data*

We procured extensive, high-frequency data that include prices, 19409013 prominence assignments, and 803849 sales on 47403 products sold by 63620 merchants from 2018-08-26 until 2020-03-23. We provide summary statistics and discuss our data in more detail in Supplemental Appendix 1 (Lee and Musolff (2026)).

Our data are sourced from a company that offers "repricing" services; therefore, these data arise directly from Amazon's APIs. For each monitored product, Amazon automatically notifies the company when there is any change in the number or content of the offers (e.g., a price hike or the entry of a new competitor). Importantly, we see which offer Amazon assigns prominence to *just after* any such change. However, prominence assignments may change frequently, and we receive no notification when the only change is which offer is assigned prominence.

We further observe, via a different API, all sales for *a single merchant* in each market—the one using the repricing company's services. These data come with their own challenges



(a) Unresidualized (b) Residualized

FIGURE 2.—Evidence of canceled prominence (binned scatter plots). This binscatter illustrates the relationship between canceled prominence and the price of the cheapest offer on a product relative to MSRP; dots are binned averages and the gray shadow provides a 95% confidence band. The underlying observations are product-date pairs for which the lowest-priced offer is priced between 50% below and 50% above MSRP. The left panel shows the unresidualized relationship; the right panel depicts the relationship after residualizing on product fixed effects. In both panels, as the cheapest offer's markup over MSRP increases, canceled prominence becomes more likely.

(e.g., having to proxy for the market size), which we address with our model and estimation strategy in Sections 4 and 5.[3]

3.2. *Canceled Prominence*

We first investigate when the platform cancels prominence, leaving the Buybox empty as in Figure 1b. This helps us understand the platform's incentives in shaping the prominence algorithm.

While an empty Buybox may lower sales, the threat of canceled prominence can discipline prices. Indeed, before 2019, Amazon's U.S. marketplace required its third-party sellers to offer lower prices than on other platforms (theverge.com). These "most-favored nations clauses" (MFN) were deemed anticompetitive by European antitrust agencies, leading Amazon to drop them in Europe by 2013 (German Competition Authority (Bundeskartellamt) (2013, 2015), UK Competition Authority (Office of Fair Trading) (2013)). Following similar pressure, Amazon also dropped the U.S. MFN clauses by 2019.

Nonetheless, canceled prominence may achieve similar effects to explicit MFN. To demonstrate, we combine data on the price of the lowest-priced offer on product j at date τ with data on Buybox ownership and information on the manufacturer's suggested retail price (MSRP), which we use as a proxy for the price at which a product could be obtained outside of Amazon. Figure 2 displays a binned scatter plot of Buybox cancellation versus the percentage distance to MSRP. As expected, the cancellation fraction increases with the cheapest offer's price. However, the increase is sharper if the offer price exceeds MSRP.

3.3. *Prominence Descriptives*

How does the platform choose which offer to feature? Table I investigates determinants of prominence. For instance, the prominent offer is the cheapest in 41% of cases. How-

[3]Yet, to our knowledge, no past papers employ sales data across many merchants to estimate demand on Amazon Marketplace. Instead, employing sales ranks as a proxy has become common.

TABLE I

DETERMINANTS OF PROMINENCE.

	Fraction of Prominence Assignments		
	Overall	FBA	Amazon
Amazon Offer Exists	19.25%	19.83%	100.00%
FBA Offer Exists	97.07%	100.00%	100.00%
Prominent Offer Is ...			
Lowest Priced	40.60%	40.24%	45.27%
Second Lowest Priced	26.42%	26.73%	30.05%
Highest Feedback Count	29.32%	29.52%	66.67%
Highest Feedback Rating	19.17%	19.15%	8.90%
Fastest Shipping	25.71%	25.74%	31.22%
Lowest Priced FBA	50.31%	51.83%	64.96%
FBA	85.98%	88.57%	97.18%
Amazon	12.83%	13.22%	66.67%
Lowest Fastest Shipping	51.92%	52.36%	65.46%

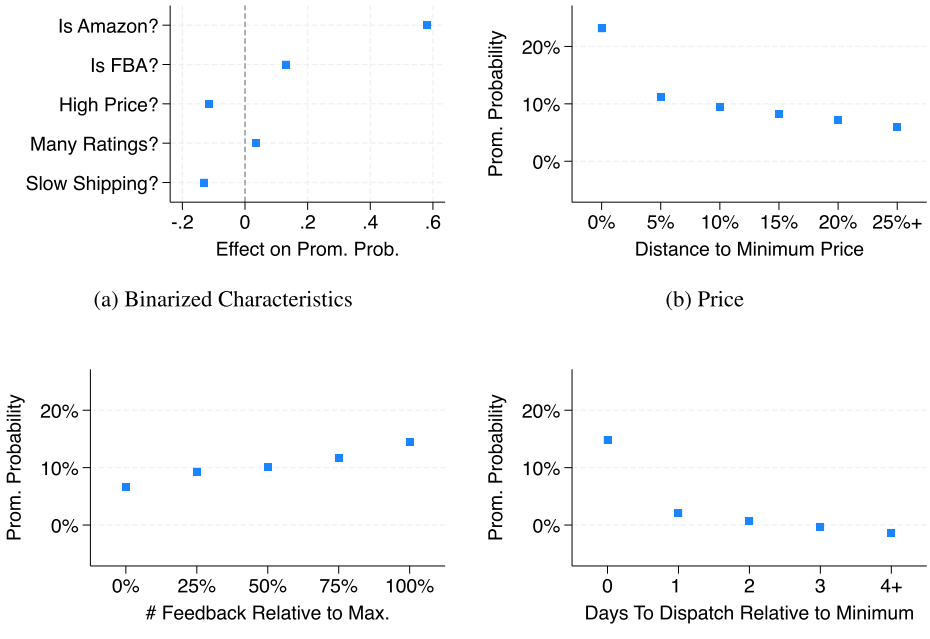
Note: Each entry in this table gives the fraction of prominence assignments satisfying the criteria listed in the first column for various subsets of the data. The second column ("Overall") is based on the entire set of observations. The third column ("FBA") employs only observations for which there is an offer that is fulfilled by Amazon, and the fourth column ("Amazon") uses only observations for which there is an offer by Amazon itself. All ties are broken randomly.

ever, Fulfilled by Amazon (FBA) offers are given prominence in 86% of cases, possibly because they dispatch much faster. Most interestingly, on products where Amazon has an offer, its offer is featured 67% of the time.

Figure 3 provides more detail on how offer characteristics matter. The top-left panel shows coefficients from bivariate regressions of a prominence dummy on (binarized) offer characteristics; continuous characteristics are split at the median. This figure shows that being sold by Amazon most strongly predicts prominence; however, as the underlying regressions are bivariate, this could reflect Amazon's offers being more competitively priced or of higher quality. FBA offers are also more likely featured, while high price, low feedback count, and slow dispatch times all correlate with lower prominence likelihood.

The remaining panels further investigate these three continuous characteristics. Acknowledging the prominence decision's discrete-choice nature, they normalize each characteristic relative to the product's best offer (e.g., expressing price as a percentage distance relative to the cheapest offer for the same product). These panels report coefficients from regressing a prominence dummy on evenly spaced bins of the normalized characteristic, controlling for offer fixed effects. Price is an important determinant, but even offers 20% above the minimum are still featured 7% of the time. Similarly, feedback count and especially dispatch time matter but are not solely decisive. The dispatch time effect centers almost entirely on whether shipping is instant. Crucially, these results do not control for other characteristics when looking at the impact of any given characteristic—in other words, the coefficients should not be understood as causal effects. Below, we estimate a discrete-choice model to disentangle the various factors influencing prominence, and provide assumptions under which the results of that model can be interpreted as causal.

Having shown that Amazon frequently features its own offers, we now turn to demand: is it true that Amazon offers are preferred by consumers? To investigate this question without access to sales data for Amazon offers—Amazon is not a client of the repricing company that is our data source—we turn to the idea that a strong preference for



(a) Binarized Characteristics

(b) Price

(c) Feedback Count

(d) Shipping Time

FIGURE 3.—More determinants of prominence. Panel (a) shows the result of regressing a prominence dummy on binary characteristics (continuous characteristics were binarized by cutting at the median) without any further controls. Panels (b)–(d) show results from regressing the same prominence dummy on various offer characteristics—normalized by the product's best offer's value—while controlling for offer fixed effects. Standard errors are clustered at the product level (but confidence intervals are too small to be visible).

TABLE II

EFFECT OF AMAZON'S PRESENCE ON THIRD-PARTY SALES.

	Norm. Sales	Norm. Sales/Market Size			
	(1)	(2)	(3)	(4)	(5)
Amazon Offer?	−38.72	−24.70	−25.79	−21.35	−10.98
	(9.978)	(1.505)	(1.608)	(1.519)	(1.868)
# Comp. Merchants			−0.667	−1.467	−1.130
			(0.171)	(0.168)	(0.148)
Divide by Mkt. Size	No	Yes	Yes	Yes	Yes
Controls	No	No	No	Yes	Yes
Offer Fixed Effects	No	No	No	No	Yes
Observations	19,409,013	19,409,013	19,409,013	19,409,013	19,408,472

Note: We regress normalized sales for the observed merchant on a dummy for whether there is an Amazon offer on the same market (i.e., product page and time.) Sales are normalized by dividing by their mean and multiplying by 100; regression coefficients can be interpreted as percentage point changes in sales. The first column uses raw sales data (showing that the effect is not due to our market size proxy); other columns use sales divided by market size (increasing power.) The *# Comp. Merchants* row counts the number of merchants pricing within 5% of the cheapest offer (for comparison, on average Amazon prices 7% above the cheapest offer.) *Controls* refer to price divided by MSRP, whether the offer is fulfilled by Amazon, shipping time, and the logarithm of the number of reviews. The effect of an Amazon offer on sales is negative, large, and statistically significant, even when controlling for the number of competitive merchants.

Amazon offers would be reflected in reduced sales for the observed merchants whenever Amazon is present on a market. We indeed find such a pattern in Table II, which provides model-free evidence that Amazon's presence casts a large shadow on the observed merchant's sales by regressing the observed merchant's normalized sales on an Amazon presence dummy (the normalization allows interpreting the coefficient as a percentage sales change). We caution that this shadow could be caused either by a consumer preference for Amazon offers or by consumers only considering prominent offers (which we have just seen may favor Amazon offers). To disentangle these two possibilities, we next build a model of how consumers choose in the presence of steering by a prominence algorithm.

4. MODEL

To disentangle consumer preferences from the influence of prominence on choices, and to assess the welfare impact of alternative prominence algorithm designs, we build a model of consumer choice and within-platform price competition and entry. This model lets us examine whether designers of online marketplaces derive intermediation power from choosing their prominence algorithms. Our approach focuses on the platform's ability to guide search by influencing customers' consideration sets (Manski (1977), Ben-Akiva and Boccara (1995)).

This section proceeds as follows. First, we specify demand: consumers have nested logit preferences with endogenous consideration sets. Second, we consider how the platform's prominence algorithm, also a nested logit, maps offer characteristics to a prominence assignment. Finally, we specify supply. On each market, potential sellers observe their own marginal costs and market-level fixed costs, make entry decisions, and then compete in a standard Bertrand–Nash pricing game. Crucially, sellers anticipate that demand depends on prominence, which hence shapes their incentives to enter and compete.

4.1. *Consumer Choice*

Demand follows a standard nested-logit framework except that consumers form endogenous consideration sets (Goeree (2008), Dinerstein, Einav, Levin, and Sundaresan (2018)).

There are two types of consumers. A fraction ρ are "sophisticated." Sophisticated consumers ignore the prominence assignment and evaluate all available options $\mathcal{J} \cup \{0\}$. The remaining $1 - \rho$ consumers are "unsophisticated." Unsophisticated consumers consider only the prominent offer $j^r \in \mathcal{J}$ and the outside option. Thus, the prominence algorithm chooses which offer the unsophisticated consumers evaluate.[4] If the algorithm chooses not to make any offer prominent, unsophisticated consumers take their outside option.[5]

In each market t, both types of consumers have preferences over the alternatives available to them. Their mean utilities for alternative j depend on its characteristics $\mathbf{x}_{jt}$, price p_{jt}, and unobserved quality ξ_{jt}:

$$\delta_{jt} = \mathbf{x}_{jt}'\beta - \alpha_t p_{jt} + \xi_{jt}. \tag{1}$$

[4]Note that, on Amazon Marketplace, consumers that do not explicitly click through to the offer listing are indeed never informed of the availability or characteristics of non-prominent offers.

[5]Here, search effort is taken as exogenous. In practice, if consumers respond to the absence of the prominence algorithm by "searching harder," our results may overstate the value of search guidance in our setting.

When discussing identification below, we impose further restrictions on ξ_{jt}: our main estimates utilize $\xi_{jt} \equiv 0$, but we also explore robustness to a more flexible specification. We also normalize the mean utility of the outside option to zero: $\delta_{0t} = 0$.

By contrast to the usual setup where markets are repeated observations of the same market over time (Berry, Levinsohn, and Pakes (1995)) or different geographic markets with nearly identical alternatives (Nevo (2001)), in our setting markets are distinct product pages. Hence, the alternatives offered differ markedly across markets, and so could price sensitivity: the composition of consumer arrivals on the page for a pen is likely to be different from that of a laptop. Hence, we assume that consumers evaluate price relative to MSRP. If R_t is the MSRP in market t, we set $\alpha_t = \alpha/R_t$, reflecting the idea that a \$0.50 price difference matters for a \$1 pen but not a \$1,000 laptop.

The flipside of alternatives differing markedly *across* markets is that *within* markets, the inside options available to consumers are very similar: they all correspond to getting the same product, albeit delivered via different channels at varying speed from sellers of varying reputation. Hence, plausibly, the inside alternatives are more substitutable with each other than with the outside option, which includes buying other products on the platform, buying the same product on some other platform or offline, and not buying the product at all.

To reflect this difference in substitutability, we partition the set of products as $\{0\} \cup \mathcal{J}_t$ and let $g = 1$ be the index of the nest of inside options. The utility consumer i derives from alternative j in market t is then

$$v_{ijt} = \delta_{jt} + \zeta_{ig(j)t} + \lambda \epsilon_{ijt}, \tag{2}$$

where ϵ_{ijt} is distributed i.i.d. Type-1 Extreme Value, and $\zeta_{ig(j)t}$ is common to all options in the same nest. The distribution of $\zeta_{ig(j)t}$ can be specified such that $\zeta_{ig(j)t} + \lambda \epsilon_{ijt}$ has a Generalized Extreme Value distribution, yielding a nested-logit model (Cardell (1997), McFadden (1978)).

Consumers choose the utility-maximizing option in their consideration sets. Therefore, given a prominent offer j^r, the probability that a consumer chooses product $j \neq 0$ in market t depends on a prominence dummy $y^r_{jt} \equiv \mathbf{1}\{j = j^r_t\}$:

$$s_{jt}\left(y^r_{jt}; \theta^d\right) = \rho \times \frac{\left[\sum_{k \in \mathcal{J}_t} \exp(\delta_{kt}/\lambda)\right]^{\lambda}}{1 + \left[\sum_{k \in \mathcal{J}_t} \exp(\delta_{kt}/\lambda)\right]^{\lambda}} \times \frac{\exp(\delta_{jt}/\lambda)}{\sum_{k \in \mathcal{J}_t} \exp(\delta_{kt}/\lambda)}$$
$$+ (1 - \rho) \times y^r_{jt} \times \frac{\exp(\delta_{jt})}{1 + \exp(\delta_{jt})}, \tag{3}$$

where $\theta^d = (\beta, \alpha, \lambda, \rho)$ collects the parameters of the demand model.

We assume that the number of customers making a choice in market t (i.e., the number of arrivals) has a Poisson distribution with mean M_t.[6] By the thinning property of the Poisson distribution (Blitzstein and Hwang (2014, Theorem 13.12.13)), the number of customers choosing j on market t is then also Poisson with mean $M_t \times s_{jt}(y^r_{jt}; \theta^d)$.

[6]By contrast, if we assumed that exactly M_t consumers arrived, our model would not be able to rationalize markets with more sales than the average number of consumer arrivals.

4.2. *Prominence Algorithm*

A *prominence algorithm* maps offer characteristics to a prominence assignment $j^r \in \mathcal{J}$ or—if prominence is canceled as in Section 3.2—the null assignment $j^r = \{0\}$. While, in general, prominence may vary by consumer, this is not the case in our empirical application. Hence, we model the prominence algorithm as solving exactly one discrete choice problem for each market.

The inner workings of the prominence algorithm are not known. But whether it is a machine learning algorithm that maximizes an objective such as total sales or a simpler rule-based routine, what matters to us is how the prominence algorithm shapes demand, pricing, and entry. To this end, the important question is how often sellers can expect their offers to be promoted, and how this depends on their own and their rivals' characteristics, including price and whether they are competing with Amazon. As with consumer choice, we capture this dependence parsimoniously with a nested-logit model. To the extent that the algorithm trades off various goals (e.g., sales profits vs. seller surplus), counterfactuals should be understood as altering the weight given to these goals (e.g., lowering the weight on sales profits could correspond to lowering the Amazon prominence advantage).

The mean utility δ^r_{jt} of each inside alternative includes observable characteristics $\mathbf{x}_{jt}$, price p_{jt}, and the econometrician-unobservable quality ξ^r_{jt}:

$$\delta^r_{jt} = \mathbf{x}'_{jt}\beta^r_t - \alpha^r \frac{p_{jt}}{R_t} + \xi^r_{jt}. \tag{4}$$

As with demand, when discussing identification below, we impose further restrictions on ξ^r_{jt}: our main estimates utilize $\xi^r_{jt} \equiv 0$, but we also explore robustness to a more flexible specification. We also normalize the mean utility of the outside option to zero: $\delta^r_{0t} = 0$.

Whether due to deliberate randomization or mistakes in evaluation, we assume prominence decisions are based on a shocked version v^r_{jt} of this mean utility. In particular, shocks follow a nested-logit structure like with consumer choice.[7] Like before, all inside alternatives share a common nest, and the outside option is in a separate nest; unlike before, this nesting here reflects the fact that cancelling prominence is likely to lower intermediation fee revenue, and hence is qualitatively different from just assigning prominence to a different offer. As the shock structure is perfectly analogous to consumer choice, we will not repeat the details here. Instead, writing the nesting coefficient of the Buybox as λ^r, we skip ahead to the implied probability of prominence:

$$r_{jt}(\theta^r) = \frac{\left[\sum_{k\in\mathcal{J}_t} \exp\left(\delta^r_{kt}/\lambda^r\right)\right]^{\lambda^r}}{1 + \left[\sum_{k\in\mathcal{J}_t} \exp\left(\delta^r_{kt}/\lambda^r\right)\right]^{\lambda^r}} \times \frac{\exp\left(\delta^r_{jt}/\lambda^r\right)}{\sum_{k\in\mathcal{J}_t} \exp\left(\delta^r_{kt}/\lambda^r\right)}, \tag{5}$$

where $\theta^r = (\beta^r, \alpha^r, \lambda^r)$ collects the parameters of the prominence model.

[7] One might assume Amazon's prominence algorithm is *deterministic*, implying the variance of the logit shock ϵ is (near) zero. However, industry sources and our web scraping indicate stochasticity via "Buybox rotations": the platform shares arrivals among similarly priced merchants. Therefore, ϵ contains said rotations and other idiosyncratic factors affecting prominence.

An alternative hypothesis is that the algorithm's ϵ reflects informational asymmetries. For instance, if a seller is having trouble fulfilling orders, then the algorithm could steer consumers clear of these offers, thereby improving customer satisfaction.

4.3. *Firm Choices*

Finally, we specify a model of entry and price competition. Each market $t \in \mathcal{T}$ is associated with a fixed cost F_t and a set of potential entrants $j \in \mathcal{N}_t$.

Fixing a market t, each putative seller j draws a type $\omega_j \sim G(\cdot)$[8] and chooses a pair (χ_j, p_j) consisting of an entry and a pricing strategy. The entry rule $\chi_j : \mathcal{I}_j \to \{0, 1\}$ maps a player's pre-entry information set $\mathcal{I}_j$ into her entry decision, whereas the pricing rule $p_j : \mathcal{I}' \to \mathbb{R}^+$ maps the post-entry information set $\mathcal{I}'$ into the price the seller will charge.

As merchants receive only a share ϕ of the revenue they generate due to intermediation fees, and assuming constant marginal costs, seller j's payoff is

$$\pi_j(\boldsymbol{\omega}, \boldsymbol{\chi}, \mathbf{p}) = \sum_{t \in \{\tilde{t} | j \in \mathcal{N}_{\tilde{t}}\}} \chi_{jt} \times \left[(\phi p_{jt} - c_{jt}) s_{jt}\big(r_{jt}(\boldsymbol{\omega}, \mathbf{p}), \boldsymbol{\omega}, \mathbf{p}\big) M_t - F_t\right], \tag{6}$$

where s_{jt} maps the vectors of types and prices into demand and M_t denotes expected monthly arrivals in market t. Our solution concept is Bayesian Nash Equilibrium.

We proceed by stating several assumptions for tractability:

ASSUMPTION 1—Separable Markets: *Each firm is a potential entrant in at most one market*: $\mathcal{N}_t \cap \mathcal{N}_s = \emptyset$ *for* $t \neq s$.

ASSUMPTION 2—Full Information Pricing: *After entry, each seller knows the identities, cost draws, and quality draws of its opponents when playing the pricing game*: $\mathcal{I}'_j \equiv \mathcal{I}' = (\boldsymbol{\omega}, F)$.

ASSUMPTION 3—Platform Offer Maximizes Short-Run Profits: *On each market, the platform's decision whether to list its first-party offer is taken as given. The price of the platform offer is set market-by-market to maximize the profits accruing to the said offer.*

ASSUMPTION 4—Blind Entry: *Before entry, each seller observes only its own marginal cost draw and the common fixed cost of entry*: $\mathcal{I}_j = (c_j, F)$.

ASSUMPTION 5—Symmetric Entry: $\chi_j^*(c, F) = \chi_k^*(c, F)$ *for all* $j, k \in \mathcal{N}$ *and* $c, F \in \mathbb{R}^+$.

The separable-markets assumption allows us to solve for equilibrium market-by-market, easing the computational burden. In our empirical application, firms are small and rarely specialize in related products, so cross-market cannibalization is of secondary importance.

Full-information pricing is standard: successful entrants play a Bertrand–Nash game (Anderson, De Palma, and Thisse (1992)) and internalize the effect of prominence when setting prices,

$$\frac{\partial s_j}{\partial p_j} = \rho \frac{\partial s_j^{\text{sophisticated}}}{\partial p_j} + (1 - \rho)\left[\frac{\partial r_j}{\partial p_j} s_j^{\text{unsophisticated}} + r_j \frac{\partial s_j^{\text{unsophisticated}}}{\partial p_j}\right]. \tag{7}$$

By influencing $\frac{\partial r_j}{\partial p_j}$, the prominence algorithm therefore partially determines the price elasticity sellers face—exactly where the platform's power to guide consumer search manifests in our model.

[8] A seller's type ω collects her marginal cost draw, attractiveness to the prominence algorithm (including ξ^r), and attractiveness to consumers (including ξ).

As for the platform offer maximizing short-run profits, we make this assumption for tractability and because it is an important benchmark for our analysis. Supplemental Appendix 7.2 explores how results change if Amazon's prices are fixed in the counterfactuals.

Given a selection rule for the pricing equilibrium,[9] Assumptions 4 and 5 ensure unique, computationally tractable equilibria at the entry stage. In each market, a set of potential entrants (sellers) $\mathcal{N}$ face a fixed cost of entry F. Each potential entrant independently decides whether to enter after observing information set $\mathcal{I}_j = \{c_j, F\}$. This assumption consists of two parts. First, in the spirit of Roberts and Sweeting (2013) and to avoid the classic equilibrium multiplicity problems, we deprive sellers of knowledge about offers *other* than their own.[10] Second, before entry, each firm's only information about its type is its cost. Because (we estimate that) price is the most important determinant of prominence, this assumption will only slightly affect our results.

Each firm enters if and only if its expected profits from entering exceed the fixed cost F. Given our assumptions, this expectation depends only on a firm's own unit cost draw. As its expected profits are declining in its own unit cost c_j, each firm best responds in cutoff strategies: $\chi_j(c_j) = \mathbf{1}\{c_j < c_j^*\}$. Thus, an equilibrium can be characterized by a vector $(c_1^*, \ldots, c_n^*)$. Assumption 5 further restricts attention to symmetric equilibria. Thus, all (ex ante identical) firms share the same (ex ante) cutoff c^*. Conditional on a selection rule for pricing equilibria, the following holds:

PROPOSITION 1: *The entry game has a unique symmetric equilibrium in cutoff strategies*.

This unique cutoff c^* exactly balances expected gross profits against fixed costs if other firms enter according to our putative cutoff rule. That is, each firm enters if and only if it draws a marginal cost weakly below c^*.

5. IDENTIFICATION AND ESTIMATION

We now discuss identification and estimation of $\theta = (\theta^r, \theta^d, \theta^s)$, where $\theta^r = (\beta^r, \alpha^r, \lambda^r)$ govern the prominence algorithm; $\theta^d = (\beta, \alpha, \lambda, \rho)$ govern demand; and $\theta^s = (\theta^F, \theta^c)$ govern the fixed F and marginal cost c components of supply. We estimate θ^r and θ^d via maximum likelihood, and θ^s via GMM.

We estimate our model using data covering 19409013 prominence assignments and 803849 sales across 47403 products and 63620 merchants from 2018-08-26 to 2020-03-23. Recall that we observe individual-level sales data and discrete choice data for prominence assignments for each market t, which we define to be a "product (web)page"-notification time pair (w, τ) for the purpose of prominence and demand estimation.[11] In each market, the alternatives $\mathcal{J}_t$ are the various offers on the same product page. This market definition allows us to sidestep estimation of potentially complex preferences over product characteristics: such characteristics do not vary across offers for the same product, so

[9] In our simulations, pricing equilibria always exist but are not necessarily unique as firms differentiate by targeting sophisticated or unsophisticated consumers. When multiple equilibria exist, we select one by starting our price fixed-point search with a lower price for the seller possessing the highest prominence algorithm attractiveness. Further details are in Supplemental Appendix 3.1.

[10] We make this assumption to avoid equilibrium multiplicity, as is common in games of strategic substitutes (e.g., Tamer (2003)). We believe this assumption is innocuous because most sellers in our sample are small and anonymous; they are unaware of their competitors before entry. Even if they know each other's identities, turnover is high. Therefore, merchant entry is best modeled with firms expecting to play against random draws from the distribution of potential opponents.

[11] When considering pricing and entry, we aggregate to a "product page"-month.

they cancel out in the discrete choice problem. Nevertheless, the offers may still differ in characteristics such as dispatch speed.

5.1. *The Prominence Algorithm*

5.1.1. *Main Model: No Unobserved Quality*

We make the identifying assumption of no unobserved prominence quality; that is, we have the following:

ASSUMPTION 6: *There is no unobserved prominence quality*: $\xi^r_{jt} \equiv 0$.

Intuitively, this assumption reflects the fact that alternatives in our model are offers, and markets are product webpages. Hence, product characteristics (like color) cancel when comparing across alternatives within a market, leaving only offer characteristics (like dispatch speed) to matter in the prominence choice. However, we observe essentially all offer characteristics that the consumer observes, including dispatch speed, seller rating, whether an offer is fulfilled by Amazon (FBA), whether an offer is sold by Amazon, and the offer's price. These characteristics are accounted for in x_{jt} and do not need to be captured by the residual ξ^r_{jt}. Assumption 6 hence mainly rules out consumers using the seller's name to draw inferences about seller quality (e.g., based on prior experiences with the seller) not predicted by observables. Such a violation seems unlikely except in the case of Amazon (the merchant), which we control for by including an Amazon dummy. The remaining third-party sellers on this marketplace are numerous and often less established, and Amazon's interface makes it hard to directly search for products sold by a particular seller.[12]

We estimate the prominence parameters $\theta^r = (\beta^r, \alpha^r, \lambda^r)$ via standard maximum likelihood; the estimator $\hat{\theta}^r$ solves

$$\max_{\theta^r} l^r(\theta^r) = \sum_{t=1}^{T} \sum_{j \in \mathcal{J}_t \cup \{0\}} y^r_{jt} \ln r_{jt}(\theta^r), \tag{8}$$

where y^r_{jt} is an indicator for whether offer j was prominent in market t and $r_{jt}(\theta^r)$ is the model-predicted prominence probability from Equation (5) but with $\xi^r_{jt} \equiv 0$.

We use maximum likelihood instead of generalized method of moments for two reasons. First, our data are high-frequency and highly disaggregated; Berry's (1994) inversion would suffer severe zero-share bias (Gandhi, Lu, and Shi (2023); Quan and Williams (2018)). Second, Amazon Marketplace is highly dynamic, with frequent price changes and merchant turnover, so the time dimension provides rich, plausibly exogenous variation—especially given the frequent delegation of pricing to algorithms reacting to competitors' prices. Aggregation, the usual strategy to address zero-share bias, would thus smooth precisely over the most informative variation.

We briefly discuss identification, which is standard (Train (2009)) and also confirmed by Monte Carlo exercises in Supplemental Appendix 4.1. The assumption that $\xi^r_{jt} \equiv 0$ implies that characteristics are exogenous (i.e., $\mathbb{E}[\xi^r_{jt}|x_{jt}] = \mathbb{E}[\xi^r_{jt}]$). Hence, the weights β^r that the algorithm places on various offer characteristics are identified by the frequency

[12]Additionally, if Amazon observes fulfillment metrics beyond x_{jt}, then $\xi^r_{jt} \neq 0$; this violation biases our estimates only if those metrics are correlated with x_{jt} or ξ_{jt}.

at which offers with these characteristics are featured. Furthermore, the same assumption implies that prices are exogenous (i.e., $\mathbb{E}[\xi^r_{jt}|p_{jt}] = \mathbb{E}[\xi^r_{jt}]$), and hence the price sensitivity α^r is identified by the covariance between an offer's price and prominence. Finally, it implies that the covariance between unobserved quality across offers within the "inside options" nest is zero (i.e., $\text{Cov}(\xi^r_{jt}, \xi^r_{kt}) = 0$ for $k \neq j$), and hence the nesting coefficient λ^r is identified by variation in choice sets—for example, variation in the offer number.

5.1.2. *Robustness: Time-Invariant Unobserved Quality*

To address concerns that price may be endogenous with respect to unobserved quality (i.e., that $\mathbb{E}[\xi^r_{jt}|p_{jt}] \neq \mathbb{E}[\xi^r_{jt}]$), we exploit the panel structure of our data. In estimation, a market t will be a product webpage w at a particular time τ, that is, $t \equiv (w, \tau)$. Hence, we can re-estimate the prominence model under the following:

ASSUMPTION 7: *Unobserved prominence quality is time-invariant*: $\xi^r_{jw\tau} \equiv \xi^r_{jw}$.

This alternative assumption is plausible as essentially all variation in offer characteristics is cross-sectional rather than temporal.[13] Furthermore, it explicitly allows for some endogeneity in prices (i.e., $\mathbb{E}[\xi^r_{jw}|p_{jt}] \neq \mathbb{E}[\xi^r_{jw}]$). To operationalize the assumption, we employ offer fixed effects to concentrate out any time-invariant sources of price endogeneity.[14] However, we find below that doing so barely moves the estimate of price sensitivity, suggesting that—as expected—unobserved quality is only a minor concern in our empirical setting.

5.2. *Consumer Choice*

5.2.1. *Main Model: No Unobserved Quality*

We make the identifying assumption of no unobserved demand quality, that is, the following holds:

ASSUMPTION 8: *There is no unobserved demand quality, that is*, $\xi_{jt} \equiv 0$.

As for prominence, this assumption reflects the fact that alternatives in our model are offers, and markets are product webpages; see Section 5.1.1.[15]

We estimate the demand parameters $\theta^d = (\beta, \alpha, \lambda, \rho)$ by maximum likelihood. However, instead of observing one discrete choice per market, we observe the total number of sales $y^d_{j^o(t),t}$ for one merchant $j^o(t)$ on each market t. If we observed arrivals, we could immediately move on to the likelihood implied by our discrete choice model. As we do not observe arrivals, we utilize our assumption that arrivals are Poisson distributed, which implies that sales are also Poisson distributed. To make progress, we need to know the

[13]The R^2 from regressing time until dispatch, log feedback count, being fulfilled by Amazon, and being sold by Amazon on offer fixed effects are 0.97, 1.00, 0.98, and 1.00, respectively.

[14]The details of introducing these fixed effects into a maximum likelihood procedure without causing an incidental-parameters problem are relegated to Supplemental Appendix 2.1. Essentially, we apply the conditional logit approach of Chamberlain (1980) to the conditional logit model itself (see also Rasch 1960, 1961). This approach yields an estimator that exclusively exploits within-offer price variation to identify the price coefficient α^r.

[15]Assumption 8 may be violated if a manufacturer sells its own product and consumers infer higher quality; we address this concern in Section 5.2.2.

expected number of arrivals M_t on each market $t \equiv (w, \tau)$; to reduce noise, we assume arrival rates do not vary over time, so $M_{w,\tau} \equiv L_{w,\tau} \times \text{Arr}_w$, where $L_{w,\tau}$ measures a market's duration in months (which is observed) and Arr_w is the number of monthly arrivals. We exploit data on each market t's sales rank, translating monthly average ranks into monthly predicted sales using estimates from *AMZScout*, a "leading market intelligence provider for Amazon sellers" (Gutiérrez (2021)).[16] We average across months to generate average monthly predicted sales for product webpage w. To turn predicted sales into arrivals $\hat{\text{Arr}}_w$, we use data on conversion rates of sponsored product listings, predicting each webpage's conversion rate based on its suggested retail price (higher prices convert at lower rates).[17]

Conditional on the number of arrivals $\hat{M}_t$ and the true prominence dummy y^r_{jt}, sales y_{jt} for offer j in market t follow a Poisson distribution with mean $\hat{M}_t s_{jt}(y^r_{jt}; \theta^d)$. Here, $s_{jt}(y^r_{jt}; \theta^d)$ is as specified in Equation (3) but with unobserved quality ξ_{jt} set to zero. We thus have a log-likelihood of

$$\sum_{t=1}^{T} \sum_{j=1}^{J_t} \left\{ y_{j,t} \ln\left(\hat{M}_t s_{j,t}\left(y^r_{jt}; \theta^d\right)\right) - \hat{M}_t s_{j,t}\left(y^r_{jt}; \theta^d\right) - \ln \Gamma(y_{j,t} + 1) \right\}, \tag{9}$$

where the last term is constant in θ^d and hence dropped below.

As written, we cannot utilize this likelihood because (i) we only observe the sales y_{jt} for one merchant $j^o(t)$ on each market t, and (ii) y^r_{jt} is only observed at the time of the last change to the set of offer characteristics $\{x_{jt}, p_{jt}\}_{j \in J_t}$ and may have changed since then.[18]

We now tackle these issues in turn. First, we note that thinning a Poisson process by classifying each arrival via a multinomial draw yields independent Poisson processes for each category (i.e., merchant j) (He and Wang (1990)). Hence, we can factor out and integrate over the unobserved merchants' sales to yield an integrated likelihood

$$\sum_{t=1}^{T} \left\{ y_{j^o(t),t} \ln\left(\hat{M}_t s_{j^o(t),t}\left(y^r_{j^o(t),t}; \theta^d\right)\right) - \hat{M}_t s_{j^o(t),t}\left(y^r_{j^o(t),t}; \theta^d\right) \right\}. \tag{10}$$

Crucially, via $s_{j^o(t),t}(y^r_{j^o(t),t}; \theta^d)$, this integrated likelihood still depends on the full set of offer characteristics $\{x_{jt}\}_{j \in J_t}$ including for $j \neq j^o(t)$—for example, the likelihood can account for sales being lower for observed offers when they are in the presence of an attractively priced unobserved competitor by raising the coefficient on price.

Second, we address potential measurement error in the prominence dummy y^r_{jt}, which may be stale by the time of sale. Following the measurement error literature, we use an instrument that satisfies an exclusion restriction—affecting the outcome only through the

[16] These estimates are based on the power law relationship between sales and ranks that has been confirmed repeatedly in the academic literature (Chevalier and Goolsbee (2003), Chevalier and Mayzlin (2006), Brynjolfsson, Hu, and Smith (2006)) but also account for the steepening of the slope in the tail that recent papers have found (Brynjolfsson, Hu, and Simester (2011)). We are grateful to German Gutiérrez for providing us with these data.

[17] The predicted conversion rates range from 10.86% (at a retail price of \$99.75) to 20.56% (at \$5.29). We are grateful to Chuan Yu for providing us with these data.

[18] In auxiliary analyses graciously provided by Andreas Haupt, we find suggestive evidence that within fifteen minutes of the most recent change, the prominence assignment had already changed again in 52.39% of cases. Most of these high-frequency changes are unobserved by us as we only observe prominence assignments when there is a change to the offer set (e.g., a price change).

mismeasured covariate (Carroll, Ruppert, Stefanski, and Crainiceanu (2006)). In a maximum likelihood framework, this instrument informs the conditional distribution of the (true) prominence dummy, which we integrate over to obtain a full information maximum likelihood (FIML) estimator (Wooldridge (2003, Section 15.7.3)).

We operationalize this idea by combining (i) the previous nested-logit prominence model, and (ii) an instrument z_{jt} for the prominence dummy that is excluded from consumer preferences: prominence eligibility (with $z_{jt} = 1$ if offer j is eligible to be featured on market t). Jointly, these yield a predicted probability $\hat{r}_{jt} = z_{jt} \times r_{jt}(\hat{\theta}^r)$ that offer j is assigned prominence on market t. We then integrate over the remaining uncertainty in prominence to yield a new log-likelihood.[19]

Putting everything together, the maximum likelihood estimator $\hat{\theta}^d$ solves

$$\begin{aligned} &\max_{\theta^d} l^d(\theta^d) = \sum_{t=1}^{T} \ln\left(\hat{r}_{j^o(t),t} \exp\left(\ell^d_{1,t}\right) + \left(1 - \hat{r}_{j^o(t),t}\right) \exp\left(\ell^d_{0,t}\right)\right), \\ &\text{where } \ell^d_{r,t} = y_{j^o(t),t} \ln\left(\hat{M}_t s_{j^o(t),t}\left(r; \theta^d\right)\right) - \hat{M}_t s_{j^o(t),t}\left(r; \theta^d\right) \text{ for } r \in \{0, 1\}, \end{aligned} \tag{11}$$

that is, $\ell^d_{1,t}$ and $\ell^d_{0,t}$ are the log-likelihood contributions when the observed offer $j^o(t)$ is, respectively, prominent or not. Supplemental Appendix 4.3 confirms that this estimator performs well in Monte Carlo simulations.

Formally, the first-order conditions of the MLE require

$$\begin{aligned} \sum_{t=1}^{T} \Bigg[&\mathbb{P}(r_{j^o(t),t} = 1 | y_{j^o(t),t}) \left(\frac{y_{j^o(t),t}}{s_{j^o(t),t}\left(1; \theta^d\right)} - \hat{M}_t \right) \frac{\partial s_{j^o(t),t}\left(1; \theta^d\right)}{\partial \theta^d} \\ &+ \mathbb{P}(r_{j^o(t),t} = 0 | y_{j^o(t),t}) \left(\frac{y_{j^o(t),t}}{s_{j^o(t),t}\left(0; \theta^d\right)} - \hat{M}_t \right) \frac{\partial s_{j^o(t),t}\left(0; \theta^d\right)}{\partial \theta^d} \Bigg] = 0, \end{aligned} \tag{12}$$

where $\mathbb{P}(r_{j^o(t),t} = 1 | y_{j^o(t),t})$ is the posterior probability of the observed offer being prominent given $\hat{r}_{j^o(t),t}$ (the prior) and $y_{j^o(t),t}$ (the number of units sold).[20]

Intuitively, the estimator tries to find parameters such that the weighted average prediction errors are orthogonal to the derivative of predicted shares with respect to the parameters. In practice, this means that consumer preferences over characteristics β and price α are identified from the covariance between an offer's sales and its value of that characteristic relative to its competition. For instance, if the observed offer never sells when a faster dispatch offer exists, we infer high dispatch speed sensitivity. Similarly, the preference for Amazon offers is identified by the shadow that the presence of an Amazon offer casts on the sales of observed merchants: if these sales are lower whenever there is an Amazon offer, and if this pattern cannot be explained by prominence alone, then consumers must prefer Amazon offers: the underlying variation used is the same as in Table II.

[19] As prominence is binary, its (conditional) distribution is fully pinned down by the predicted probability $\hat{r}_{jt}$. By contrast, for a continuous variable, we would require full knowledge of the conditional distribution, not just the conditional mean. Formally, this procedure yields a consistent estimator if $\hat{r}_{jt}$ is well-calibrated, that is, $\mathbb{E}[y^r_{jt} | \hat{r}_{jt}] = \hat{r}_{jt}$. This calibration is established by Figure 2a in Supplemental Appendix 5. Practically, the procedure performs well in Monte Carlo simulations in Supplemental Appendix 4.3.

[20] Formally, $\mathbb{P}(r_{j^o(t),t} = 1 | y_{j^o(t),t}) = \frac{\hat{r}_{j^o(t),t} \exp(\ell^d_{1,t})}{\hat{r}_{j^o(t),t} \exp(\ell^d_{1,t}) + (1 - \hat{r}_{j^o(t),t}) \exp(\ell^d_{0,t})}$, and analogously for $r_{j^o(t),t} = 0$.

We caution that this identification strategy is not without limitations. First, if Amazon systematically enters markets with *negative* market-wide demand shocks, we will attribute the lower sales on those markets to a preference for Amazon offers. However, such counterproductive entry behavior seems unlikely,[21] and in any case we control for time-invariant market-wide shifters in the main demand table. Second, one may worry that we cannot separately identify the preference for Amazon offers from the substitutability between Amazon offers and observed offers: if observed offers are much more substitutable with Amazon offers than with other offers, even a slight preference for Amazon offers will cast a large shadow on the observed merchant's sales. Conversely, strong preference with lower substitutability could cast the same shadow. However, such a difference in substitutability would be identified by differing responses of the observed offer's sales to the price of the Amazon offer versus other offers. Furthermore, it seems unlikely that Amazon offers are particularly substitutable with observed offers relative to other offers.

Third, with unsophisticated consumers, attributing preference correctly (to consumers vs. the prominence algorithm) requires a consistent estimate of unsophisticated consumers' prevalence; underestimating this prevalence could misattribute prominence advantage to consumer preference.

To identify the fraction of sophisticated consumers ρ, the procedure described above effectively instruments offer prominence with prominence eligibility, which sellers must purchase with a small monthly payment.[22] The estimator thus effectively targets $\text{Cov}(y^d_{j^o(t),t}, \hat{r}_{jt})$, where $\hat{r}_{jt} = r_{jt}(\hat{\theta}^r) \times z_{jt}$ is the predicted probability of being featured from the prominence model $r_{jt}(\hat{\theta}^r)$ multiplied by a dummy for prominence eligibility z_{jt}. Instrumenting in this way avoids attenuation bias from outdated prominence assignments, and also avoids potential bias from correlation between ξ_{jt} and ξ^r_{jt} (Ursu (2018), Donnelly, Kanodia, and Morozov (2024)): the prominence model assumes $\xi^r_{jt} \equiv 0$ and hence $\hat{r}_{jt}$ does not contain ξ^r_{jt}.

The instrument is relevant by construction (first stage $F = 8367$). Exclusion could be violated if an unobserved offer characteristic valued by consumers is also required for eligibility, that is, if $\mathbb{E}[\xi_{jt}|x_{jt}, z_{jt}] \neq \mathbb{E}[\xi_{jt}|x_{jt}]$. For instance, if merchants build a reputation for on-time delivery with customers and on-time performance matters for prominence eligibility, we may misattribute the reputation-driven increase in sales to being featured. However, recall that we observe essentially all offer characteristics that consumers observe, including seller feedback, dispatch speed, whether an offer is fulfilled by Amazon, whether an offer is sold by Amazon, and its price. Hence, as before, the only plausible channel for endogeneity goes through customers drawing inferences from sellers' names as they remember a particular seller's performance (e.g., for on-time delivery); this seems highly unlikely given the setting.

Finally, as with prominence, the nesting coefficient λ is identified by the slope at which sales decline as a function of the number of offers.

[21] If Amazon enters positive shock markets, we instead underestimate the Amazon preference. Hence, a correction would not affect our conclusion that Amazon is not self-preferencing.

[22] To be eligible, a seller must sign up for a professional account and pay a monthly subscription fee of $39.99 (junglescout.com), then become a Featured Merchant by "hit[ting] certain targets on metrics like excellent customer ratings, competitive pricing, low order defect rate, and more... [but] Amazon won't give you specific numbers" (repricerexpress.com). Empirically, initial transitions into eligibility are sticky (raising eligibility by 76.42%) and not associated with large changes in prices or time to ship (which drop by 0.53% and 0.99 hours, respectively).

5.2.2. *Robustness: Time-Invariant Unobserved Quality*

To address the common concern that price may be endogenous with respect to quality (i.e., that $\mathbb{E}[\xi_{jt}|p_{jt}] \neq \mathbb{E}[\xi_{jt}]$), we exploit the panel structure of our data. In estimation, a market t is defined as a product webpage w observed at a particular time τ, that is, $t \equiv (w, \tau)$. Hence, we can re-estimate the demand model under the following:

ASSUMPTION 9: *Unobserved demand quality is time-invariant, that is,* $\xi_{jw\tau} \equiv \xi_{jw}$.

As with prominence, this assumption is plausible because observed offer characteristics are essentially time-invariant. In Supplemental Appendix 2.2, we show how the high-frequency nature of the data allows us to condition out fixed effects even when only one merchant's sales are observed. In particular, we focus on pairs of periods between which all except the observed merchant's offer characteristics (including prices) stay fixed. Then, we compare the observed merchant's sales in periods where her price is high to her sales in periods where her price is low. However, because sales are infrequent, this comparison mostly captures substitution between the observed merchant and the outside option.

5.3. *Estimating Wholesale and Fixed Costs*

To ensure there are no negative realizations of costs, we assume both fixed and marginal costs follow a lognormal distribution, that is, the following holds:

ASSUMPTION 10—Costs DGP: *For all third-party offers*,

$$\begin{bmatrix} (c_{jt})_{j\in\mathcal{N}_t} \\ F_t \end{bmatrix} = \begin{bmatrix} \left(\theta_0^c + x_t'\theta_x^c\right)\left(\epsilon_{jt}^c\right)_{j\in\mathcal{N}_t} \\ \left(\theta_0^F + x_t'\theta_x^F\right)\epsilon_t^F \end{bmatrix}, \quad \textit{with } \left(\begin{bmatrix} \left(\epsilon_{jt}^c\right)_{j\in\mathcal{N}_t} \\ \epsilon_t^F \end{bmatrix} | x_t\right) \sim \text{Log}\,N\left(0, \begin{bmatrix} \theta_\sigma^c \mathbf{I}_{|\mathcal{N}_t|} & 0 \\ 0 & \theta_\sigma^F \end{bmatrix}\right). \tag{13}$$

For the platform's offer,

$$c_{Amz,t} = \left(\theta_{0,Amz}^c + x_t'\theta_{x,Amz}^c\right)\epsilon_{Amz,t}^c, \textit{ with } \epsilon_{Amz,t}^c \sim LogN\left(0, \theta_{\sigma,Amz}^c\right). \tag{14}$$

This parameterization explicitly allows Amazon to face marginal costs that differ from those of third-party sellers. We collect the supply parameters in $\theta^s = (\theta^c, \theta^F)$, where $\theta^c = (\theta_0^c, \theta_x^c, \theta_\sigma^c, \theta_{0,\text{Amz}}^c, \theta_{x,\text{Amz}}^c, \theta_{\sigma,\text{Amz}}^c)$ and $\theta^F = (\theta_0^F, \theta_x^F, \theta_\sigma^F)$.

Identification of θ^c follows the usual approach (Bresnahan (1981)). Together, the demand model, the assumption of Nash–Bertrand competition, and variation in prices jointly identify wholesale cost parameters. In particular, the mean price across markets provides a moment that speaks to θ_0^c; its covariance with market characteristics x_t gives us θ_x^c; and the within-market price variance tells us about θ_σ^c. Their analogs for the platform offer speak to $\theta_{0,\text{Amz}}^c$, $\theta_{x,\text{Amz}}^c$, and $\theta_{\sigma,\text{Amz}}^c$, respectively.

Meanwhile, identification of θ^F follows conventional practice in static entry models (Berry and Reiss (2007)). At the entry stage, fifteen potential entrants simultaneously decide whether to enter, each knowing only its own wholesale costs and the fixed entry cost. Variation in the expected variable profits implied by the demand model—especially as driven by market size—allows us to draw conclusions about the fixed cost parameters θ^F. The mean number of entrants is informative about θ_0^F; its covariance with market characteristics tells us about θ_x^F; its variance speaks to θ_σ^F.

In theory, marginal cost shocks across (merchant, market) pairs ϵ^c_{jt} may depend on fixed cost shocks for that market ϵ^F_t. Following the literature on empirical two-stage games,[23] the covariance matrix in Assumption 10 imposes that, conditional on observables (e.g., MSRP, number of arrivals), marginal cost shocks are mean independent of fixed cost shocks. Conditioning on observables rules out many possible sources of correlation.[24]

As pricing and entry decisions are unlikely to be made at higher frequencies—and as θ^c and θ^F do not need to be identified from high-frequency variation—we aggregate to the monthly level, defining a market t as a product webpage in a given month.

We estimate $\hat{\theta}^s$ by GMM. In particular, the estimated $\hat{\theta}^s$ sets the simulated moment vector $\hat{G}(\theta^s)$ close to zero:

$$\hat{\theta}^s = \arg\min_{\theta^s \in \Theta^s} Q(\theta^s) = \hat{G}(\theta^s)' W \hat{G}(\theta^s), \quad \hat{G}(\theta^s) = \frac{1}{T}\sum_{t=1}^{T}\left(y_t - \widehat{\mathbb{E} f_t(\theta^s)}\right). \tag{15}$$

The weight matrix W is obtained by bootstrapping a covariance matrix of the underlying moments.

The moments $y_t = (y^F_t, y^c_t)$ are partitioned into "fixed cost" and "wholesale cost" moments, respectively, with $\widehat{\mathbb{E} f_t(\theta^s)}$ their simulated model analogs, where the expectation $\mathbb{E}$ here is taken over simulation draws. The "fixed cost" moments y^F_t are specified as

$$y^F_t = \left(\bar{J}_t, \mathrm{Var}(J_t), \overline{J_t \times A_t}, \overline{J_t \times Z_{\ln(\mathrm{MSRP}),t}}, \left(\overline{\mathbb{1}\{J_t = j\}}\right)_{j=1,2,14,15}\right)', \tag{16}$$

which includes the mean and variance of the number of entrants across markets; the mean, across markets, of the number of entrants on a product multiplied by its market size A_t; an analogous mean but with standardized MSRP (in logs) $Z_{\ln(\mathrm{MSRP}),t}$; and finally the fractions of markets with 1, 2, 14, and 15 entrants.

Meanwhile, the wholesale cost moments, for both third-party offers and Amazon offers, take the form

$$y^c_t = (\hat{\mu}_P, \hat{\delta}_P, \hat{\sigma}_P, \hat{\mu}_{P,\mathrm{AMZ}}, \hat{\delta}_{P,\mathrm{AMZ}}, \hat{\sigma}_{P,\mathrm{AMZ}})' \tag{17}$$

and are obtained via indirect inference. In both the data and model-simulated markets, we first obtain $\hat{\sigma}_P$ by computing the standard deviation of log prices across offers. Next, we obtain $\hat{\sigma}_{P,\mathrm{AMZ}}$ by computing the standard deviation of the price of the Amazon offer, relative to MSRP, across markets. Finally, for $(\hat{\mu}_P, \hat{\delta}_P)$, we run regressions (at the market level) predicting mean third-party offer prices as a function of a constant and product-level covariates, that is,

$$\bar{p}_t = \hat{\mu}_P + x^{P\prime}_t \hat{\delta}_P + \epsilon^P_t, \tag{18}$$

and we minimize the distance between the model-implied regression coefficients and the data-implied ones. The covariates for third-party offers only include the MSRP (standardized across products). To identify $(\hat{\mu}_{P,\mathrm{AMZ}}, \hat{\delta}_{P,\mathrm{AMZ}})$, we run an analogous regression

[23] See, for example, Eizenberg (2014), Holmes (2011), Houde, Newberry, and Seim (2023), Kuehn (2018), Rossetti (2018), Wollmann (2018). The exposition in Bontemps, Gualdani, and Remmy (2023) is particularly salient in our setting: "Introducing such a correlation is possible in principle, but numerically challenging, because it would break the separation between the two stages in the estimation procedure. We view the absence of correlation as a reasonable simplification."

[24] Supplemental Appendix 4.4 performs Monte Carlo simulations to validate that our procedure recovers the true parameters if the data-generating process (DGP) is correctly specified.

for first-party offers, but include the price of the market's cheapest offer (possibly Amazon's own offer) as an additional covariate. This relationship between Amazon's price and the cheapest offer's price is a useful proxy for the intensity of competition.[25]

Estimation is computationally intensive. To evaluate the objective for a given guess of θ^s, we must solve the entry game for each market. However, to solve a single entry game, we must find its associated wholesale cost cutoff c_t^*. Doing so requires computing equilibrium profits under many candidate wholesale cost values. Finally, we must evaluate many candidate parameters θ^s during our outer GMM optimization procedure. We employ several techniques to lighten the computational burden; see Supplemental Appendix 3 for technical details.

6. RESULTS

6.1. *Prominence Algorithm*

We present results in Table III, noting that all reported coefficients on offer characteristics have been normalized by dividing through by λ (i.e., they measure preferences between inside offers). While (1) assumes a (non-nested) logit model (i.e., imposes $\lambda = 1$), (2) and (3) employ a nested logit model. Finally, (4) uses the weaker identifying Assumption 7, yielding estimates robust to arbitrary correlation of an offer's average unobserved quality with its price.

We begin by discussing the price sensitivity of prominence assignments. First, correlation in the inside goods' utilities matters: moving from logit (1) to nested logit (2) doubles the estimated price coefficient. Second, including observable quality measures in (3) further increases our estimate. Third, however, additionally including offer fixed effects in (4) does not affect our estimate much.

We emphasize a few additional findings about our preferred model (3). First, the implied price elasticity for the prominence algorithm is very high, at about -21. This finding is corroborated by our discussions with sellers on the platform.[26] Second, as $\hat{\lambda} \approx 0.09 < 1$, Amazon considers offers more substitutable with each other than with the outside option. This is expected, as the outside option yields no intermediation fees. Finally, quality also matters. FBA offers have a 12% price advantage, and Amazon's offers have an additional advantage equivalent to 4.5% (on top of being FBA themselves). Slow shippers are penalized: a one-day longer dispatch time carries a penalty similar to Amazon's advantage.

In the Supplemental Appendix, we assess the out-of-sample fit of our preferred model (3) (Figure 2, Appendix 5.1) and conclude it fits well. We also explore heterogeneity (Table VI, Appendix 6.1), finding the algorithm similar across categories but potentially less price-sensitive for Office and Pet products.

6.2. *Consumer Choice*

Our consumer choice estimation results are in Table IV, noting that all reported coefficients on offer characteristics have been normalized by dividing through by λ (i.e., they measure preferences between inside offers).

[25] While this moment is not strictly necessary for identification, in practice, we find that it improves power: it helps the estimator distinguish prices that are low because of low costs from prices that are low due to competition.

[26] Even more intriguingly, on its seller interface, Amazon explicitly tells the seller which of her "listings [...] are priced not more than 5% above the Buybox." This statement is consistent with our finding that offers priced more than 5% above the cheapest offer are (almost) never featured.

TABLE III

PROMINENCE ALGORITHM ESTIMATION RESULTS.

	(1)	(2)	(3)	(4)
Price / MSRP	−7.60	−15.24	−25.17	−21.18
	(0.162)	(0.316)	(0.790)	(1.577)
Time to Ship (in Days)			−1.10	
			(0.090)	
100 × Log(# Feedback)			5.23	
			(0.545)	
FBA			2.91	
			(0.161)	
Amazon			1.14	
			(0.073)	
Inside	9.55	4.39	4.39	
	(0.201)	(0.128)	(0.118)	
Nesting (λ)		0.14	0.09	
		(0.009)	(0.006)	
Offer FE?	✗	✗	✗	✓
Impl. Elasticity	−6.87	−13.30	−20.92	n/a

Note: Estimates from maximum likelihood estimation on 0.1% sample. Reported coefficients measure the effect of price (normalized by the MSRP), an extra day of time until dispatch, a percentage increase in feedback count, an offer being fulfilled by Amazon, an offer being sold by Amazon, the mean utility of the inside option (relative to the outside option), and the amount of nesting (with $\lambda = 1$ corresponding to no nesting.) The reported elasticity is computed by numerical differentiation and averaged across all observations. It cannot be computed for (4) because we condition out fixed effects. Standard errors are clustered at the product level.

The first five columns restrict $\rho = 1$; that is, they assume all consumers are sophisticated. By comparing models (1) to (2) and (3) to (4), we conclude that the inclusion of observed quality covariates only slightly moves the estimated price coefficient. This finding suggests there is little endogeneity of price with respect to these covariates, implying that correlation between unobserved quality and price is also likely small.

To further assess this hypothesis, we compare a specification with offer-level fixed effects (5) to those without ((1)–(4)); model (5) is the sole model that uses the alternative identifying Assumption 9 (instead of 8). Just as for prominence, we find that including fixed effects, if anything, reduces the estimated price sensitivity.[27]

The remaining columns (6)–(8) freely estimate ρ, and agree that $\hat{\rho} < 1$, that is, a substantial fraction of unsophisticated consumers only considers the prominent offer. However, (6) uses *observed* prominence assignments y^r_{jt} in place of *predicted* $\hat{r}_{jt}$; as expected, the resulting measurement error severely attenuates the estimate of the fraction of unsophisticated consumers. When correctly accounting for the measurement error in (7), our estimate of the fraction of sophisticates falls from 71% to 5%. Intuitively, this estimate should be compared with the fraction of product-page visits in which a customer clicks through to the Offer Listing Page; while we do not observe that statistic, the European Commission states that "the percentage of total customer visits to the page where the customer clicked on the link to the Offer Listing Page remained below [0–10]% in the period

[27]Among the non-nested models, (1) and (2) yield higher estimates of price sensitivity than (5). This result goes opposite to what we would expect with the usual endogeneity concern (where price is positively correlated with quality).

TABLE IV

DEMAND ESTIMATION RESULTS.

	No Unsophisticates					With Unsophisticates		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price / MSRP	−1.47	−1.50	−9.70	−12.40	−0.58	−11.89	−2.90	−2.99
	(0.118)	(0.097)	(0.535)	(0.576)	(0.195)	(0.628)	(0.397)	(0.376)
Time To Ship (in Days)		−0.37		−0.80		−0.63	−0.05	−0.07
		(0.053)		(0.070)		(0.079)	(0.161)	(0.152)
100 × Log(# Feedback)		8.66		−5.63		−5.68	−0.67	−0.83
		(1.330)		(1.617)		(1.875)	(3.219)	(3.116)
FBA		0.30		1.50		1.38	1.39	1.37
		(0.119)		(0.158)		(0.179)	(0.247)	(0.230)
Amazon		5.27		2.34		2.65	1.56	1.66
		(0.159)		(0.194)		(0.218)	(0.267)	(0.270)
Inside	−4.15	−4.74	−2.99	−2.97		−2.97	−3.28	−3.30
	(0.134)	(0.207)	(0.089)	(0.082)		(0.081)	(0.070)	(0.076)
× Ever Has Amz								2.22
								(2.709)
Nesting (λ)			0.11	0.10		0.09	0.04	0.04
			(0.013)	(0.009)		(0.010)	(0.041)	(0.038)
Frac. Soph. (ρ)						0.71	0.05	0.05
						(0.009)	(0.003)	(0.003)
Offer FE?	✗	✗	✗	✗	✓	✗	✗	✗
Prom. Elig. Instrument?						✗	✓	✓
Elasticity (Fix. Prom.)	−1.56	−1.60	−8.68	−10.63	n/a	−9.55	−0.98	−1.01
Elasticity (Endog. Prom.)						−11.99	−13.70	−13.72

Note: Estimates from maximum likelihood estimation; (6) uses observed prominence dummies y^r_{jt} in place of predicted $\hat{r}_{jt}$; (7) and (8) use predicted $\hat{r}_{jt}$ and hence account for measurement error in prominence. Reported coefficients have already been normalized by dividing through by λ (i.e., they measure preferences between inside offers.) They measure the effect of price (normalized by the MSRP), an extra day until dispatch, a percentage increase in feedback amount, an offer being fulfilled by Amazon, the mean utility of the inside option, and the change in that mean utility for markets which Amazon ever is present on. λ is the nesting coefficient for sophisticated consumers and ρ is the fraction of consumers that are sophisticated. Elasticity Fix. Prom. refers to the average price elasticity of demand assuming that price does not influence consideration sets; Elasticity Endog. Prom. allows for this influence. All standard errors are clustered at the product level.

2017–2019" (European Commission (2022, p. 10)), giving us confidence in our estimate from (7).

Finally, specification (8) allows the average attractiveness of inside options to differ depending on whether an Amazon offer is ever present on a product. This is a robustness check to ensure our estimated preference for Amazon is not driven by market-level unobservables. This is the case: if anything, the preference for Amazon offers strengthens in this robustness check.

Our preferred specification (7) incorporates our consideration set model (leaving ρ unrestricted) and observable quality covariates. The price elasticity is -14, but drops to just -1 if prominence is held fixed as prices vary (i.e., if $\partial r_{jt}/\partial p_{jt}$ is artificially set to zero). This large difference stems from our estimate that 95% of customers consider only the prominent offer. However, customers considering all offers view them as close substitutes

TABLE V

FRACTION OF SALES VIA BUYBOX.

	OLS/2SLS	MLE
Naive Estimation	51.18%	41.7%
	(1.69%)	(0.65%)
Buybox Eligibility Instrument	90.74%	96.09%
	(6.05%)	(0.21%)
FTC Allegation	≈98%	

Note: This table reports the implied fraction of sales via the Buybox. Column 1 estimates this using linear regressions of sales on Buybox status, with and without an eligibility instrument. Column 2 uses our discrete choice model—after estimation, we compare predicted choice probabilities when exogenously varying Buybox status. The fraction is calculated by taking the ratio of predicted sales when in the Buybox to average sales, multiplied by the fraction of time the Buybox is occupied.

($\hat{\lambda} \approx 0.04$).[28] Regarding non-price characteristics, we seem underpowered to detect quality effects beyond FBA and Amazon offers themselves, which are roughly equal: being sold by Amazon is worth 1.12 times as much as being FBA (intriguingly, for prominence this ratio was only 0.39). Again, note that Amazon offers are also FBA, so both advantages accrue to them.

In the Supplemental Appendix, we assess the out-of-sample fit of our preferred model (Figure 3, Appendix 5.1), concluding that it fits well. We also explore heterogeneity (Table VI, Appendix 6.1), finding that markets with more offers have more sophisticated consumers, and that the Amazon advantage is strongest for products in the Pet, Baby, and Office categories.

6.2.1. *Demand Estimates Validate Regulatory Findings*

The European Commission notes that (in France and Germany) "sales made directly through the Featured Offer accounted for [70–80] to [80–90] % of the overall sales" (European Commission (2022, p. 10)).[29] For the United States, the FTC claims that "[n]early 98% of all purchases on Amazon are made using [...] the Buy Box" (Federal Trade Commission and State Attorneys General (2023, p. 29)).

We validate our estimates against these claims in Table V. To this end, we estimate the causal effect of prominence on sales in two ways. First, we regress sales on Buybox status—both directly ("Naive Estimation") and using Buybox eligibility as an instrument. Each regression yields a reduced-form estimate of the effect of Buybox status on sales, allowing us to compute

$$\text{Fraction of Sales via Buybox} = \frac{\widehat{\mathbb{E}}(\text{sales}|\text{Buybox}) \times \widehat{\Pr}(\text{Buybox})}{\widehat{\mathbb{E}}(\text{sales})} \tag{19}$$

to obtain the values in column 1 of Table V.

[28]If the price associated with the outside option (including competing products on Amazon) covaries with offer prices beyond what is captured by the MSRP, this estimate may be biased downwards, which may lead us to under-state the importance of entry below. However, as the inside alternatives are all offers for the same product, it is natural to expect a lot of correlation in the utility of the inside options.

[29]The EC also cites Amazon's own claim that "90%+ of sales [are] com[ing] from the Buy Box" (European Commission (2022, p. 10)).

TABLE VI

FIXED AND MARGINAL COST PARAMETERS.

Parameter	Estimate	Std. Err.	Description
θ_0^F	4.640	(0.015)	mean of log fixed costs
θ_R^F	1.935	(0.029)	.. interacted with MSRP
θ_σ^F	−0.047	(0.019)	variance of log fixed cost
θ_0^c	2.481	(0.019)	marginal cost mean shift
θ_R^c	0.765	(0.004)	marginal cost coefficient on MSRP
θ_σ^c	0.157	(0.002)	variance of marginal cost
$\theta_{0,\text{Amz}}^c$	4.037	(0.042)	Amazon: marginal cost mean shift
$\theta_{R,\text{Amz}}^c$	0.443	(0.011)	Amazon: marginal cost coefficient on MSRP
$\theta_{\sigma,\text{Amz}}^c$	0.101	(0.014)	Amazon: variance of marginal cost

Note: This table displays wholesale and fixed cost parameter estimates obtained from our SMM procedure. The estimated parameters imply an interquartile range for fixed costs of [US$0.37, US$2.83]. For the product at the median MSRP ($21), the marginal cost estimates imply a mean wholesale cost of $18.38 for potential entrants, but only wholesale costs of approximately $17.09 for *successful* entrants. All standard errors are clustered at the product level.

Second, we use our discrete choice model of consumer behavior, comparing model-predicted consumer choice probabilities when an offer is prominent to those when it is not. Averaging over offers, this yields two structural estimates of the effect of prominence on sales (one instrumented, one not), again allowing us to compute (19), but now using the predicted choice probabilities from our model.

Omitting the instrument severely attenuates the share of sales that go through the Buybox. However, with our eligibility instrument, even the linear model recovers a sales fraction of 91%. By imposing the structure of our model, we estimate that 96% of sales on Amazon "go through" the Buybox—very close to the FTC's figure of "[n]early 98%" (ibid.).

6.3. *Marginal and Fixed Costs*

We present fixed and marginal cost estimates in Table VI. For successful entrants, we find wholesale costs of 81% of the manufacturer's suggested retail price (MSRP), which seems plausible for small retailers. Crucially, our estimates suggest Amazon faces significantly lower marginal costs than third-party merchants.

We estimate an interquartile range for fixed costs of [US$0.37, US$2.83]. While these fixed costs appear small, they are consistent with the low barriers to entry on Amazon. For example, 59% of merchants set prices *above* MSRP, with 15% pricing at higher than 1.5 times MSRP. Since the prominence algorithm often directs consumers to the cheapest offer, these high-pricing merchants take very little of the Buybox (and hence market) share.[30] However, to account for the entry of these merchants on the competitive fringe, the model settles on low fixed costs: any larger fixed cost would render entry unprofitable for merchants with such high marginal costs.

That said, given the empirical sales frequency, our estimated fixed costs seem reasonable. These values are comparable to storage costs on Amazon. We observe roughly one

[30]This phenomenon of surprisingly high prices is common on e-commerce websites: Dinerstein, Einav, Levin, and Sundaresan (2018, p. 1855) discuss a similar finding in the context of eBay, and eventually settle on allowing their model to infer high marginal costs for these merchants. Similarly, we allow our model to infer a high variance of marginal cost shocks.

sale per month for the median offer in our sample. Assuming the median merchant is storing a single item (per product), and that this item is less than one cubic foot in volume, storing it in an Amazon fulfillment center would cost about $0.69 per product-month (junglescout.com).

7. COUNTERFACTUALS

With estimates in hand, we investigate three questions. First, we examine the value of improving search guidance. This exercise builds intuition for how the prominence algorithm affects outcomes via its impact on the allocation of offers to unsophisticated consumers, prices, and entry. Second, this naturally leads to the question of how price-sensitive prominence should be, especially given that the platform may have an incentive to prevent sellers from competing with each other to keep intermediation fees high. Finally, we examine if steering toward Amazon's offers constitutes self-preferencing in the sense of harming consumers.

All counterfactuals assume that all markets feature a prominent offer,[31] and refer to an estimation sample of 47,249 products, representing approximately US$161 million in revenue and 465 million arrivals per month.[32] We distinguish two time horizons in our counterfactuals. In the short run, we disallow sellers from adjusting prices or their entry decisions. In the long run, sellers may change both their prices and entry decisions.

7.1. *What Is the Value of Search Guidance?*

We evaluate the impact of search guidance by comparing outcomes under the estimated prominence algorithm to those in a counterfactual where unsophisticated consumers only consider the cheapest offer.[33] While we do not claim this fully reflects consumer behavior in the absence of search guidance, it provides a useful benchmark for illustrating how algorithmic prominence influences market outcomes such as pricing, entry, and welfare.

As shown in Table VII, search guidance generates notable consumer gains even in the short run: relative to a scenario where unsophisticated consumers focus solely on price, the algorithm increases consumer surplus by $19 million (1.7% of status quo consumer surplus). This gain reflects more efficient matching: unsophisticated consumers are steered toward offers with favorable non-price attributes, such as FBA offers, that are otherwise hard to discover. In practice, as we had conjectured, the algorithm steers consumers to the first-party offer: though average sales increase by 1.72 units per product per month, third-party surplus falls slightly while first-party surplus rises.

Once prices and entry are allowed to adjust, the reduced emphasis on price in the Status Quo algorithm, as well as its increased emphasis on non-price dimensions of seller quality, leads entry to increase and prices to fall. Though these effects are modest, they both contribute to an additional consumer surplus increase, raising the total gains from search guidance to $29 million (2.6% of status quo consumer surplus). Adjustments to prices

[31] While we can allow for canceled prominence, doing so leads to other welfare results being overwhelmed by the negative effects of showing an unsophisticated consumer zero offers.

[32] Based on public disclosures, market research firms have estimated that the total revenue attributable to third-party sellers on Amazon Marketplace is $300 billion (marketplacepulse.com). Thus, one could extrapolate our results to the entirety of Amazon Marketplace by multiplying our numbers by $300B/$161M $\approx$ 1,863 to highlight the magnitude of the issues at stake.

[33] To ensure best-response iteration convergence, we assume prominence responds exclusively to price but at a price sensitivity that is only 3x the estimated sensitivity.

TABLE VII

THE VALUE OF STATUS QUO SEARCH GUIDANCE (VS. SEEKING LOWER PRICES).

	Short-Run	Long-Run
Prices & Entry Adjust?	✗	✓
Δ Consumer Surplus (CS)	\$19,106,388	\$28,820,565
Δ Consumer Surplus (Naive)	\$19,106,388	\$28,034,454
Δ Consumer Surplus (Soph.)	\$0	\$786,111
Δ 3P Producer Surplus	−\$312,736	\$365,485
Δ 1P Producer Surplus + Fees	\$2,665,025	\$2,661,831
Δ Welfare	\$21,458,676	\$31,847,881
Δ Mean # Sales/Month	1.72	2.47
Δ Mean Min Price (% MSRP)		−3.02%
Δ Mean % Entrants		3.02%

Note: This table provides the difference in various outcomes, on our sample of 47,249 products, that can be attributed to the estimated prominence algorithm's performance relative to a non-prominence baseline where unsophisticated consumers seek lower prices. This baseline is operationalized by tripling the price coefficient in the algorithm ($\alpha^{r\prime} = 3\hat{\alpha}^r$) and switching off the dependence of the algorithm on observable non-price characteristics. In the short run, sellers cannot change their prices or entry decisions. In the long run, sellers can change both their prices and their entry decisions. Outcomes that cannot change (e.g., the number of entrants in the short run) are omitted for clarity. 3P refers to Third-Party while 1P denotes First-Party, that is, the former outcomes are computed using only non-Amazon offers, while the latter use only Amazon offers. Fees refer to intermediation fees collected by the platform.

and entry allow third-party offers to capture more sales, which increase by 2.47 units per product per month; the platform's surplus is left largely unchanged.

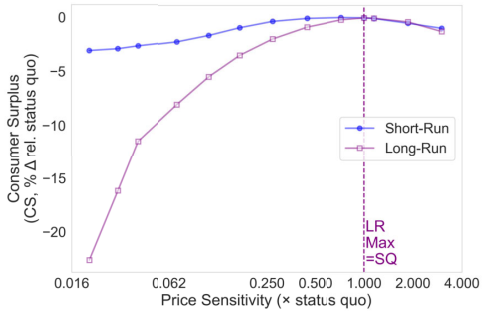
7.2. *How Price-Sensitive Should Prominence Be?*

We found above that the prominence algorithm is highly price-sensitive, with a price elasticity of −21. Furthermore, when comparing the status quo to a world in which unsophisticated consumers only consider the cheapest offer, we found that price sensitivity plays a crucial role in redistributing surplus between producers and consumers through its impact on prices. This effect on prices flows through to Amazon's revenue via ad valorem intermediation fees, suggesting that Amazon may be incentivized to choose an insufficiently price-sensitive prominence algorithm. However, as competition becomes more intense, participation on the platform becomes less attractive for merchants. This tradeoff raises the natural question: how price-sensitive should prominence assignments be?

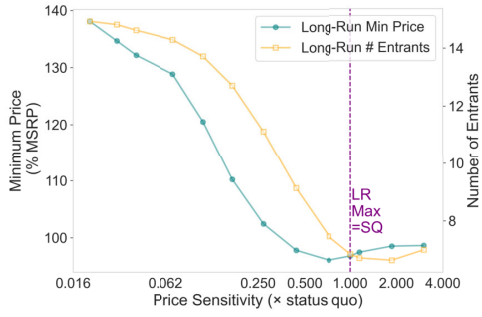
We now consider how a social planner would choose prominence assignment price sensitivity to maximize consumer surplus, accounting for endogenous pricing and entry. To do so, we vary price sensitivity (measured as a multiple of the status quo sensitivity) and illustrate the impacts on consumer surplus in Figure 4a; the corresponding impacts on prices and entry in Figure 4b; the effect on producer surplus in Figure 4c; and the effect on Amazon's profits (revenue from intermediation fees plus profits from first-party sales) in Figure 4d.

As Figure 4a shows, the current price sensitivity is approximately optimal from a long-run consumer surplus perspective. Our results thus do not support the view that Amazon's algorithm is inefficiently price-insensitive. Indeed, further raising price sensitivity reduces consumer surplus in both the short and long run.

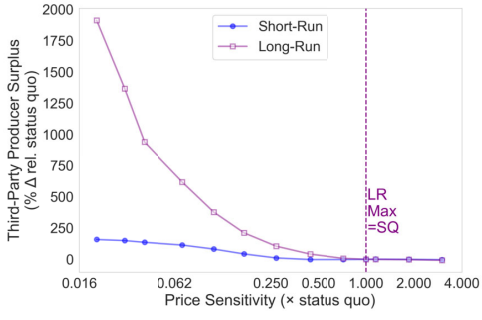
How can higher price sensitivity reduce consumer surplus? First, offers on online marketplaces differ greatly in quality and service dimensions beyond price. Therefore, effectively matching the roughly 95% of unsophisticated consumers to attractive options



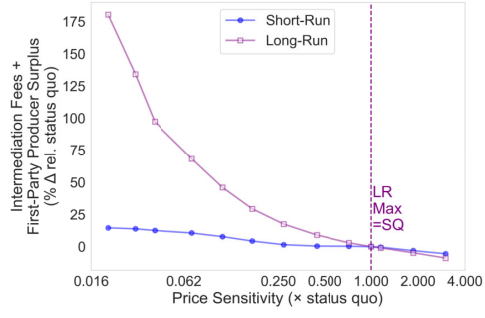
(a) Consumer Surplus



(b) Price and Number of Entrants



(c) Third-Party Producer Surplus



(d) Amazon Fees + First-Party Profit

FIGURE 4.—How price-sensitive should the prominence algorithm be? The top left panel shows the impact on consumer surplus as we vary the weight that Amazon's prominence algorithm places on price, expressed as a multiple of the status quo estimate $\hat{\alpha}^r$. We measure the consumer surplus impact as the % change in consumer surplus relative to the status quo. The top right panel shows the corresponding effects on prices and the average number of entrants. The bottom left and right panels show corresponding impacts on third-party producer surplus and Amazon profits, respectively. The vertical dashed line labeled 'LR Max = SQ' represents the multiple of $\hat{\alpha}^r$ that maximizes the long-run consumer surplus, which coincides with the status quo price sensitivity.

requires the algorithm to weigh factors other than price. This "less-than-infinite" price sensitivity enables, for example, featuring Amazon-fulfilled offers despite slightly higher prices, thus benefiting consumers through better overall matches. This mechanism drives the short-run decline in consumer surplus as price sensitivity grows beyond the status quo (Figure 4a).

Second, driven by the decline in seller surplus illustrated in Figure 4c, an increase in the price sensitivity of prominence also decreases the number of entrants, as illustrated in Figure 4b. This effect is more pronounced at levels of price sensitivity below the status quo, where it mostly reduces the available choices. At higher levels of price sensitivity, it can also make competition less intense, reducing consumer surplus gains from price decreases that would result if entry had been held fixed. This second channel is responsible for the small additional decline in the long-run consumer surplus (relative to the short run) as we move rightwards from the status quo in Figure 4a.

While increasing price sensitivity slightly reduces consumer surplus, decreasing price sensitivity does so much more dramatically, especially in the long run, where it signifi-

cantly raises prices. This is concerning because it provides a mechanism through which Amazon could, in principle, exploit its intermediation power to increase profits at consumers' expense. As on-platform offers are much more substitutable with each other than with off-platform offers, prominence-driven price increases would not immediately prompt consumer switching. This "loyalty" creates room for Amazon to significantly increase revenue from intermediation fees at the expense of consumer surplus. Indeed, Figure 4d shows that decreasing price sensitivity would dramatically increase Amazon's profits, especially in the long run as prices adjust upwards. However, our estimates indicate that, during our sample period, Amazon did not choose to increase its profits by picking such a low weight on price. Furthermore, if Amazon were to regularly feature uncompetitive offers, consumers' willingness to switch to other platforms could eventually rise (e.g., through reduced Prime membership). Such shifts in consumers' substitution patterns fall beyond the scope of our analysis.

7.3. *Is Amazon Self-Preferencing?*

Our prominence algorithm estimates point to the algorithm steering consumers toward Amazon's own offers. However, we also found that consumers prefer Amazon's offers. Therefore, without further analysis, it is unclear whether Amazon's behavior constitutes 'self-preferencing' in the sense of harming consumers.

To investigate, we run counterfactuals that vary the weight that Amazon's prominence algorithm places on its own offers.[34] We report our results in Figure 5. Figure 5a shows the impact on consumer surplus, while Figure 5b shows the corresponding effects on the prominence-weighted price and Amazon's prominence assignment share. Figures 5c and 5d show the effects on producer surplus and Amazon's profits, respectively. The span of weights on Amazon we consider ranges from those implying an Amazon prominence assignment share from a minimum of 62% up to 87%.[35] We indicate both the status quo—in which the additional prominence weight on Amazon offers is 0.39 times that of FBA offers—and the value of $\hat{\beta}^r_{\text{Amazon}}/\hat{\beta}^r_{\text{FBA}}$ that maximizes long-run consumer surplus.

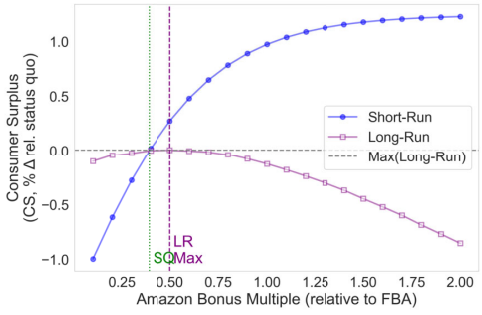
As expected, we find that third-party sellers unambiguously suffer from Amazon's prominence advantage—as illustrated in Figure 5c, their preferred prominence weight on Amazon offers is zero.

By contrast, the short-run consumer surplus (the blue line in Figure 5a) is locally increasing in the prominence weight on Amazon offers near the status quo (marked by the vertical line labeled 'SQ'): as we estimate strong consumer preferences for Amazon offers, featuring these offers more frequently is beneficial. Beyond the plotted span, this effect eventually levels off and reverses as excessive steering toward Amazon offers diverts consumers away from cheaper offers. This reversal, however, does not occur until the prominence weight on Amazon offers reaches about 5.6 times that of FBA offers. Thus, had we ignored pricing and entry decisions and looked only at the short run, we would have concluded that Amazon's current prominence weight on its offers is much too low.

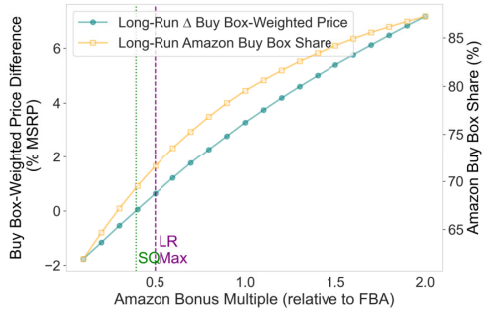
In the long run, however, this finding is not the case. As we see in the left panel, consumer surplus is maximized at a prominence weight on Amazon offers of about 0.5 times that of FBA offers, just slightly above the status quo. For even higher values of

[34] Note we report results on products with an Amazon offer only as other markets are unaffected.

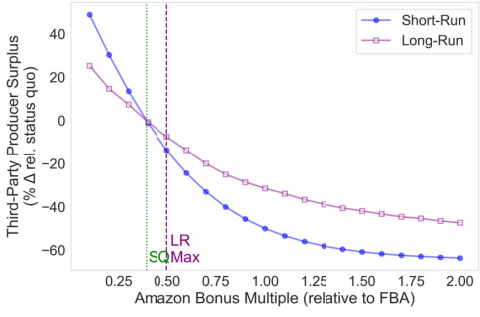
[35] Recall Amazon offers can be featured even if $\beta^r_{\text{Amz}} = 0$, for example, due to competitive prices.



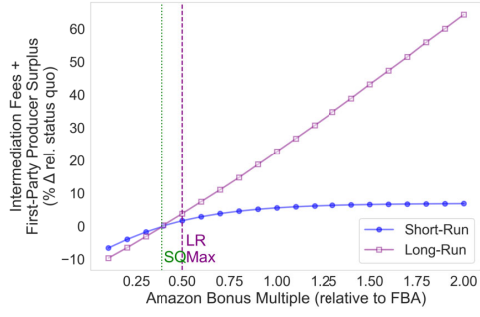
(a) Consumer Surplus



(b) Price & Amazon Rec. Share



(c) Third-Party Producer Surplus



(d) Amazon Fees + First-Party Profit

FIGURE 5.—How much should Amazon guide consumers to its own offers? The top left panel shows the impact on consumer surplus as we vary the weight that Amazon's prominence algorithm places on Amazon's own offers. We measure the consumer surplus impact as the % change in consumer surplus relative to the status quo. The top right panel shows the corresponding effects on prices and the average share of the Amazon offer in prominence assignments, conditional on the Amazon offer being present. The bottom left and right panels show corresponding impacts on third-party producer surplus and Amazon profits, respectively. The vertical dashed line labeled 'LR Max' represents the value of $\hat{\beta}^r_{\text{Amazon}}/\hat{\beta}^r_{\text{FBA}}$ that maximizes the long-run consumer surplus, while the vertical dotted line labeled 'SQ' represents the status quo.

$\beta^r_{\text{Amazon}}/\beta^r_{\text{FBA}}$, consumer surplus is decreasing in the prominence weight on Amazon offers, indicating that Amazon's current prominence weight is about right: Amazon is not self-preferencing in the sense of harming consumers.

What drives the difference between the short and long run? It is not entry, which remains largely stable (and is hence not plotted). Accordingly, our analysis does not support the common theory of harm that Amazon's prominence advantage suppresses entry. Instead, Figure 5b shows that the prominence-weighted price across offers rapidly increases as Amazon features its own offers more frequently. This change is almost exclusively driven by Amazon increasing its prices to exploit its market power: our model assumes that the platform's retail division maximizes short-run profits, as do other merchants. This assumption generates a strong pricing response by the platform to an increase in the prominence weight on Amazon offers. Similarly, in Figure 5d, we see that Amazon's profits increase more rapidly in the long run than the short run as we increase the prominence weight on Amazon offers.

However, this implication is in tension with Jeff Bezos's (admittedly conflicted) public statements that Amazon "[does] price elasticity studies, and every time the math tells us to

raise prices[, but we do not]" (Rose (2013)). Still, Supplemental Appendix 7.2 shows that fixing Amazon's prices at their status quo levels shifts the consumer-surplus-maximizing prominence weight on Amazon offers upward, leaving our conclusion, that Amazon is not self-preferencing, intact.

We make three cautioning observations. First, these results rely on our estimates of the extent to which consumers prefer Amazon offers to those of third-party merchants. As we cannot access sales data for Amazon offers, we must infer these sales from market size estimates and sales on other, non-Amazon offers. Although this procedure yields no formal identification issues, future work could improve on our estimates with better sales data.

Second, we cannot rule out the possibility that Amazon could self-preference via multiple prominence algorithms: when it makes its own offer prominent, the platform could also improve the associated product's visibility on the search results page, inflating sales (and estimated preferences) relative to third-party merchants. However, given recent findings of little to no self-preferencing in search results, this possibility is unlikely to be a major concern (Farronato, Fradkin, and MacKay (2025)).

Third, our model estimates are based on data from 2018 to 2020; however, Amazon can change its algorithm at will. Hence, while Amazon may not have been self-preferencing at this point in time, the company could hypothetically be self-preferencing at present, or in the future.

8. CONCLUSION

Regulators worry that platforms matching buyers and sellers may influence market outcomes by guiding consumer search through algorithmic prominence. We explore this issue by building and estimating a model of intermediation power. Our model captures algorithmic influence on consideration sets yet remains flexible enough to let the data speak to the extent of this influence. Our findings demonstrate the power of steering: on Amazon Marketplace, 95% of customers only consider prominent offers. Through their choice of search, ranking, and prominence algorithms, platforms are indeed able to influence market outcomes, guide consumers toward their own offers, and choose the intensity of price competition. In short, algorithmic prominence gives platforms intermediation power.

However, the mere existence of intermediation power does not imply that platforms exploit it to consumer detriment. Indeed, we find that Amazon's prominence algorithm favors the platform. Yet the self-preferencing question is more nuanced: according to our estimates, consumers also prefer Amazon's offers; our estimated model implies that Amazon is not self-preferencing in the sense of harming consumers. While featuring its own offers more prominently does lead to higher prices, it does not affect entry incentives much; consumers benefit from seeing offers they are likely to purchase. On net, we find that the frequency with which Amazon's algorithm features its own offers is approximately optimal from a consumer surplus perspective: featuring Amazon's offers more frequently would lower consumer surplus via higher prices, whereas featuring them less would steer consumers away from their preferred offers.

Our model also allows us to assess a recently prominent theory of harm: that self-preferencing acts as a barrier to entry. We find no evidence supporting this theory. Nevertheless, entry matters: by choosing the price sensitivity of prominence assignments, the platform can trade off competition among incumbents against entry incentives. While some fear the platform is insufficiently motivated to foster competition because of its ad valorem intermediation fees, our results suggest that it has approximately chosen the consumer-surplus-maximizing price sensitivity.

We discuss future directions. On the demand side, richer sales could reveal *which* consumers are helped or hurt by platform steering, as well as how much they benefit from a "faster" shopping process versus making "better" choices. Future work could also assess the extent to which our attempt to compensate for the shortcomings of our data by, for example, proxying for market size has affected our conclusions. On the supply side, our analysis can be adapted to speak to innovation and the growth of platforms over the long term. In the very long run, merchant quality and investment matter; how these dynamics interact with prominence algorithms is an open question. For instance, Amazon's prominence algorithm incentivizes merchants to adopt FBA—does this incentive eventually harm consumers through higher prices or benefit them through better service? Finally, even large platforms face competitors. Although our entry margin implicitly accounts for merchants' ability to switch platforms, a richer model could explicitly allow both consumers and merchants to multi-home.

REFERENCES

Anderson, Simon P., Andre De Palma, and Jacques-Francois Thisse (1992), *Discrete Choice Theory of Product Differentiation*. MIT Press. [2092]

Bailey, Joseph P. (1998), "Intermediation and Electronic Markets: Aggregation and Pricing in Internet Commerce." Ph.D. thesis. [2083]

Baye, Michael R., John Morgan, and Patrick Scholten (2004), "Price Dispersion in the Small and in the Large: Evidence From an Internet Price Comparison Site." *The Journal of Industrial Economics*, 52 (4), 463–496. [2083]

Ben-Akiva, Moshe and Bruno Boccara (1995), "Discrete Choice Models With Latent Choice Sets." *International Journal of Research in Marketing*, 12 (1), 9–24. [2082,2089]

Berry, Steven, James Levinsohn, and Ariel Pakes (1995), "Automobile Prices in Market Equilibrium." *Econometrica*, 63 (4), 841–890. [2090]

Berry, Steven T. (1994), "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics*, 25 (2), 242–262. [2094]

Berry, Steven T. and Peter C. Reiss (2007), "Empirical Models of Entry and Market Structure." In *Handbook of Industrial Organization*, Vol. 3 (Mark Armstrong and Robert H. Porter, eds.), 1845–1886, Elsevier. [2099]

Blitzstein, Joseph K. and Jessica Hwang (2014), *Introduction to Probability*. Chapman & Hall/CRC Texts in Statistical Science, CRC. Press/Taylor & Francis Group. [2090]

Boldrini, Michela and Francesco Clavorà Braulin (2025), "Taming Tech Giants' Algorithms: An Analysis of Consumers' Choices on Amazon." Working Paper. [2084]

Bontemps, Christian, Cristina Gualdani, and Kevin Remmy (2023), "Price Competition and Endogenous Product Choice in Networks: Evidence From the US Airline Industry." Working Paper. [2100]

Bresnahan, Timothy F. (1981), "Departures From Marginal-Cost Pricing in the American Automobile Industry: Estimates for 1977–1978." *Journal of Econometrics*, 17 (2), 201–227. [2099]

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Duncan Simester (2011), "Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales." *Management Science*, 57 (8), 1373–1386. [2096]

Brynjolfsson, Erik, Yu Jeffrey Hu, and Michael D. Smith (2006), "From Niches to Riches: Anatomy of the Long Tail." *Sloan Management Review*, 47 (4), 67–71. [2096]

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello (2025), "Artificial Intelligence, Algorithmic Recommendations and Competition." SSRN Working Paper No. 4448010. [2083]

Cardell, Scott (1997), "Variance Components Structures for the Extreme-Value and Logistic Distributions With Application to Models of Heterogeneity." *Econometric Theory*, 13 (2), 185–213. [2090]

Carroll, Raymond J., David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu (2006), "Measurement Error in Nonlinear Models: A Modern Perspective." Chapman and Hall/CRC, 2 ed. [2097]

Chamberlain, Gary (1980), "Analysis of Covariance With Qualitative Data." *The Review of Economic Studies*, 47 (1), 225–238. [2095]

Chen, Nan and Hsin-Tien Tsai (2024), "Steering via Algorithmic Recommendations." *The RAND Journal of Economics*, 55 (4), 501–518. [2084]

Chevalier, Judith and Austan Goolsbee (2003), "Measuring Prices and Price Competition Online: Amazon.com and BarnesandNoble.com." *Quantitative Marketing and Economics*, 1, 203–222. [2096]

Chevalier, Judith A. and Dina Mayzlin (2006), "The Effect of Word of Mouth on Sales: Online Book Reviews." *Journal of Marketing Research*, 43 (3), 345–354. [2096]

Ciotti, Fabrizio and Leonardo Madio (2023), "Competition for Prominence." Mimeo. [2084]

Crawford, Gregory, Matteo Courthoud, Regina Seibel, and Simon Zuzek (2022), "Amazon Entry on Amazon Marketplace." CEPR Discussion Paper No. 17531. [2084]

Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu (2018), "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Econometrica*, 86 (3), 891–954. [2083]

De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest (2012), "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review*, 102 (6), 2955–2980. [2083]

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan (2018), "Consumer Price Search and Platform Design in Internet Commerce." *American Economic Review*, 108 (7), 1820–1859. [2082,2083,2089, 2105]

Donnelly, Robert, Ayush Kanodia, and Ilya Morozov (2024), "Welfare Effects of Personalized Rankings." *Marketing Science*, 43 (1), 92–113. [2098]

Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan (2015), "Assessing Sale Strategies in Online Markets Using Matched Listings." *American Economic Journal: Microeconomics*, 7 (2), 215–247. [2083]

Eizenberg, Alon (2014), "Upstream Innovation and Product Variety in the US Home PC Market." *The Review of Economic Studies*, 81 (3), 1003–1045. [2100]

Ellison, Glenn and Sara Fisher Ellison (2009), "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, 77 (2), 427–452. [2083]

Etro, Federico (2024), "E-Commerce Platforms and Self-Preferencing." *Journal of Economic Surveys*, 38 (4), 1516–1543. [2084]

European Commission (2022), "Commission Decision of 20.12.2022 Relating to Cases AT.40462 and AT.40703." Tech. rep., European Commission, accessed: 2025-03-31. [2081,2103,2104]

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay (2023), "Self-Preferencing at Amazon: Evidence From Search Rankings." *AEA Papers and Proceedings*, 113, 239–243. [2085]

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay (2025), "Vertical Integration and Consumer Choice: Evidence From a Field Experiment." Working Paper. [2084,2085,2111]

Federal Trade Commission and State Attorneys General (2023), "Complaint [Public Redacted Version]." Tech. Rep. 2:23-cv-01495-JHC, Federal Trade Commission. [2081,2084,2104]

Gandhi, Amit, Zhentong Lu, and Xiaoxia Shi (2023), "Estimating Demand for Differentiated Products With Zeroes in Market Share Data." *Quantitative Economics*, 14 (2), 381–418. [2094]

German Competition Authority (Bundeskartellamt) (2013), Decision B 6–46/12. [2086]

German Competition Authority (Bundeskartellamt) (2015), Decision B 9–121/13. [2086]

Goeree, Michelle Sovinsky (2008), "Limited Information and Advertising in the US Personal Computer Industry." *Econometrica*, 76 (5), 1017–1074. [2083,2089]

Gutiérrez, Germán (2021), "The Welfare Consequences of Regulating Amazon." Working paper, UW Foster. [2084,2096]

Hagiu, Andrei, Tat-How Teh, and Julian Wright (2022), "Should Platforms Be Allowed to Sell on Their Own Marketplaces?" *The RAND Journal of Economics*, 53 (2), 297–327. [2083]

Hartzell, Olivia and Andreas Haupt (2025), "Platform Preferencing and Price Competition I: Evidence From Amazon." SSRN Working Paper No. 5126918. [2084]

He, Shengwu and Jiagang Wang (1990), "Thinning of Point Processes—Martingale Method." *Probability in the Engineering and Informational Sciences*, 4 (1), 117–129. [2096]

Holmes, Thomas J. (2011), "The Diffusion of Wal-Mart and Economies of Density." *Econometrica*, 79 (1), 253–302. [2100]

Houde, Jean-François, Peter Newberry, and Katja Seim (2023), "Nexus Tax Laws and Economies of Density in e-Commerce: A Study of Amazon's Fulfillment Center Network." *Econometrica*, 91 (1), 147–190. [2100]

Johnson, Justin P., Andrew Rhodes, and Matthijs Wildenbeest (2023), "Platform Design When Sellers Use Pricing Algorithms." *Econometrica*, 91 (5), 1841–1879. [2082]

Khan, Lina M. (2016), "Amazon's Antitrust Paradox." *Yale Law Journal*, 126, 710–805. [2084]

Klingler, Désirée, Jonathan Bokemeyer, Benjamin Della Rocca, and Rafael Nunes Bezerra (2020), "Amazon's Theory of Harm." Working paper. [2081]

Kuehn, Joseph (2018), "Spillovers From Entry: The Impact of Bank Branch Network Expansion." *The RAND Journal of Economics*, 49 (4), 964–994. [2100]

Lam, Honn Tai (2021), "Platform Search Design and Market Power." Working Paper. [2084]

Lee, Kwok Hao and Leon Musolff (2026), "Supplement to 'Two-Sided Markets Shaped by Platform-Guided Search'." *Econometrica Supplemental Material*, 94, https://doi.org/10.3982/ECTA22288SUPP. [2085]

Luco, Fernando and Guillermo Marshall (2020), "The Competitive Impact of Vertical Integration by Multiproduct Firms." *American Economic Review*, 110 (7), 2041–2064. [2083]

Malmendier, Ulrike and Young Han Lee (2011), "The Bidder's Curse." *American Economic Review*, 101 (2), 749–787. [2083]

Manski, Charles F. (1977), "The Structure of Random Utility Models." *Theory and Decision*, 8 (3), 229–254. [2082,2089]

McFadden, Daniel (1978), "Modeling the Choice of Residential Location." In *Spatial Interaction Theory and Planning Models* (Anders Karlqvist, Lars Lundqvist, Folke Snickars, and Jörgen Weibull, eds.), 75–96, North-Holland. [2090]

McManus, Brian, Aviv Nevo, Zachary Nolan, and Jonathan W. Williams (2020), "Steering Incentives of Platforms: Evidence From the Telecommunications Industry." Tech. Rep. 27083. National Bureau of Economic Research. [2083]

Nevo, Aviv (2001), "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, 69 (2), 307–342. [2090]

Quan, Thomas W. and Kevin R. Williams (2018), "Product Variety, Across-Market Demand Heterogeneity, and the Value of Online Retail." *The RAND Journal of Economics*, 49 (4), 877–913. [2094]

Rasch, Georg (1960), *Studies in Mathematical Psychology: I. Probabilistic Models for Some Intelligence and Attainment Tests,* Nielsen & Lydiche. [2095]

Rasch, Georg (1961), "On General Laws and the Meaning of Measurement in Psychology." In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 4, 321–333. [2095]

Raval, Devesh (2022), "Steering in One Click: Platform Self-Preferencing in the Amazon Buy Box." Working paper. [2084]

Reimers, Imke and Joel Waldfogel (2023), "A Framework for Detection, Measurement, and Welfare Analysis of Platform Bias." Tech. Rep. 31766. National Bureau of Economic Research. [2083]

Roberts, James and Andrew Sweeting (2013), "When Should Sellers Use Auctions?" *American Economic Review*, 103 (5), 1830–1861. [2093]

Rose, Charlie (2013), "Amazon." Television broadcast, 60 Minutes. Aired December 1 on CBS. [2111]

Rossetti, Joseph Anthony (2018), "Product Variety in the US Yogurt Industry." Ph.D. thesis. The Ohio State University. [2100]

Schneider, Henry S. (2016), "The Bidder's Curse: Comment." *American Economic Review*, 106 (4), 1182–1194. [2083]

Smith, Michael D. and Erik Brynjolfsson (2001), "Consumer Decision-Making at an Internet Shopbot: Brand Still Matters." *The Journal of Industrial Economics*, 49 (4), 541–558. [2083]

Stigler Committee on Digital Platforms (2019), "Final Report." Tech. rep., Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business. [2081,2085]

Stigler, George J. (1961), "The Economics of Information." *Journal of Political Economy*, 69 (3), 213–225. [2083]

Tamer, Elie (2003), "Incomplete Simultaneous Discrete Response Model With Multiple Equilibria." *The Review of Economic Studies*, 70 (1), 147–165. [2093]

Teh, Tat-How and Julian Wright (2022), "Intermediation and Steering: Competition in Prices and Commissions." *American Economic Journal: Microeconomics*, 14 (2), 281–321. [2083]

Train, Kenneth E. (2009), *Discrete Choice Methods With Simulation*. Cambridge University Press. [2094]

U.S. Congress (2021), "American Choice and Innovation Online Act." In *H.R. 3816, 117th Congress*. [2081, 2085]

U.S. Senate Judiciary Committee (2020), "Investigation of Competition in Digital Markets." Tech. rep., Congress. [2081,2084,2085]

UK Competition Authority (Office of Fair Trading) (2013), "Case CE/9692/12." [2086]

Ursu, Raluca M. (2018), "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions." *Marketing Science*, 37 (4), 530–552. [2082,2098]

Waldfogel, Joel (2024), "Amazon Self-Preferencing in the Shadow of the Digital Markets Act." Tech. Rep. 32299. National Bureau of Economic Research. [2084,2085]

Wollmann, Thomas G. (2018), "Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." *American Economic Review*, 108 (6), 1364–1406. [2100]

Wooldridge, Jeffrey M. (2003), *Econometric Analysis of Cross Section and Panel Data*. MIT Press. [2097]

Yu, Chuan (2024), "The Welfare Effects of Sponsored Product Advertising." Job Market Paper, 2024, Stanford University. [2082]

Co-editor Francesco Trebbi handled this manuscript.

Manuscript received 26 September, 2023; final version accepted 1 June, 2026; available online 5 June, 2026.

The replication package for this paper is available at https://doi.org/10.5281/zenodo.19416797. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices. Given the highly demanding nature of the algorithms, the reproducibility checks were run on a simplified version of the code, which is also available in the replication package.

All authors assume responsibility for all aspects of the paper.