The Welfare Consequences of Regulating Amazon*

Germán Gutiérrez†

November 23, 2021

Click Here for latest version

Abstract

Amazon acts as both a platform operator and seller on its platform, designing rich fee policies and offering some products direct to consumers. This flexibility may improve welfare by increasing fee discrimination and reducing double marginalization, but may decrease welfare due to incentives to foreclose rivals and raise their costs. This paper develops and estimates an equilibrium model of Amazon's retail platform to study these offsetting effects, and their implications for regulation. The analysis yields four main results: (i) Optimal regulation is product- and platform-specific. Interventions that increase welfare in some categories, decrease welfare in others. (ii) Fee instruments are substitutes from the perspective of the platform. Interventions that ban individual instruments may be offset by the endogenous response of (existing and potentially new) instruments. (iii) Regulatory interventions have important distributional effects across platform participants. (iv) Consumers value both the Prime program and product variety. Interventions that eliminate either of the two decrease consumer as well as total welfare. By contrast, interventions that preserve Prime and product variety but increase competition – such as increasing competition in fulfillment services – may increase welfare.

*I am very grateful to Luis Cabral, Thomas Philippon, Chris Conlon and Johannes Stroebel for their generous guidance, support, and encouragement. I am also grateful to the Smith Richardson Foundation, The Stern Center for Global Economy and Business and the Washington Center for Equitable Growth for financial support; and to Viral Acharya, Milena Almagro, Hannah Halaburda, Boyan Jovanovic, Quinn Maingi, Ariel Pakes, Paul Scott, Robert Seamans, Ryan Stevens and seminar participants at NYU Stern (IO and Finance) for helpful comments and discussions.

†New York University

1 Introduction

Amazon is the world's largest retail e-commerce platform with total US sales reaching nearly $300 billion in 2019 or 7.7% of total retail sales. This represents 44% of US e-commerce, compared to 7% for Walmart, 5% for eBay, and 2% for target, the next three (distant) followers. Amazon is also America's second-largest private employer, after Walmart.

An important aspect of Amazon's business model is that it acts as both a marketplace, where third-party sellers can list products for a fee; and as a reseller, where Amazon purchases products from manufacturers and sells them on to retail customers. Most other e-commerce firms tend to specialize as either a reseller (like Target, and Costco) or marketplace (like eBay and AliBaba).[1] One major advantage Amazon possesses is its distribution network, which enables it to deliver packages throughout the United States in two days or less as part of its "Amazon Prime" program. Somewhat unique to Amazon is that it also operates a hybrid business, "Fulfilled by Amazon" (FBA), where third-party sellers set prices while taking advantage of Prime, and for which Amazon charges sellers additional fees for warehousing inventory and shipping orders. It is this third, hybrid method that represents Amazon's fastest growing retail business segment.

The success of Amazon as an e-commerce platform has attracted scrutiny from both antitrust agencies (e.g., Khan, 2016) and lawmakers. Following a 16-month investigation, the House Antitrust, Commercial, and Administrative Law Subcommittee issued a report on Competition in Digital Markets (Congress, 2020). The report concluded that Amazon holds monopoly power over third-party sellers and raised concerns regarding Amazon's ability to: (a) introduce products and services in direct competition with third parties; (b) steer buyers and sellers towards more profitable products and services (e.g., to steer buyers towards Amazon's private label products or to steer sellers towards fulfillment-by-Amazon); and (c) extract "excessive" rents (through high prices or unfair terms) given its (potential) role as consumer gatekeeper. Many of these concerns relate to Amazon's dual role as both seller and platform operator and the inherent "conflicts of interest" and "self-preferencing" issues that arise.

Following its investigation, Congress has put forth several pieces of legislation.[2] If implemented in their strictest form, these bills would force Amazon to either: (a) close its marketplace and become a pure reseller (like, for example, Costco); (b) divest its first-party and logistics operations and become a pure marketplace (like, for example, eBay); or (c) be broken up along business lines (i.e., "structural separations").Amazon's purported response, outlined in recent disclosures, would be to become a pure reseller, effectively foreclosing third-party resellers altogether (Amazon, 2020).

This paper attempts to quantify the welfare effects of Amazon's dual role as both seller and platform operator, and analyze the welfare implications of proposed regulations. To do so, we build and estimate an equilibrium model where consumers choose products on the Amazon platform, while third-party sellers and

[1] Other dual trade platforms include e-commerce players such as AJD.com and Walmart.com; as well as Apple's Appstore, Google's Playstore, Window's Apps, Intuit's Quickbooks Apps, Salesforce's AppExchange, and videogame consoles like Nintendo Switch, all of which sell their own apps alongside third-party apps.

[2] The American Innovation and Choice Online Act prohibits discriminatory conduct by dominant platforms, including a ban on self-preferencing. The Ending Platform Monopolies Act, in turn, prevents big tech companies from "selling products in marketplaces they control." Three additional bills relate to Big Tech broadly and aim to (a) restrict acquisitions, (b) increase data portability and (c) increase funding for antitrust agencies, respectively.

Amazon endogenously set prices of products and platform fees. This allows us to predict how prices, fees and welfare respond to various policy interventions.[3]

Modeling and Estimation. We begin our empirical analysis by estimating a multi-product demand system for each product category, in the tradition of Industrial Organization (Berry, 1994; Berry et al., 1995; Nevo, 2001). The model encodes consumer preferences across products as well as distribution methods (i.e., Prime vs non-Prime). We estimate the model using data from Keepa.com, which tracks product level prices, distribution methods and sales ranks over time.[4]

The size and scale of Amazon's operation, as well as the nature of our dataset present several challenges for demand estimation. First, we observe a single national market with no data on geographic or consumer heterogeneity, so that recovering cross-elasticities through something like Berry et al. (1995) is not feasible. We instead follow Almagro and Manresa (2021) and recover cross-elasticities by grouping products into nests via K-means clustering. Second, estimation is complicated by the high prevalence of low or zero market share products in e-commerce (Quan and Williams, 2018). We use the methodology of Gandhi et al. (2020) to recover unbiased parameters. Last, we confront various endogeneity issues regarding prices and Prime eligibility.

The estimated own and cross-price elasticities are high: the quantity-weighted average price elasticity ranges from -3 to -6 across categories, with the median being even higher. The within-nest correlation parameters range from 0.7-0.8, which implies that consumers have strong preferences for both the platform and their nests. The implied outside good diversion ratios — which measure the portion of consumers that leave the platform, rather than switch to an inside product when prices rise — fall in the 15-30% range, implying relatively inelastic category-level demands between -0.5 and -1.5.

The strategic interactions between Amazon and third-party sellers are more complicated. From the perspective of third-party sellers, Amazon levies a tax on revenue. Products fulfilled by Amazon see both higher demand (via consumer preferences for the Prime program) and an additional per unit "tax". Perhaps most importantly, the level of these "taxes" varies significantly across categories and within categories over time. Ad-valorem referral fees, for example, range from 6% to 45% depending on the product category. With regards to third-party sellers, Amazon's problem mirrors the literature on optimal commodity taxation and depends critically on the extent to which platform fees are passed-through by sellers to retail prices.

In most product categories, Amazon also acts as a horizontal competitor to third-party sellers either by purchasing products from manufacturers and selling them directly, or through its own private-label brands (such as AmazonBasics). Our model captures the incentives Amazon has to (a) raise third-party rivals costs, (b) steer consumers to Amazon's own listings, and (c) use house-brands to discipline seller market power and reduce "double marginalization".[5]

[3]While the focus of the paper is on Amazon, the insights from this paper apply more broadly to both traditional and trade platforms.

[4]Buyers use Keepa to track the price of desired items, while sellers use it to identify new product markets to enter or study competing products. Our complete dataset covers nearly 4,000 product categories and includes over 8 million products across 8 countries.

[5]These kinds of incentives have been recently studied in other vertical arrangements. In particular, Crawford et al. (2018) investigates the welfare effects of vertical integration of regional sports networks. They show that welfare consequences are ambiguous

Our estimates suggest that, since relatively few consumers leave the platform when Amazon raises seller fees, Amazon might be able to increase short-run profits by increasing fees further. One likely explanation is that Amazon continues to invest in consumer loyalty today, but could begin to "harvest" customers in the future. In order to capture these dynamic incentives, we allow Amazon to place a positive weight on consumer surplus in their objective function.[6]

We estimate that Amazon places seven times more weight on consumer surplus than on its own surplus, on average. Nonetheless, our estimated Amazon mark-ups and fees are substantial: an average profit margin of 15% on third-party sales and 28% on first-party sales, compared to 31% for third-party sellers and 45% for wholesalers. Interestingly, we find that Amazon rarely charges a mark-up on fulfillment services. This is because estimated pass-through on unit taxes is higher than on ad-valorem taxes (Anderson et al., 2001), so the platform would rather subsidize unit fees and set higher ad-valorem fees (especially when the weight on consumer surplus is high).

Counterfactuals. Equipped with an estimated model, we perform a series of counterfactuals that cast light on the welfare consequences of proposed regulation.[7]

The first counterfactual estimates the new equilibrium when Amazon turns into a pure marketplace (like eBay). Product variety is maintained but the gains from the Prime program are lost. The loss of utility from Prime induces Amazon to lower its fees for most categories.[8] The decline in fees is not enough to offset the loss of prime, however, so that consumer and platform welfare falls. Third-party sellers benefit from (a) lower fees and (b) the ability to offer products previously sold by Amazon. Wholesalers lose since they previously benefitted from Prime.

The second counterfactual estimates the new equilibrium when Amazon becomes a pure reseller. The loss of product variety from foregone third-party sellers lowers consumer preferences for the platform which, again, induces Amazon to lower its fees. However, even with lower fees consumer welfare is lower, as is the platform's. Exiting third-party sellers (and their manufacturers) lose, while manufacturers of products that remain on the platform benefit both from reduced competition and lower retailer markups.

The third counterfactual corresponds to banning reselling but preserving Prime by expanding the Fulfillment-by-Amazon program. The welfare consequences depend on the platform's investment incentives: when Amazon is investing in (resp. harvesting) consumers, it uses reselling to lower (resp. raise) prices. Therefore, banning reselling decreases (resp. increases) consumer surplus.

Finally, the last counterfactual considers the (forced) introduction of a Seller Fulfilled Prime program, which requires Amazon to give the Prime checkmark to any seller who consistently meets pre-defined performance metrics. This preserves the gains from Prime but enables competition in fulfillment services.[9] We

and depend on the market structure.

[6] A similar weight is used in (Conlon and Mortimer, 2021). Intuitively, it captures the fact that Amazon considers the possibility that lower buyer surplus may lead platform participants to switch to a different platform.

[7] Here we focus on the main counterfactuals. Additional ones are presented in Section 7.

[8] In selected categories, reselling was used to raise rather than lower prices. Banning reselling, then, results in an increase in Amazon fees on third-party products. This further exacerbates the welfare decline.

[9] Such a program previously existed, but it has been progressively made more and more stringent by Amazon over the past few years.

model this by assuming that the increased competition pushes Amazon's mark-up on fulfillment services to zero. Since Amazon would often prefer to set a negative mark-up on unit fees (see above), this policy is non-binding for most categories. Welfare increases in the categories where the policy binds but, interestingly, referral fees endogenously increase after the regulation, offsetting a significant portion of the gains from intervention.

The rest of the paper is organized as follows. Section 2 provides a literature review. Section 3 provides a brief overview of the Amazon retail platform. Section 4 and 5 develop the model and estimation strategy, respectively. Section 6 describes the data used to estimate the model. Section 7 discusses the estimated parameters and their implications for regulation. Section 8 discusses the counter-factual analyses and results. Section 9 concludes.

2 Related Literature

This paper contributes to a long literature on platform theory (Rohlfs, 1974; Katz and Shapiro, 1985, 1994; Farrell and Katz, 2003; Rochet and Tirole, 2003; Armstrong, 2006). In fact, our marketplace model can be viewed as a special case of the canonical framework of Rochet and Tirole (2006), where we explicitly model the trade of goods under symmetrically differentiated competition; and we allow the platform to set a general set of fees. Explicitly modeling the trade of goods helps us clarify how platform incentives relate to supply and demand primitives; and allows us to discipline the model using traditional tools from empirical Industrial Organization.

Within the platform literature, this paper is closest to a growing literature focused on the business models, behavior and regulation of trade platforms. Hagiu (2009) and Anderson and Bedre-Defolie (2017) study optimal pricing and variety provisioning by marketplace platforms. Hagiu and Wright (2015); Hagiu et al. (2020); Etro (2020) study the trade-offs and determinants of marketplace vis a vis reseller intermediation, with a focus on platform cost and selling advantages. Zhu and Liu (2018) investigate this question empirically, showing that transition into Prime decreases sales ranks. Jiang et al. (2011), Etro (2020) and Anderson and Bedre-Defolie (2020) analyze platforms that face a trade-off between extracting rents and motivating innovation by third-party complementors. Last, a few papers study intermediary steering incentives (Cornière and Taylor, 2019; Chen and Tsai, 2021; Cure et al., 2021; Johnson et al., 2020), sometimes in the context of data sharing (Kirpalani and Philippon, 2021); as well as the profit and welfare implications of platform price-parity agreements (Gomes and Mantovani, 2020; Liu et al., 2021).

Different features of the platform relate to four other literatures. (a) The marketplace model is based on insights from the incidence and optimal taxation literatures (see, for example, Marshall, 1890; Anderson et al., 2001; Hamilton, 2009; Weyl and Fabinger, 2013; Peitz and Reisinger, 2014; Adachi and Fabinger, 2020), though we focus on a profit-maximizing platform as opposed to a welfare-maximizing government. (b) The reseller model thinks of Amazon as a multi-product intermediary, as in Forbes (1988); Mulhern and Leone (1991); Rhodes (2014); Rhodes et al. (2021). Like Rhodes et al. (2021), it features incentives for the platform to lower prices on products with higher pass-through and raise prices on products with lower pass-through. (c) The paper can also be interpreted as contrasting the welfare consequences of wholesaler vs.

marketplace relationships in the tradition of Villas-Boas (2007). (d) Last, since the platform uses alternate fee structures and distribution methods to optimally price discriminate, the paper relates to a long literature on price discrimination (see, for example, Varian, 1985; Aguirre et al., 2010).

From a methodological perspective, this paper contributes to literatures on supply and demand estimation. Following the seminal work of Berry (1994), the demand estimation literature has mostly focused on random coefficients models estimated as in Berry et al. (1995); Nevo (2001). We take a different approach due to data limitations, and instead estimate a nested logit model. Rather than making arbitrary assumptions for nests, however, we exploit the Grouped Fixed Effects estimator of Bonhomme et al. (2019) and Almagro and Manresa (2021) to recover the nest structure from the data. In addition, we deal with the zeroes of demand problem by following Gandhi et al. (2020). Dubé et al. (2020) and Li (2019) present alternate approaches for dealing with this problem.[10]

On the supply side, our paper relates to a long literature that empirically analyzes the price and welfare implications of vertical relationships (Villas-Boas, 2007; Bonnet and Dubois, 2010; Crawford et al., 2018; Conlon and Mortimer, 2021), often in the context of private label products (e.g., Chintagunta et al. (2002); Ellickson et al. (2018)). The closest paper is perhaps Crawford et al. (2018) which emphasizes the trade-off between potential efficiency gains from vertical integration against potential welfare losses from foreclosure incentives.

At a high level, this paper also speaks to the rise and welfare consequences of e-commerce. The initial literature focused on contrasting online and offline commerce and quantifying the gains from e-commerce. Hortacsu and Syverson (2015) provide a good introduction to aggregate trends. Brynjolfsson and Smith (2000) and Cavallo (2017) compare online and offline prices. Brynjolfsson and Smith (2001) study the importance of online retailer brands in the presence of price comparison sites. Brynjolfsson et al. (2011) discuss how the aggregation of consumers into a single national platform, combined with powerful search tools and recommendation systems gave rise to the "long tail" – the impressive variety of online offerings. Brynjolfsson et al. (2003) estimate gains from product variety in the books category. Quan and Williams (2018) emphasize the importance of across-market demand heterogeneity and the "zeroes of demand" problem for estimates of gains from variety. Accounting for these factors reduces gains from variety in the shoes category by 45%. Dolfen et al. (2019) revisit these estimates using credit card data, while also emphasizing gains from convenience. More recent papers emphasize the welfare costs of winner-take-all dynamics (Khan, 2016, 2018, 2019).

3 Overview of Amazon's Retail Platform

We begin by providing a brief overview of the Amazon platform, with a primary focus on distribution channels and fees that guide model development.

[10] Dubé et al. (2020) introduce a pairwise-difference approach that constructs moment conditions based on differences in demand between pairs of products. Li (2019) uses a parametric empirical Bayes estimator, which uses information in other markets to generate strictly positive posterior estimates of the purchase probabilities.

Figure 1: *Overview of Amazon's Retail Platform*

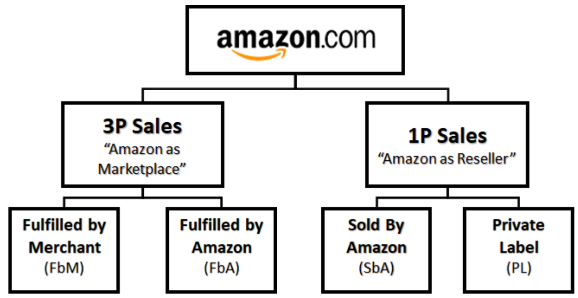

		FbM	FbA	SbA	PL
AMZ activities	Design + Mfg				✓
	Marketing/price-setting			✓	✓
	Inventory			✓	✓
	Fulfilment		✓	✓	✓
	Matching + processing	✓	✓	✓	✓
Prime-eligible?		No	Yes	Yes	Yes
Sales share (2020)		14%	45%	36%	5%

Distribution Methods. Figure 1 provides an overview of Amazon's distribution channels.[11] Consumers purchase items through the platform, either from independent third-party sellers (3P Sales) or directly from Amazon (1P Sales). Since the regulatory concerns differ across narrower distribution methods, we further divide 3P sales into sales fulfilled by merchants (FbM) and fulfilled by Amazon (FbA); and 1P sales into externally branded products sold by Amazon (SbA) and private label products (PL).

Under FbM, Amazon functions as a pure marketplace, like Ebay under fixed price offerings. In exchange for matching buyers and sellers and processing the transaction, Amazon charges an ad-valorem "referral fee" ranging from 6% to 45% of the selling price (including shipping), depending on the product category. The seller sets prices and controls the inventory and fulfillment processes – meaning they stock products on their own warehouses and ship them directly to the buyer upon sale.[12] Since sellers ship the items, FbM products are *not* eligible for 2-day Prime or Free Super Saver shipping (for non-prime customers). Instead, sellers may choose to charge a shipping fee or provide free shipping. They also provide the customer service and set their own refund and return policy.[13] Seller behavior is governed by the Amazon Seller Agreement, which sets limits on pricing, and outlines minimum levels of customer service and return policies. The Seller Agreement also prohibits the sale of counterfeits and requires sellers to abide by Minimum Advertised Prices (MAP) set by manufactures, though complaints of violations abound (e.g., Stone (2013)).

Under FbA, sellers contract Amazon for fulfillment. They send their inventory in bulk to Amazon fulfillment centers, where they are stored until sale. Upon sale, Amazon handles the shipping, customer service, refunds, and returns for the products – following Amazon's own processes and policies. Thus, FbA is essentially equivalent to SbA from the *buyers* perspective: FbA products are eligible for 2-day Prime shipping as well as Free Super Saver Shipping and, since Amazon handles the returns, returns are as easy as

[11] We exclude other businesses such as Whole Foods, Prime Video and Amazon Web Services, since they are outside the scope of this paper.

[12] Sellers may enlist a third-party to fulfill orders on their behalf, though Amazon has very specific policies around this practice.

[13] According to Amazon, most sellers on the platform offer a returns policy equivalent to Amazon's, which offers free returns for items within 30 days of receipt of shipment (link, accessed on June 4, 2020).

Table 1: *Fees for a Typical 3rd Party Transaction*

		3P sales		1P sales	
		FbM	**FbA**	**SbA**	**PL**
Consumer Price		**100**	**100**	**100**	**100**
Avg. Fees	Referral fee (ad valorem)	15	15		
	Fulfillment fees (unit+fixed)		20		
	Advertising fee (per-click)	5	5		
Wholesale cost				80	60
Marginal cost		**50**	**40**	**60**	**60**
3P gross margin		**30**	**20**	**20**	**0**
AMZ gross margin		**20**	**40**	**30**	**40**

Notes: Average fees (in red) represent our best estimates. Referral fees based on Amazon's publicly available referral fee structure. Fulfilment fees based on data from Keepa. Advertising fees based on Amazon investor disclosures. Remaining values purely representative.

for SbA products. Amazon even informs sellers that "to ensure a great customer experience, we may accept returns beyond the time-frame stated in these policies" (link, accessed on June 4, 2020). In addition to the referral fee for third party sales, Amazon collects fulfillment fees.

Under SbA, Amazon functions as an online retailer: it purchases products at wholesale prices from producers and sells them directly to consumers. Amazon collects the full selling price, so that profits are driven by a mark-up over marginal cost. Amazon takes charge of the pricing, shipping, customer service, refunds, and returns for those products – following it's own policies. SbA products are eligible for 2-day Prime and Free Super Saver Shipping.

Last, under PL, Amazon not only sells the products but it also designs and markets them under private label or exclusive brands. Most PL sales are in consumer electronics – such as the Kindle, the Fire TV and Alexa-enabled speakers – but Amazon also offers thousands of products under hundreds of brands such as AmazonBasics, AmazonCollection, Pinzon and Mama Bear.

Fee and Income Structure by Distribution Method. Table 1 summarizes the average fee and income structure by distribution channel. Amazon monetizes 1P sales based on a mark-up over wholesale costs. It monetizes 3P transactions by charging three types of fees: referral fees for matching and processing, fulfillment fees for FbA and advertising fees. Matching and processing fees apply to all 3P transactions and are set at 15% of the selling price (including shipping), on average. Fulfillment fees apply only to FbA products. They involve a monthly inventory storage fee (which varies throughout the year) plus a per-unit fulfillment fee based on the dimensions and weight of the item. On average, we estimate that fulfillment fees account for ∼20% of the selling price. Last, advertising fees apply on a product-by-product basis depending on seller decisions. They are estimated to account for ∼5% of the selling price on 3P sales.

Sales Ranks by Distribution Methods. Table 2 shows the distribution of selling methods by sales rank. SbA is primarily used for high-selling (i.e., low ranked) products. FbA is used for mid- and high-selling items. FbM is more prevalent in the long tail.

Table 2: *Cross-Sectional Distribution of Selling Methods*

	% of ASINs		
Sales rank	**SbA**	**FbA**	**FbM**
<99	71%	27%	3%
100-999	60%	39%	4%
1k-10k	48%	47%	7%
10k-50k	35%	52%	14%
>50k	19%	19%	49%

Source: estimated by author based on Keepa data

Figure 2: *Share of Sales by Selling Method*

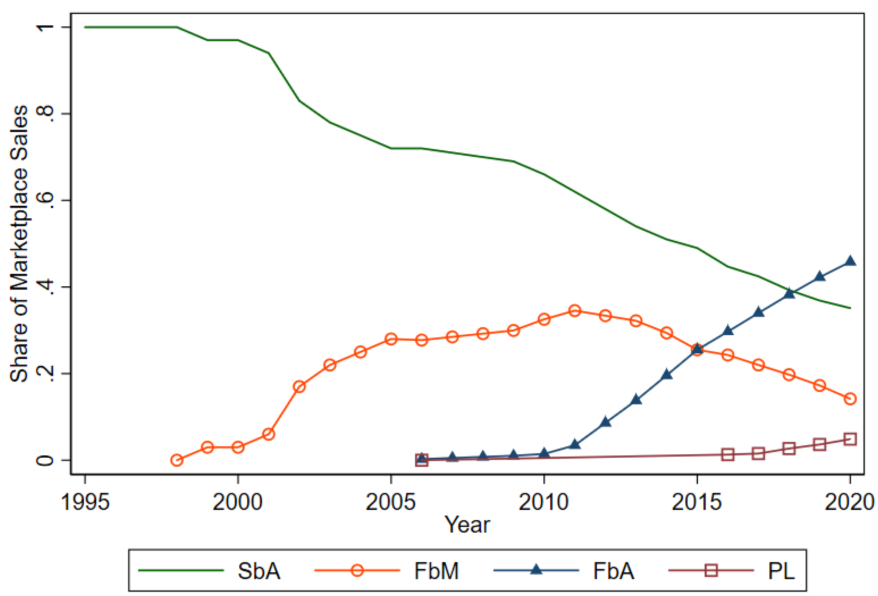

Source: compiled from AMZ disclosures (Sba, Fba, 3P) and Suntrust analyst reports (PL)

Time-series Evolution of Distribution Channels. Figure 2 shows the time-series evolution of distribution channels, as a share of sales. Amazon initially acted as a retailer, selling all products (books) directly. 100% of sales were SbA. In 2000, Amazon opened it's platform to third party sellers with the introduction of the Amazon Marketplace. These sellers became a critical part of the Amazon ecosystem and now account for 58% of sales. 3P sales were initially all FbM, but following it's introduction in 2006, the FbA program has grown rapidly and now accounts for 45% of total sales (75% of 3P sales). Private Label products were introduced in 2007 with the Kindle, and now account for 4% of sales (approximately 10% of SbA sales).

Time-series Evolution of Fees. Figure 3 shows the evolution of 3P fees over time, by type of fee. 3P fees nearly doubled since 2009, from 17% to 32%. Referral fees remained largely stable, so that the increase is almost entirely explained by a rise in Fulfillment and Advertising fees. Total fulfillment fees increased from less than 1% in 2009 to nearly 12% in 2020. Most of the increase is driven by a composition effect that accounts for the growth of the FbA program while holding FbA fees fixed at their initial level. However, a

Figure 3: *3P Fees as a % of 3P Sales, overall and by distribution channel*

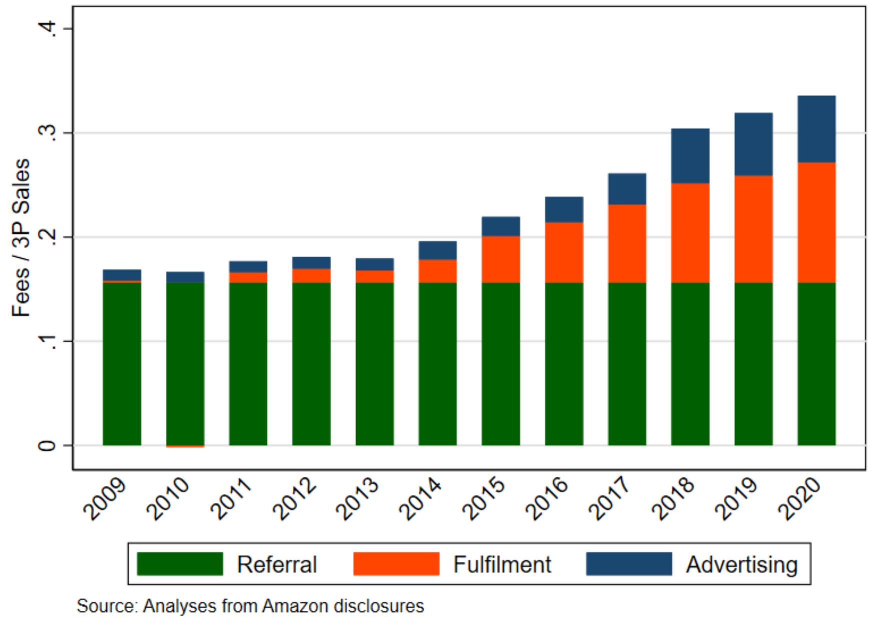

Source: compiled by the author from Amazon disclosures.

significant portion of the increase is due to a near doubling of FbA fees since 2009. Last, advertising fees increased from less than 1% of 3P sales to nearly 5%.

Competition between Sellers. Multiple merchants may offer the same item on Amazon. When there are multiple offers, Amazon uses an algorithm to rank them based on the price, seller ratings, and fulfillment methods to determine who wins the "BuyBox" – the white box on the right side of the Amazon product detail page, where customers can add items to their cart.[14] More than 80% of Amazon sales go through the BuyBox, so winning it is imperative for sellers to capture sales. Prime-eligible products have a higher chance of winning the BuyBox, *ceteris paribus*, which has be the subject of substantial criticism (e.g., Mitchell and Sussman, 2019). Essentially, this implies that Amazon preferences products that are Prime eligible and therefore steers sellers to FbA.

4 Model

We now develop a model of Amazon that incorporates the salient features described above.

4.1 Demand

Define a market as a choice set, so that consumers $i = 1, \ldots, I$ facing a choice between products $j = 1, \ldots, J$ at time $t = 1, \ldots, T$ are in the same market. We will consider national markets that include all

[14] Under rare circumstances, Amazon may suppress the BuyBox.

products sold within an Amazon subcategory k over a particular week t – where subcategories follow the Amazon product tree. There are more than 44,000 subcategories in the US including, for example, 'irons' and 'baby wipes and refills'.

Assume the utility obtained by consumer i from purchasing product j in subcategory k and market t is given by[15]

$$u_{ijt}^k = V_{jt}^k + \varepsilon_{ijt},$$

where

$$V_{jt}^k = \alpha^k p_{jt} + \zeta^k 1\{Prime_{jt}\} + x_{jt}\beta^k + \xi_{jt}.$$

The subscript t highlights that choices are market-dependent, while the superscript k highlights that parameters vary across subcategories. We omit both of these in the rest of the paper to simplify notation.

α^k measures the consumer price sensitivity in subcategory k. ζ^k measures consumer preferences for Prime fulfillment – interpreted as a bundle of services including 2-day shipping and Amazon's return policy, for example. Since several regulatory interventions involve banning Prime, this parameter will feature prominently in counter-factuals.[16] x_{jt} is a $1 \times K$ vector of additional product and seller characteristics for product j in market t including, for example, product ratings and reviews.

$\xi_{jt} = \xi_j + \xi_t + \Delta\xi_{jt}$ captures the unobservable product quality which, as pointed out by Berry (1994), is typically correlated with prices p_{jt} and, in our case, with Prime distribution. ξ_{jt} implies that there is something about a pair of *Nike* shoes that consumers prefer over *Sketchers* beyond what is captured by the observable x_{jt} characteristics (color, size, rating, etc.) that affect prices and choice of distribution methods. We include year-month fixed effects to control for trends and seasonality in the demand; and product fixed effects to control for the persistent component of unobserved quality (Nevo, 2001). Since only $\Delta\xi_{jt}$ remains, the interpretation of the unobservable quality term changes to represent the month-specific deviation from the average product quality. Estimation requires instruments z_{jt} that are correlated with prices and Prime eligibility but uncorrelated with unobserved quality so that $E[\Delta\xi_{jt}|z_{jt}]= 0$.

We instrument for prices using a manually collected history of Amazon referral and fulfillment fees. Fees are a valid instrument for prices because they affect marginal costs, but are set without knowledge of individual product shocks. While there is a long history of using fees (or, really, government taxes) as instruments (e.g., Conlon and Rao (2015)), Amazon fees have the unique feature of being widely scalable across product categories.

We instrument for Prime eligibility using seller fixed effects, as a proxy for heterogeneous preferences for distribution channels across sellers due to, for example, differences in inventory costs or shipping times. For example, Chinese sellers may prefer fulfillment by Amazon given the long shipping times from China. Seller fixed effects are a valid instrument as long as seller identity affects demand only by affecting Prime, after controlling for product fixed effects, seller ratings and reviews.

[15] This utility function can be easily derived from a quasilinear utility function in which consumers buy one (or zero) of the available products.

[16] This parameter is reminiscent of the dominant firm advantages parameter in Cabral (2018). We assume that the estimated preferences for Prime fully translate into utility gains. This may be an over-estimation if Amazon's algorithms steer consumers towards Prime-eligible products (e.g., through the BuyBox). We hope to explore this in future versions.

Last, $\varepsilon_{ijt} = \eta^k_{g(j)t} + \tilde{\varepsilon}_{ijt}$ is an idiosyncratic taste shock assumed to vary unobservably over individuals i. $g(j)$ denotes the nest to which product j belongs so that the idiosyncratic taste shock is correlated among products in the same nest. This induces a correlation structure in the covariance matrix across products, so that products in the same nest exhibit higher cross-elasticities than products across nests.

Assuming that ε_{ijt} has the appropriate distribution so that the resulting model is the nested logit, we obtain closed-form solution probabilities

$$P_j = \frac{e^{V_j/\lambda_g}\left(\sum_{k\in J_g} e^{V_k/\lambda_g}\right)^{\lambda_g-1}}{1+\sum_{h=1}^{G}\left(\sum_{k\in J_h} e^{V_k/\lambda_h}\right)^{\lambda_h}}, \tag{1}$$
$$= \underbrace{\frac{e^{V_j/\lambda_g}}{\left(\sum_{k\in J_g} e^{V_k/\lambda_g}\right)}}_{P_{ij|g}} \cdot \underbrace{\frac{\left(\sum_{k\in J_g} e^{V_k/\lambda_g}\right)^{\lambda_g}}{1+\sum_{h=1}^{G}\left(\sum_{k\in J_h} e^{V_k/\lambda_h}\right)^{\lambda_h}}}_{P_{ig}},$$

where we assume the outside good belongs to it's own nest and has utility of 0. This shows that the Nested Logit model can be thought of as a sequential choice: consumers first choose a nest, P_{ig}, and then a product within the nest, $P_{ij|g}$. The key parameter (roughly) governing within group correlation is λ_g. When $\lambda_g = 1$, the model simplifies to the multinomial logit and cross-elasticities across all products exhibit Independence of Irrelevant Alternatives. As $\lambda \to 0$, consumers always stay within their nest.

The expected surplus from choosing nest g is given by the logit inclusive value of the nest, adjusted for λ_g:

$$A_g = \lambda_g \log \sum_{k\in J_g} e^{V_k/\lambda_g}.$$

The aggregate consumer surplus, CS, is then obtained by choosing across nests:

$$CS = \log \sum_g e^{A_g}$$

Inverting the system, we obtain the traditional estimating equation:

$$\ln P_{jt} - \ln P_{0t} = x_{jt}\beta - \sigma \ln\left(s_{j|gt}\right) + \xi_{jt}, \tag{2}$$

which highlights the need for a last set of instruments for within-nest shares. We follow standard practice and use the number of products in the nest.[17]

4.2 Supply

Let us now describe the supply model.

[17]This is a valid instrument because the number of products is set prior to the realization of individual product shocks.

4.2.1 Platform Problem

We assume that Amazon maximizes

$$\Pi^A = \sum_{j \in 1P} \underbrace{(p_j - \hat{w}_j)\, s_j(\boldsymbol{p})}_{\text{1P profits}} + \sum_{l \in 3P} \left(\underbrace{u_l s_l(\boldsymbol{p})}_{\text{Unit fees}} + \underbrace{v_l p_l s_l(\boldsymbol{p})}_{\text{Ad-valorem fees}} \right) + \underbrace{\gamma CS(\boldsymbol{p})}_{\text{Consumer surplus}}, \tag{3}$$

where $s_j(\boldsymbol{p})$ denotes the share (in quantities) of product j given a vector of prices $\boldsymbol{p}$, which depends on the distribution method of product j according to consumer preferences for Prime, ζ. The first term captures profits from 1P products, which depend on a mark-up over effective marginal cost given by the sum of wholesale and Amazon fulfillment costs, $\hat{w}_j = w_j + f_j^A$. The next two terms capture revenues from unit and ad-valorem fees, $\boldsymbol{u}$ and $\boldsymbol{v}$, which are specified below. The last term puts weight $\gamma \in (0, \infty)$ on consumer surplus, to account for dynamic investment incentives of the platform (including, for example, buyer entry costs). When γ is high, platform and consumer interests are aligned and the platform sets lower mark-ups and lower fees. As γ falls, the platform begins to harvest it's customer base through higher prices and fees.

Amazon's objective function has important implications for equilibrium prices and fees. For a general price- or fee-setting instrument, f_m, the platforms first-order-condition can be written as

$$\frac{\partial AS(\boldsymbol{p})}{\partial f_m} + \gamma \frac{\partial CS(\boldsymbol{p})}{\partial f_m} = 0.$$

Dividing by $\frac{\partial AS(\boldsymbol{p})}{\partial f_m}$ and defining the Marginal Cost of Consumer Loss (MCCL) as

$$MCCL_m = -\frac{\frac{\partial CS(\boldsymbol{p})}{\partial f_m}}{\frac{\partial AS(\boldsymbol{p})}{\partial f_m}},$$

we obtain

$$MCCL_m = {}^1\!/\!_\gamma, \tag{4}$$

which shows that the platform equates the marginal consumer welfare losses across all instruments. When γ is high, the platform stomachs only small consumer welfare losses and the equilibrium $MCCL_m$ is low (for all instruments). As γ falls, however, the platform accepts higher consumer welfare losses, and raises fees across all instruments. Increasing γ then, serves to discipline the platform across all fee instruments; while decreasing γ results in pervasive increases in fees across all instruments. This may explain why a wide range of complaints against Amazon were raised essentially simultaneously.

Equation (4) also implies that the platform prefers instruments with lower $MCCL_m$. Such instruments either reallocate "effective" fees from high pass-through products towards low pass-through products (Aguirre et al., 2010), or have a lower pass-through and therefore lower incidence on consumers (Adachi and Fabinger, 2020). At the extreme, instruments with negative $MCCL_m$ allow the platform to create consumer surplus while also increasing revenues.

Importantly, instruments need not be "explicit" fees. Vertical integration, for example, may lead to a decrease in equilibrium prices in four ways: first, reselling (may) mitigate double marginalization. Specifically, when Amazon replaces an independent third-party seller by purchasing directly from the manufacturer, it turns a triple-marginalization problem (Amazon, seller, producer) into a double-marginalization problem (Amazon, producer). This improves consumer, producer and platform welfare at the expense of seller welfare. Second, reselling transfers pricing power to Amazon. This allows Amazon to optimally set prices for the individual product, instead of relying on shared fee policies for monetization. It can be viewed as an example of vertical integration for price discrimination (Varian, 1985). Third, reselling changes the order of marginalization. Under marketplace intermediation, the platform sets fees first, depending on *seller* pass-through. Under reselling, the wholesaler sets mark-ups first, depending on *platform* pass-through. To the extent that platform and seller pass-through differ from one and from each other, transitions to and from reselling re-allocate welfare across participants. Last, reselling transfers responsibility of the selling and fulfillment process to Amazon. If Amazon has selling or cost advantages vis a vis third parties, this increases welfare (and vice versa).

When γ is high, then, the platform can be expected to use vertical integration to increase consumer welfare. As γ falls, however, incentives to foreclose rivals and raise their costs become stronger, and vertical integration may be used to increase prices.

4.2.2 Price-setting

Let us now derive the optimal price-setting first-order-conditions.

Third-Party Sales. Third-party sellers, denoted by s, take the fees as given and set prices in order to maximize profits

$$\begin{aligned}\Pi^{3P} &= \sum_{j\in\mathcal{J}_s}\Big((1-v_j)p_j - (c_j+u_j^l)\Big)\, s_j(\boldsymbol{p}),\\ &= \sum_{j\in\mathcal{J}_s}(1-v_j)\,(p_j-\hat{c}_j)\, s_j(\boldsymbol{p}),\end{aligned}$$

where $\hat{c}_j = \frac{c_j+u_j}{1+v_j}$ denotes the "effective" marginal cost accounting for unit and ad-valorem fees; and $\mathcal{J}_s$ denotes the products sold by seller s. The first order condition with respect to p_j is

$$(1-v_j)\, s_j(\boldsymbol{p}) + \sum_{k\in\mathcal{J}_s}(1-v_k)\,(p_k-\hat{c}_k)\,\frac{\partial s_k}{\partial p_j} = 0,$$

which shows that the seller trades-off increased revenues from higher prices (first term) against decreased sales, accounting for the possibility of consumers switching to other goods offered by the seller (second term). Stacking these conditions into a matrix, we obtain the optimal 3P pricing policy

$$(1-\boldsymbol{v})\odot(\boldsymbol{p}-\hat{\boldsymbol{c}}) = (\boldsymbol{A}\odot\boldsymbol{\Omega})^{-1}\left((1-\boldsymbol{v})\odot s(\boldsymbol{p})\right), \tag{5}$$

where $\odot$ is the element-wise or Hadamard product of two matrices, Ω is a matrix of share-price derivatives

$$\Omega_{(j,k)} = -\frac{\partial s_k}{\partial p_j},$$

and A denotes product ownership: $A_{(j,k)} = 1$ if (j,k) have the same seller and 0 otherwise.

For single-product firms, the above reduces to the standard pricing condition:

$$p_j - \hat{c}_j = -\frac{s_j}{\frac{\partial s_j}{\partial p_j}} \Rightarrow \frac{p_j - \hat{c}_j}{p_j} = \frac{1}{\varepsilon_j},$$

where $\varepsilon_j = \frac{p_j}{s_j}\frac{\partial s_j}{\partial p_j}$. Relative to this, $\boldsymbol{A}$ accounts for cross-ownership of multi-product sellers and $(\boldsymbol{1} - \boldsymbol{r})$ accounts for seller incentives to steer buyers towards products with lower fees.

First-Party Sales. Amazon takes wholesale prices w_{jt} as given and sets prices. The first order condition of equation (3) with respect to $p_j : j \in SbA$ is then given by:

$$s_j(\boldsymbol{p}) + \sum_{j \in SBA} (p_k - \hat{w}_k) \frac{\partial s_k}{\partial p_j} + \sum_{l \in 3P} \left(u_l \frac{\partial s_m}{\partial p_j} + v_l p_l \frac{\partial s_l}{\partial p_j} \right) + \gamma \frac{\partial CS(\boldsymbol{p})}{\partial p_j} = 0,$$

The first two terms are standard for a multi-product retailer. The platform trades-off increased revenues from higher prices against decreased revenues from lower quantities, accounting for cross-elasticities across products sold by Amazon. The next two terms are new, however. The first one accounts for unit and ad-valorem fees collected on 3P products, to the extent that consumers switch to purchase them. This leads to higher mark-ups vis a vis a pure reseller. The last term measures the consumer welfare loss from raising prices, and is weighted by the platform's investment incentive parameter, γ. It implies that the platform will optimally set lower prices when (a) γ is high or (b) products contribute more to consumer surplus.

Stacking these conditions into a matrix, we obtain the optimal pricing policy

$$\boldsymbol{p} - \hat{\boldsymbol{w}} = (\boldsymbol{A} \odot \boldsymbol{\Omega})^{-1} \left(s(\boldsymbol{p}) - (\boldsymbol{B} \odot \boldsymbol{\Omega})(\boldsymbol{u} + \boldsymbol{r} \odot \boldsymbol{p}) + \gamma \frac{\partial CS(\boldsymbol{p})}{\partial \boldsymbol{p}} \right), \tag{6}$$

where $\boldsymbol{\Omega}$ and $\boldsymbol{A}$ are defined as before; and $\boldsymbol{B}$ identifies third-party products collecting fees

$$\boldsymbol{B} = \begin{cases} 1 & \text{for}(j,k) : j \in \mathcal{J}_{SBA}, k \in \mathcal{J}_{3P} \\ 0 & \text{otherwise} \end{cases}.$$

To gain some intuition, consider the case of symmetric products distributed directly by Amazon, and assume that $\gamma = 0$. With a slight abuse of notation, let p, $s(p)$ and ρ denote the symmetric price, quantity and pass-through, respectively. In that case, the pricing condition simplifies to

$$p - \hat{w} = -\frac{s(p)}{\rho \sum_j \frac{\partial q_i}{\partial p_j}}.$$

Denoting the the outside good diversion – which measures the share of consumers that leave the platform when the price of a product increases – as θ:

$$\theta = \frac{\frac{\partial q_i}{\partial p_i} - \sum_{j \neq i} \frac{\partial q_j}{\partial p_i}}{\frac{\partial q_i}{\partial p_i}},$$

we obtain:

$$\frac{p - \hat{w}}{p} = \frac{1}{\theta \varepsilon_j}.$$

This condition shows that optimal mark-ups decrease with the aggregate elasticity, $\theta\varepsilon_j$, instead of the firm elasticity, ε_j. Intuitively, when the retailer offers all the products, it cares about the rate at which consumers stop purchasing *any* product – instead of the rate at which they stop purchasing *a particular* product. It is easy to show that, in this case, mark-ups increase with product variety and substitutability.

Wholesaler price-setting. Wholesalers take Amazon price-setting as given, and set a price w to maximize

$$\Pi_t^m = \sum_{j \in \mathcal{J}_m} (w_j - c_j)\, s_j(\boldsymbol{p}),$$

where $\mathcal{J}_m$ denotes the products sold by wholesaler m. The first order condition with respect to w_j is

$$0 = s_j(\boldsymbol{p}) + \sum_{k \in \mathcal{J}_M} (w_k - c_k) \frac{\partial s_k}{\partial w_j},$$

where

$$\frac{\partial s_k}{\partial w_j} = \sum_l \frac{\partial s_k}{\partial p_l} \frac{\partial p_l}{\partial w_j}.$$

Stacking these terms into a matrix, we can write the change in quantities relative to an increase in wholesale prices as

$$\frac{\partial \boldsymbol{s}}{\partial \boldsymbol{w}} = \boldsymbol{\Omega} \boldsymbol{\rho}_{\boldsymbol{w}},$$

where Ω is as defined above and $\boldsymbol{\rho}_{\boldsymbol{w}}$ denotes the wholesale price pass-through matrix:

$$(\boldsymbol{\rho}_{\boldsymbol{w}})_{jk} = \frac{\partial p_j}{\partial w_k}.$$

$\boldsymbol{\rho}_{\boldsymbol{w}}$ measures how prices of *all* products change when wholesale prices of product k increase. This is a rich object that depends on buyer and seller elasticities and cross-elasticities (Weyl and Fabinger, 2013).

Optimal wholesale mark-ups are then given by

$$\boldsymbol{w} - \boldsymbol{c} = (\boldsymbol{A} \odot \boldsymbol{\Omega} \boldsymbol{\rho}_{\boldsymbol{w}})^{-1} s(\boldsymbol{p}). \tag{7}$$

Wholesale prices are inversely proportional to (a) consumer elasticity of demand and (b) the *platform's* pass-through. Since $\boldsymbol{\rho_w}$ depends on γ, this introduces an important interaction between wholesaler mark-ups and platform investment incentives. When γ is high, the platform aims to provide a better bundle for consumers, passing through a smaller share of cost increases (i.e., $\boldsymbol{\rho_w}$ is low). Wholesalers internalize this and set higher mark-ups. As γ falls platform pass-through increases and wholesaler mark-ups fall.

4.2.3 Fee-setting

Let us now discuss our modeling of fees.

Modeling. We assume that unit and ad-valorem fees depend on a rich set of fee instruments designed and implemented by the platform. Specifically, the net unit and ad-valorem fees are given by

$$\boldsymbol{v} = \boldsymbol{\tau}^r \boldsymbol{V} \quad \text{and} \quad \boldsymbol{u} = \boldsymbol{\tau}^u \boldsymbol{U}$$

where $\boldsymbol{V} = [\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{N^v}]$ and $\boldsymbol{U} = [\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{N^u}]$ are $J \times N^v$ and $J \times N^u$ matrices of ad-valorem and unit fee instruments, respectively; and $\boldsymbol{\tau}^r$, $\boldsymbol{\tau}^u$ are $1 \times N^v$ and $1 \times N^u$ vectors of loadings on the corresponding instruments. For example, when only a common ad-valorem fee of 15% is applied to all products, $\boldsymbol{V}$ is a vector of ones and $\tau^r = 0.15$. When fully heterogeneous fees are allowed, $\boldsymbol{V}$ and $\boldsymbol{U}$ are identity matrices.[18]

In theory, the platform would choose to set fully heterogeneous fees in order to extract maximum surplus. In practice, however, platforms are often restricted in the set of fee instruments they use. Even today, many platforms use a single ad-valorem fee for all products and categories (e.g., Etsy), perhaps due to operational (e.g., it may be hard to set or enforce heterogeneous fees) or strategic concerns (e.g., fees may have anchored seller and regulatory expectations, so that changing them may too costly).

Flexibility in fee structures has important implications on the welfare consequences of regulation because it governs the platform's ability to (a) price discriminate on third-party products, (b) load on low pass-through instruments and (c) endogenously respond to regulatory interventions. To closely mirror the platform – in both estimation and counter-factuals – we take the set of fee instruments $\boldsymbol{V}$, $\boldsymbol{U}$ from the data and only allow the platform to adjust the optimal weighting vectors $\boldsymbol{\tau}^r$ and $\boldsymbol{\tau}^u$ in response to regulation.

Unit fees. The first order condition of equation 3 with respect to a particular unit fee instrument u_m is:

$$\sum_l u_l s_l(p) + \sum_{k \in 1P} (p_{kt} - \hat{w}_{kt}) \frac{\partial s_k(\boldsymbol{p})}{\partial u_m} + \sum_{l \in 3P} \left(\tau^u_m u_{m,l} \frac{\partial s_l(\boldsymbol{p})}{\partial u_m} + v_l p_l \frac{\partial s_l(\boldsymbol{p})}{\partial u_m} + v_l s_l \frac{\partial p_l}{\partial u_m} \right) + \gamma \frac{\partial CS(\boldsymbol{p})}{\partial u_m} = 0$$

[18]We focus on explicit unit and ad valorem fees in this paper. However, as described by Weyl and Fabinger (2013) and Adachi and Fabinger (2020), many interventions ranging from exogenous competition to selling advantages enter third-party sellers' problem as a mixture of unit and ad-valorem taxation. They could, therefore, be easily incorporated into our framework.

where

$$\frac{\partial s_k}{\partial u_m} = \sum_l \frac{\partial s_k(\boldsymbol{p})}{\partial p_l}\frac{\partial p_l}{\partial u_m}.$$
$$\frac{\partial CS(\boldsymbol{p})}{\partial u_m} = \sum_l \frac{\partial CS(\boldsymbol{p})}{\partial p_l}\frac{\partial p_l}{\partial u_m}.$$

Stacking the above conditions across fee instruments and using matrix notation, optimal Amazon unit fee loadings are

$$\boldsymbol{\tau}^u = (\boldsymbol{\Omega_f U})^{-1}\left(\boldsymbol{U'}s(\boldsymbol{p}) - (\boldsymbol{C}\odot\boldsymbol{\Omega_f})(\boldsymbol{p}-\hat{\boldsymbol{w}}) - \boldsymbol{\Omega_f}(\boldsymbol{r}\odot\boldsymbol{p}) + \boldsymbol{\rho_f}(\boldsymbol{r}\odot s(\boldsymbol{p})) + \gamma\boldsymbol{\rho_f}\frac{\partial CS(\boldsymbol{p})}{\partial \boldsymbol{p}}\right), \quad (8)$$

where $\boldsymbol{C}$ is a $N^u \times J$ matrix that identifies 1P products ($C_{jk} = 1$ if $k \in \mathcal{J}_{SBA}$); and

$$\boldsymbol{\Omega_f} = \boldsymbol{\rho_f\Omega}$$

is a $N^u \times J$ matrix of demand derivatives with respect to changes in fees which depends on price-elasticity of demand, $\boldsymbol{\Omega}$, and a $N^u \times J$ unit fee pass-through matrix $\boldsymbol{\rho_w}$(defined analogously to $\boldsymbol{\rho_w}$). This condition is closely related to equation (6) above, except that (a) the impact of fees on consumer prices is now moderated by $\boldsymbol{\rho_f}$ and (b) fees now apply across multiple products – as governed by the matrix $\boldsymbol{U}$ – and therefore depend on the weighted average, rather than product-specific own and cross-elasticities.[19]

Referral fees. Finally, through a similar process, we can recover Amazon's referral fee first-order condition as:

$$\boldsymbol{\tau}^v = ([\boldsymbol{\Omega_r}\odot\boldsymbol{p} - \boldsymbol{\rho_r}\odot s(\boldsymbol{p})]\,\boldsymbol{V})^{-1}\left(\boldsymbol{V'}(s(\boldsymbol{p})\odot\boldsymbol{p}) - (\boldsymbol{C}\odot\boldsymbol{\Omega_r})(\boldsymbol{p}-\hat{\boldsymbol{w}}) - \boldsymbol{\Omega_r u} + \gamma\boldsymbol{\rho_r}\frac{\partial CS(\boldsymbol{p})}{\partial \boldsymbol{p}}\right), \quad (9)$$

where all matrices are defined as above. Again, pricing depends crucially on pass-through and cross-elasticities across products.

5 Estimation

We estimate the model in three steps.

5.1 Step 1: Assign Products to Nests

In the first step, we exploit the Grouped Fixed Effects estimator of Bonhomme et al. (2019) to recover the nesting structure $g(j)$ from the data – as suggested by Almagro and Manresa (2021).

[19] The pricing condition in the case of symmetric products with a single unit fee is $\frac{t}{p} = \frac{1}{\rho\theta\varepsilon}$,which shows that unit fees are highest when (i) fees are not passed on to consumers, (ii) consumers switch to inside goods in response to higher prices or (iii) consumers have a low price-elasticity.

We use this methodology to recover cross-elasticities, as opposed to a random coefficients model as in Berry et al. (1995), due to data limitations: we observe a single national market with no data on geographic or consumer heterogeneity; and have only a few years of data so that projecting coefficients to time-varying demographic characteristics of the US offers limited power. The grouped fixed effects approach allows us to exploit comparatively rich product-level data to recover cross-elasticities by essentially grouping together products that are "similar" according to observable metrics.[20] This methodology is particularly appealing in our setting because (a) it is easily scalable over the thousands of categories in our sample and (b) it can be extended to include external measurement of types based on unstructured data (including, for example, textual data in product titles and descriptions).

Two assumptions are needed to ensure we recover the right nest structure. First, the unobservable groups must be of low-dimension. This holds in our setting, as long as consumers make decisions consistent with the nested logit model. In that case, we can factorize unobservable shocks $\eta_{g(j)t}$ into low-dimensional vectors identifying product nests, $\zeta_{j0}^{'}$, and market-specific shocks ι_{t0} so that $\eta_{g(j)t} = \zeta_{j0}^{'}\iota_{t0}$. Second, we must identify a set of moments from which the underlying types can be approximated. We construct these moments using the panel data, as described below. In the presence of endogenous parameters, Almagro and Manresa (2021) show how to extend the approach to account for endogeneity by following a control function approach, as in Petrin and Train (2010). Depending on the number of groups G, the K-means algorithm delivers different product partitions. We choose the number of nests that maximizes the Silhouette Score (a common proxy in the Machine Learning literature).

5.2 Step 2: Estimate demand parameters, taking nests as given

In the second step, we estimate the parameters of the demand model, taking the nests as given. The estimation is complicated by the large share of low or zero market share products prevalent in e-commerce. We begin by introducing this "zeroes of demand" problem, before discussing the estimation approach of Gandhi et al. (2020).

Zeroes of Demand Problem. Equation 2 shows the estimating equation for the nested logit model. Under an asymptotic framework that rules out zeroes as the number of consumers goes to infinity, this can be estimated through a two-stage least square regression where P_{jt} and P_{0t} are replaced with their empirical counterparts s_{jt} and s_{0t}.[21] This asymptotic framework is not realistic for e-commerce, however. E-commerce sellers are less constrained by physical space than traditional sellers, and the marginal cost of offering a new product is close to zero. As a result, many products are offered even if they sell only a few times per year.

When the market share is zero, the left hand size of equation 2 is not defined. While it is common practice to simply drop these observations, this severely biases the parameters (Gandhi et al., 2020). Even if s_{jt} is non-zero but small, measurement error can lead to large bias.[22]

[20] Other settings where a similar approach may be useful include, for example, mutual find markets in Finance where the market is fundamentally national and data on investor heterogeneity is relatively limited (Koijen and Yogo, 2014).

[21] Berry et al. (2004), for example, assumes that $|s_{jt} - P_{jt}|/P_{jt} \to_p 0$ for logit class models.

[22] This is particularly a problem in our case, since sales rank are less informative for products that sell only a few items per year:

A more realistic asymptotic framework for e-commerce is to let $n_t \to \infty$, while simultaneously allowing a non-negligible fraction of products to drift to zero at the rate $1/n_t$. Intuitively, as more consumers join the platform, the long tail grows longer so that some products will always sell only a few times per year. In a recent paper, Gandhi et al. (2020) consider precisely this asymptotic framework, and develop an estimation strategy that recovers unbiased parameters when the fraction of zeroes is as high as 95%.

Estimator. Let us provide an overview of the approach. We begin by discussing the assumptions and the estimation problem, before providing intuition for how it works. We refer the reader to Gandhi et al. (2020) for complete derivations and proofs.

Assume that each product jt is either a "safe" product for which $P_{jt} \geq \varepsilon_0$ or a "risky" product for which $n_t P_t \geq \varepsilon_1$, for positive numbers ε_0 and ε_1.[23] We will use the safe products as a source of identification, while imposing slack moment inequalities for the risky products that do not contribute but also do not undermine identification. To do so, we assume that safe products are characterized by observable instruments $z_{jt} \in \mathcal{Z}_0$, where $\mathcal{Z}_0$ is a subset of the support of z_{jt}.

We then recover the model parameters, $\hat{\theta} := \left(\hat{\beta}', \hat{\lambda}'\right)'$, by solving:

$$\hat{\theta}_T := \left(\hat{\beta}'_T, \hat{\lambda}'_T\right)' = \arg\min_{\theta \in \Theta} \widehat{Q}_T(\theta), \tag{10}$$

where

$$\begin{aligned}
\widehat{Q}_T(\theta) &= \sum_{g \in \mathcal{G}} \mu(g) \left\{ [\bar{m}^u_T(\theta, g)]^2_- + [\bar{m}^l_T(\theta, g)]^2_+ \right\}, \text{ with} \\
\bar{m}^u_T(\theta, g) &:= T^{-1} \sum_{t=1}^{T} \sum_{j=1}^{J} (\delta^u_{jt}(s_t, \lambda) - x'_{jt}\beta) g(z_{jt}); \\
\bar{m}^l_T(\theta, g) &:= T^{-1} \sum_{t=1}^{T} \sum_{j=1}^{J} (\delta^l_{jt}(s_t, \lambda) - x'_{jt}\beta) g(z_{jt});
\end{aligned}$$

$\mu(g) : \mathcal{G} \to [0, 1]$ is a probability mass function on $\mathcal{G}$; $[x]_- = \min\{0, x\}$ and $[x]_+ = \max\{0, x\}$. $\mathcal{G}$ is a countable collection of instrumental indicator functions $g : R^{d_z} \to \{0, 1\}$. Last, $\delta^u_{jt}(s_t, \lambda)$ and $\delta^l_{jt}(s_t, \lambda)$ are functions that bound $\delta_{jt}(\pi_t, \lambda)$ from above and below, *on average*:

$$\begin{aligned}
\delta^u_{jt}(s_t, \lambda) &= \log\left((n_t s_{jt} + \iota_u)/n_t\right) + \delta_{jt}(\tilde{s}_t, \lambda) - \log(\tilde{s}_{jt}), \\
\delta^l_{jt}(s_t, \lambda) &= \log\left((n_t s_{jt} + \iota_l)/n_t\right) + \delta_{jt}(\tilde{s}_t, \lambda) - \log(\tilde{s}_{jt}),
\end{aligned}$$

where ι_u and ι_l are fixed numbers and $\tilde{s}_t$ is a slight modification of s_t to take it off the boundary of the probability simplex. We follow Gandhi et al. (2020) and let $\iota_l = 2^{-52}$, $\iota_u = 2$, and $\tilde{s}_{jt} = s_{jt} + 1/n_t$.

the purchase of a single item can lead to sizable declines in the sales rank, which slowly revert over long periods with no sales.

[23] Neither ε_0, ε_1, nor the identity of safe products is known – except for the outside good, which is always assumed to be safe product.

Two non-standard features of problem (10) ensure that we recover the true value of β. First, the bounds δ^u_{jt} and δ^l_{jt} are used instead of a point estimate of δ_{jt}. These bounds are specifically designed to collapse to each other and to δ_{jt} for safe products, while remaining lax for risky ones. Second, the moments enter the criterion function $\widehat{Q}_T(\theta)$ through a negative part and a positive part, which implies that the criterion function is small when evaluated at the true value. Combined, these features mean that the criterion function behaves like a standard GMM problem for safe products.

Formally, let $\mathcal{G}_0$ denote the subset of $\mathcal{G}$ containing the g's with support lying in $\mathcal{Z}_0$, i.e.,

$$\mathcal{G}_0 = \{g \in \mathcal{G} : g(z) = 0 \,\forall\, z \notin \mathcal{Z}_0\}.$$

Then, we can re-write $\widehat{Q}_T(\theta)$ as

$$\begin{aligned}\widehat{Q}_T(\theta) &= \sum_{g\in\mathcal{G}} \mu(g) \left\{ [\bar{m}^u_T(\theta, g)]^2_- + [\bar{m}^l_T(\theta, g)]^2_+ \right\} \\ &= o_p(1) + \sum_{g\in\mathcal{G}_0} \mu(g) \left(T^{-1} \sum_{t=1}^{T} \sum_{j=1}^{J} (\delta_{jt} - x'_{jt}\beta) g(z_{jt}) \right)^2,\end{aligned}$$

which behaves as a GMM criterion function where δ_{jt} is used directly and is bounded away from zero for β bounded away from β_0.

The challenge, then, is ensuring that $\mathcal{Z}_0$ is rich enough and $\mathcal{G}_0$ contains enough functions on $\mathcal{Z}_0$. Intuitively, we need to find an instrument or a combination of instruments that have the ability of indicating high demand (i.e., safe) products. Thankfully, the instruments can be control variables included in x_{jt} (i.e., they do not need to be excluded instruments). Product dummies that identify products with steadily high demand, for example, can serve to indicate safe products. This is our strategy, which ensures the non-emptiness of $\mathcal{Z}_0$.

Given non-emptiness, the richness of $\mathcal{G}_0$ is ensured by the construction of $\mathcal{G}$. Andrews and Shi (2013) show that, if $\mathcal{G}_0$ contains all indicator functions of hypercubes $B_g \subseteq \mathcal{Z}_0$, it is rich enough to preserve the information provided by the richness of $\mathcal{Z}_0$. With $\mathcal{Z}_0$ unknown, a simple way to ensure that is to let $\mathcal{G}$ contain all indicator functions of hypercubes in the support of z_{jt}. The number of hypercubes increases quickly with instruments, however, so it is not feasible to implement for thousands of products. The appendix of Andrews and Shi (2013) shows that a countable reduction of the set of all indicators of hypercubes works just as well, so we use the latter.

Formally, we divide the instrument vector z_{jt} into discrete instruments, $z_{d,jt}$, and continuous instruments $z_{c,jt}$. Without loss of generality assume $z_{c,jt}$ lies in $[0, 1]^{d_{z_c}}$. Let the set $\mathcal{Z}_d$ be the discrete set of values that $z_{d,jt}$ can take. Then, the set of instruments, $\mathcal{G}$, is defined as

$$\begin{aligned}&\mathcal{G} = \left\{ g_{a,r,\zeta}\left(z_c, z_d\right) = 1\left((z'_c, z'_d) \in C_{a,r,\zeta}\right) : C_{a,r,\zeta} \in \mathcal{C} \right\}, \text{ where} \\ &\mathcal{C} = \left\{ \left(\times_{u=1}^{d_{z_c}} (a_u - 1)/2r, a_u/2r \right) \right\} \times \{\zeta\} : a_u \in \{1, 2, \ldots, 2r\}, \text{ for } u = 1, \ldots, d_{z_c}, \\ &r = r_0, r_0 + 1, \ldots, \text{ and } \zeta \in \mathcal{Z}_d.\end{aligned}$$

Like Gandhi et al. (2020), we truncate r at a finite value $\bar{r}_T$. For $\mu(\cdot)$, we use

$$\mu(\{g_{a,r,\zeta}\}) \propto (100+r)^{-2}(2r)^{-d_{z_c}} K_d^{-1},$$

where K_d is the number of elements in $\mathcal{Z}_d$.

5.3 Step 3: Estimate supply parameters, taking demand parameters as given

In the last step, we estimate supply parameters, taking the demand structure as given. To do this, we need to first specify the fee structure. We focus on the interaction between ad-valorem referral fees and unit fulfillment fees, and therefore abstract from advertising fees.[24] In addition, we assume that all referral fees are ad valorem and all fulfillment fees are unit fees even though, in reality, ad-valorem fees include a small closing unit fee and fulfillment fees include a quasi-fixed storage cost. We then take the structure of the fees from the data. Specifically, let a vector $\boldsymbol{r} = (r_1, \ldots, r_J)$ denote the referral fees applicable for all products in a given market t. We let $\boldsymbol{v} = \tau^r \boldsymbol{r}$ so that $\tau^r = 1$ in the baseline. We then hold $\boldsymbol{r}$ fixed, and only allow the platform to change τ^r in response to regulation. We follow the same process for unit fulfillment fees.

To solve the model, we search for the set of mark-ups and platform's weight on consumer surplus that rationalizes the observed prices and referral fees (taken from the data), while jointly satisfying Amazon's first-order-conditions; as well as wholesaler and third-party price-setting FoCs. This estimation methodology (and the model, itself) assumes that observed prices and fees satisfy the platform's first-order-conditions. This is a standard assumption, and a useful benchmark that is certainly plausible for prices. It may be less plausible for referral fees, however.[25] We plan to explore alternate estimation assumptions in the future.

Ultimately, supply estimation recovers five objects: (i) Amazon's weight on consumer surplus; (ii) Amazon's mark-up on FbA services; (iii) Amazon's mark-up on 1P products; (iv) third-party mark-ups; and (v) wholesaler mark-ups. We then use the estimated mark-ups to recover marginal and wholesale costs as well as Amazon's fulfillment costs, which serve as inputs to counter-factuals.

Given that there are hundreds of products, solving the supply model is computationally difficult. We compute pass-through analytically to speed up calculations. Let $\boldsymbol{Z}$ be a matrix that collects the price-setting first-order conditions

$$Z = \begin{cases} (\boldsymbol{A} \odot \boldsymbol{\Omega})(\boldsymbol{p} - \hat{\boldsymbol{c}}) + s(\boldsymbol{p}) & \text{if 3P} \\ (\boldsymbol{A} \odot \boldsymbol{\Omega})(\boldsymbol{p} - \hat{\boldsymbol{w}}) + s(\boldsymbol{p}) - (\boldsymbol{B} \odot \boldsymbol{\Omega})(\boldsymbol{r}\boldsymbol{p} - \boldsymbol{u}) + \gamma \frac{\boldsymbol{\partial CS_t}}{\boldsymbol{\partial p}} & \text{if SbA} \end{cases}$$

[24] We abstract from advertising primarily due to data limitations, but hope to explore counter-factual analyses accounting for these fees in the future.

[25] Such fees exhibit limited heterogeneity and have remained largely fixed for many years. They may be difficult to adjust due to operational (e.g., they may be hard to set or enforce) or strategic concerns (e.g., fees may have anchored seller and regulatory expectations, so that changing them may too costly). If so, instead of raising referral fees, the platform may have chosen to introduce new fee instruments. This would lead to an over-estimation of the estimated weight on consumer surplus, but would not affect the economic mechanisms.

Using the implicit function theorem, the pass-through rate matrix can be derived as

$$\frac{\partial \boldsymbol{p}}{\partial \boldsymbol{u_m}} = -\left(\frac{\partial \boldsymbol{Z}}{\partial \boldsymbol{p}}\right)^{-1}\left(\frac{\partial \boldsymbol{Z}}{\partial \boldsymbol{u_m}}\right),$$

which depends on the specific tax instrument $\boldsymbol{u}_m$ (or $\boldsymbol{v}_m$). The challenge, then, is to compute these matrices given a demand and supply structure.

6 Data

Product-level data. In order to estimate the model, we obtain product-level data from Keepa.com (hereafter Keepa). Keepa is a private online market intelligence service for Amazon buyers and sellers that has been scraping Amazon since 2011. Buyers use Keepa to track the price of desired items, while sellers use it to identify new product markets to enter or study competing products.

Keepa can potentially track all Amazon products except for e-books. Once a product enters the database, tracking begins and continues indefinitely – even if the product is no longer sold on Amazon (in which case it appears with no associated price). Keepa's database is constantly growing as new products are added on a daily basis. At any given point, Keepa tries to track all best-sellers within each Amazon product category and, if anyone searches for a product that is not being tracked, then Keepa starts tracking it. As of November 2021, the database includes over 2.3 billion products sold on Amazon in 11 countries (USA, UK, Germany, France, Japan, Canada, Italy, China, Spain, Mexico and India).

Keepa stores two types of information for each product. First, descriptive information including product codes (e.g., ASIN, UPC and EAN), product characteristics (e.g., brand, manufacturer, size), display information (e.g., product title and description on Amazon), and a product category tree (e.g., Automotive >...>Windshield Sunshades). Second, time-series information for several measures of prices (e.g., BuyBox prices, Amazon prices, FbA and FbM prices), several proxies for quantities called sales rank as well as product ratings, number of reviews, number of sellers, and seller identifiers. Keepa updates it's data by repeatedly scraping Amazon's product pages. Some time-series fields (e.g., Amazon prices and sales ranks) are updated several times per day but other fields critical to our analyses (e.g., the BuyBox price and seller identifiers) are updated only a few times per week.[26]

Due to budget constraints, we cannot download the full dataset. Instead, we take a $\sim$5% random sample of product subcategories in the US and gather all information available in Keepa for the underlying products. We then map these subcategories across countries by searching for the corresponding products and picking the subcategories with the highest match rates. The resulting dataset is far larger than nearly all other datasets used in the literature including, for example, Cavallo (2018). [27]

[26] Keepa continuously improves its data gathering capabilities. In 2021, it began to collect BuyBox information several times a day.

[27] The closest dataset in terms of scale is perhaps Nielsen. Keepa and Nielsen, however, differ in three important ways. First, Keepa provides much broader coverage across product categories, while Nielsen focuses on grocery products. Amazon includes 44 thousand product subcategories, compared to only one thousand for Nielsen. Second, Keepa provides data across multiple countries, while Nielsen mostly focuses on the US. Last, Keepa provides data for a single national market, while Nielsen provides far more geographic and consumer heterogeneity – including store – and, for a smaller sample, consumer-level variation.

We restrict the analysis to Aug-2018 - March-2020 where coverage of distribution methods and BuyBox prices improves; and sales ranks use a common definition.[28] We then use the full sample for aggregate analyses, and a few selected product subcategories for model estimation.

Estimating sales quantities. Keepa contains only sales ranks not sales quantities. Sales ranks are ordinal rankings describing the quantity sold of a product within a given category. The item ranked as 1 is the best-seller in its category. Sales ranks are updated hourly, and range from 1 to several million – depending on the number of items in the category. Amazon has not disclosed the algorithm used to compute sales ranks but it is clear that more recent purchases are weighed more heavily than older ones.

Several marketing papers study the relationship between sales rank and sales quantities at the root-category level – primarily for books (e.g., Chevalier and Goolsbee, 2003; Chevalier and Mayzlin, 2006; Brynjolfsson et al., 2006). Chevalier and Goolsbee (2003) use a combination of actual sales and sales rank data from one publisher, as well as experiments where they purchase individual items and track how the sales rank changes to argue that rank data fits a power law. Brynjolfsson et al. (2011) revisit this relationship using data for books collected in 2000 and 2008. They find that sales of high-ranked products have increased (i.e., the long tail has grown longer over time) and argue that, "while power laws are a good first approximation for the rank-sales relationship, the slope is not constant for all book ranks, becoming progressively steeper for more obscure books." In other words, splines fit the relationship between sales ranks and sales quantities better.

Two leading market intelligence services for Amazon sellers (AMZScout and JungleScout) operationalize these methods to provide estimates of sales quantities based on sales ranks to sellers. AMZScout appears to use a Power law, while JungleScout uses splines. We download a sample of product sales ranks and estimated quantities by root category and fit a Pareto and spline relationship, respectively. We then use these relationships to predict sales quantities from sales ranks.[29]

Constructing a market. Last, we use the data to construct "markets". Specifically, we define a market as a subcategory x week to roughly match the update frequency of BuyBox prices and selling methods. We assume that all products in a subcategory belong to the same market. Whenever there are more than one offers for the same product, we assume the BuyBox-winning offer is the primary one. This is consistent with the fact that vast majority of sales on the platform occur through the BuyBox. We aggregate to the weekly level by adding daily sales and averaging end-of-day prices and other product characteristics. Last, we assume the market size is such that Amazon's market share equals its e-commerce share, as estimated by eMarketerPro. Assumptions that consider a bigger market would imply higher outside elasticities but lower weights on consumer surplus. The implications on the welfare consequences of regulation are therefore

[28] Amazon changed the "unit of analysis" for defining sales ranks in 2020, from variations to parent products which would affect the conversion to quantities.

[29] Several studies measure sales rank elasticities and convert them to sales quantity elasticities using the power law. However, it is important to note that sales ranks have only been studied at the level of root categories (e.g., books, consumer electronics). These are far too broad to be considered markets, and a power law in the aggregate need not imply a power law for all subcategories. It is therefore preferable to use aggregate relationships to convert sales ranks to sales, and then use sales estimates to measure market shares.

Table 3: *Summary statistics: Gouda Cheese*

This table shows the summary statistics across products in the "Gouda Cheese" category. The first three columns report the share of products by distribution method. The next three report the distribution of prices and fees (conditional on the fees being relevant).

	SBA	FBA	prime	prices ($)	Ref fee (%)	FbA fee ($)
count	11584	11584	11584	11584	11505	534
mean	0.01	0.05	0.05	52	0.15	3.50
std	0.08	0.21	0.22	66	0.01	1.55
min	0	0	0	5	0.07	2.41
25%	0	0	0	21	0.15	2.41
50%	0	0	0	27	0.15	3.19
75%	0	0	0	52	0.15	3.48
max	1	1	1	492	0.26	10.79

ambiguous.

Amazon fees. Last, to use as instruments in demand estimation and targets in supply estimation, we manually collect a history of Amazon fees. These data were initially gathered and implemented for a companion project, Covarrubias et al. (2022), which uses them to estimate pass-through across a large set of categories. Fees vary in complex ways within and across product categories, dimensions and weight. We replicate Amazon's policies at the product subcategory level.

7 Estimated Parameters

Let us now discuss the estimated parameters. We provide a detailed discussion of the estimation and results for a sample category and discuss estimates for a sample of 10 subcategories. In future versions, we hope to consider thousands of categories, in order to perform cross-sectional analyses.

Table 3 shows the summary statistics for our sample category, Gouda cheese. Only 1% of products are SbA and 5% of products are FbA, but combined they account for 78% of sales (see Table 5 below). The remainder is captured by FbM products. The median price for a cheese is $27. The median ad-valorem fee is 15%, and unit fulfillment fee is 3.19. Both fees vary across products, and over time. FbA fees are changed annually while referral fees on products with a price below $15 were cut from 15% to 8% in 2019.

Figure 4 shows the results of K-means clustering for the category. As shown, nests primarily depend on prices. Intuitively, this implies that consumers first decide "roughly" how much they want to spend on the product, and then consider alternatives around that price.

Table 4 shows the estimated demand parameters. The first column includes no fixed effects. This leads to a very high estimate on the consumer value of Prime. The next columns add product fixed effects. The welfare gains of Prime drop significantly, to 3%. Further using seller fixed effects as instruments, the gains from Prime increase to 9%. Last, including product ratings and reviews, the consumer gains from Prime drop to only 6%. The average within-nest correlation parameters (which differ across nests) is 0.85. This implies that consumers have strong preferences for both the platform and their nests.

Figure 4: *Sample Clustering Results: Gouda Cheese*

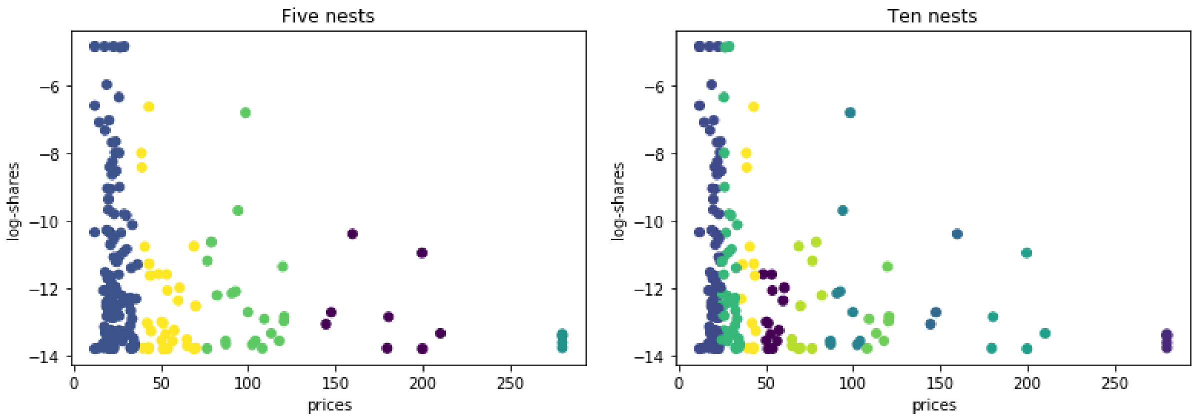

Table 4: *Estimated Demand Parameters: Gouda Cheese*

This table shows the estimated demand parameters for the "Gouda Cheese" category, as we progressively add fixed effects and instruments. All estimates are based on Gandhi et al. (2020), including product fixed effects as instruments. Prime is an indicator equal to one when the product is distributed under the Prime program. σ reports the average within nest correlation parameter (which varies across nests). The remaining parameters are standard.

	(1)	**(2)**	**(3)**	**(4)**
Constant	-3.83	-2.76	-2.74	-3.24
Prices	-0.03	-0.03	-0.03	-0.04
Prime	1.32	0.03	0.09	0.06
log(rating)				0.69
log(reviews)				0.05
$\bar{\sigma}$	0.76	0.85	0.85	0.84
Obs	11584	11584	6133	6133
Time + Prod FE	N	Y	Y	Y
Seller FE Ins	N	N	Y	Y

Last, table 5 shows the implied elasticities, as well as the estimated supply parameters (for Gouda and a few other categories in our sample). For Gouda cheese (first column), the median own elasticity is -4.32, and the implied diversion ratios are 26%. The resulting aggregate elasticity is approximately -0.4. Weighted average pass-through on referral and fulfillment fees is roughly equalized around 0.7. The remaining categories exhibit slightly higher own elasticities, with a median of -6.61 and an average of -3.9. Cross-elasticities are also high, with average outside diversion ratios of 0.18, resulting in average aggregate elasticities of -0.66.

For a monopoly platform, such low elasticities would translate to very high prices and fees. Amazon sets prices and fees that are lower than those suggested by a profit maximizer with such low elasticities, however. This suggests that consumer surplus is an important component of Amazon's objective function.[30] For Gouda, we estimate that Amazon puts 7.26x more weight on consumer surplus than it's own. On average, this parameter is 6.9.

The average estimated net fee rate charged by Amazon is 15% and the average 1P mark-up is 28%. The average 3P mark-up is 31% and the average wholesale mark-up is 45%. Interestingly, we find that Amazon rarely charges a mark-up on fulfillment services. This is because estimated pass-through on unit taxes is higher than on ad-valorem taxes (Anderson et al., 2001), so the platform would rather subsidize unit fees and set higher ad-valorem fees (especially when the weight on consumer surplus is high).

8 Counter-factuals

Equipped with an estimated model, we can now perform a series of counterfactuals that cast light on Amazon's strategy, as well as the welfare effects of proposed regulations. For each counter-factual, we adjust the market structure and estimate the new equilibrium. The results, then, measure the "short run" effects holding the number of buyers and sellers constant. In the long run, the number of buyers and sellers will adjust according to (a) buyer and seller entry elasticities and (b) the strength of network effects.

8.1 Changes in Platform Investment Incentives

We start by exploring how equilibrium quantities adjust as the platform's investment incentives change. For illustration, we present results for the Baby Scale category. Table 6 shows the results. Consistent with intuition, platform fees and prices fall as investment incentives rise. The effects are not homogeneous across instruments and activities, however. Mark-ups on 1P products react more than fees on 3P products – so that the platform uses reselling to lower prices when γ is high; and to raise prices when γ is low. Wholesalers internalize the platform's investment incentives and react by raising mark-ups as γ increases, appropriating some of the surplus. Third party sellers also slightly increase mark-ups in response to the lower fees.

[30] As a validation exercise, we solve the model with no weight for consumers. We find a very high number of products with negative marginal costs, precisely due to high Amazon fees or mark-ups. Thus, a model with no weight on consumer surplus is clearly mis-specified.

Table 5: *Summary of Estimated Parameters : Selected Categories*

This table shows key estimated demand and supply parameters and equilibrium quantities for selected Amazon categories (first five columns) and the average across a sample of ten categories. All measures report the dollar-weighted average unless otherwise noted.

Parameter			**Gouda Cheese**	**Macadamia nuts**	**Dry erase sheets**	**Baby mattress pads**	**Baby scales**	**Avg of 10 categories**
Demand	**% of Sales**	**1P**	2%	18%	13%	42%	21%	35%
		FbA	76%	79%	85%	58%	76%	60%
		FbM	22%	3%	2%	0%	3%	5%
	Prime	**Prime coeff**	0.09	0.10	0.85	0.51	0.23	0.29
	Elasticities	**Median** ε_{own}	-4.32	-5.25	-7.38	-7.82	-6.16	-6.61
		ε_{own}	-1.99	-3.71	-5.85	-4.77	-4.85	-3.93
		θ	0.26	0.18	0.09	0.16	0.18	0.18
		Agg ε	-0.43	-0.68	-0.53	-0.74	-0.86	-0.66
	Pass-through	ρ_r	0.69	1.15	0.97	0.83	0.91	0.78
		ρ_{fba}	0.71	1.04	1.11	0.99	1.01	0.83
Supply	**CS weight**	γ	7.26	5.63	7.61	3.37	10.39	6.95
	Fees	**Referral fee**	0.11	0.15	0.15	0.15	0.15	0.14
		FbA mark-up	0.95	0.00	0.00	0.00	0.00	0.27
		Total fee rate	0.16	0.15	0.15	0.15	0.15	0.15
	Prices	**Price**	17.81	27.82	18.93	17.52	49.37	25.34
		3P mark-up	0.74	0.28	0.19	0.25	0.23	0.31
		WH mark-up	0.75	1.10	0.20	0.22	0.24	0.45
		1P mark-up	0.27	0.51	0.20	0.23	0.19	0.28
	Other	**Inside share**	0.23	0.18	0.58	0.33	0.29	0.35
		% mc < 0	0.01	0.01	0.01	0.06	0.00	0.04

Table 6: *Equilibrium objects by* γ*: Baby Scale*

This table shows counter-factual fees, prices and changes in welfare as we vary the platform's investment incentive γ. The middle column represents our baseline estimate.

		γ		
		8.39	**10.39**	**12.39**
Fees	**Avg. referral fee**	0.23	0.15	0.07
	Avg. FbA mark-up	0.00	0.00	0.00
	Total fee rate	0.23	0.15	0.07
Share	**AMZ share**	0.27	0.29	0.31
Prices	**Avg. price**	52.9	49.4	45.9
	3P mark-up	0.22	0.23	0.24
	WH mark-up	0.21	0.24	0.28
	1P mark-up	0.28	0.19	0.08
ΔSurplus	**Consumers**	-4.14	0.00	3.88
	Sellers	-0.15	0.00	0.12
	Wholesalers	-0.09	0.00	0.11
	Amazon	1.12	0.00	-1.27
	Total	-3.26	0.00	2.84

8.2 Counter-factual 2: Regulatory Interventions

Next, we consider four interventions inspired by the proposed regulation.

Pure Marketplace. Our first intervention turns Amazon into a "pure" marketplace. We assume all products stay on the platform, but transition to FbM distribution. The gains from Prime disappear and the platform is left with a single ad-valorem fee instrument. In addition, we assume that all SbA products begin to be offered directly by wholesalers so that products are still subject to double marginalization (Amazon fees + seller/wholesaler). These are the most optimistic assumptions. In reality, some of the products will exit, while others will only be offered by independent third party sellers subject to triple marginalization (Amazon, seller, wholesaler).

Table 7 shows results for two sample categories. The results depend crucially on distribution channels and preferences for Prime. In the case of Tortellini, Amazon is using reselling to lower prices. The intervention, then, eliminates (a) the gains from Prime and (b) the gains from pricing pressure from direct reselling. The associated loss of utility lowers consumer preferences for the platform, which in turn induces Amazon to lower total fees from 14% to 13%. But this is not enough. The inside share falls by 5% leading to a significant decline in consumer welfare. Seller welfare rises (given the decline in competition from Amazon), while wholesaler welfare falls (since wholesalers previously benefited from Prime). Platform welfare remains relatively stable.

In the case of Baby Mattress pads, Amazon uses reselling to raise prices. Following the intervention, it reacts by raising fees. Quantities fall drastically from 33% to 23% and shift towards cheaper products that were not Prime eligible (given the comparative gain in utility). Consumer, wholesaler and platform welfare falls, while seller welfare remains stable.

Pure Reseller. Our second intervention turns Amazon into a "pure" reseller. This is Amazon's purported response to the regulation (Amazon, 2020). We assume that Amazon continues to sell SbA products and begins to sell all FbA products. These are products currently fulfilled from Amazon warehouses, hence it is feasible for Amazon to absorb them. Nonetheless, the long tail of products currently offered under FbM would exit the platform.

In the case of Gouda cheese, a significant portion of sales are FbM. The loss of variety decreases preferences for the platform, inducing Amazon to cut mark-ups and pass-through. Wholesalers internalize this and raise mark-ups slightly. Average prices fall but the decrease is not enough to offset the loss of variety, leading to a 2% decline in shares. Welfare falls for buyers and sellers (since many of these are kicked off the platform), while it increases for wholesalers who remain on the platform.

Combined, the above two conterfactuals yield two important lessons for regulators: first, consumers value both the Prime program and product variety. Interventions that eliminate either of the two are likely to decrease welfare on the platform. Second, regulatory interventions have important distributional effects between buyers, sellers, wholesalers and the platform. Since Amazon typically offers products from large established brands, turning Amazon into a reseller will hurt smaller producers and help more mature manufacturers.

Table 7: *"Pure Marketplace" Counterfactual*

This table shows the equilibrium fees, prices and welfare when we turn Amazon into a "pure" marketplace, like Ebay. We assume all products stay on the platform, but transition to FbM distribution. When computing changes in seller/wholesaler surplus, we hold the identity of agents fixed even as the distribution methods change.

		Tortellini		Baby mattress pads	
		Base	**Mktplace**	**Base**	**Mktplace**
Distn	**% SbA**	0.95	0.00	0.42	0.00
	% FbA	0.04	0.00	0.58	0.00
	Prime gains	0.2	0.2	0.51	0.51
Fees	**Avg. referral fee**	0.13	0.13	0.15	0.23
	Avg. FbA mark-up	1.45		0.00	
	Total fee rate	0.14	0.13	0.15	0.23
Share	**AMZ share**	0.36	0.31	0.33	0.23
Prices	**Avg. price**	18.7	18.4	17.5	16.1
	3P mark-up	0.23	0.35	0.25	0.29
	WH mark-up	0.35		0.22	
	1P mark-up	0.06		0.23	
ΔSurplus	**Consumers**		-3.60		-7.62
	Sellers		1.76		0.13
	Wholesalers		-2.21		-0.50
	Amazon		0.44		-0.15
	Total		-3.61		-8.14

Table 8: *"Pure Reseller" Counterfactual*

This table shows the equilibrium fees, prices and welfare when we turn Amazon into a "pure" reseller, like Costco. We assume that Amazon continues to sell SbA products and begins to sell all FbA products; yet all FbM products exit. When computing changes in seller/wholesaler surplus, we hold the identity of agents fixed even as the distribution methods change.

		Gouda cheese		Baby scale	
		Base	**Reseller**	**Base**	**Reseller**
Distn	**% SbA**	0.02	1.00	0.21	1.00
	% FbA	0.74		0.76	0.00
Fees	**Avg. referral fee**	0.11		0.15	
	Avg. FbA mark-up	0.95		0.00	
	Total fee rate	0.16		0.15	
Share	**AMZ share**	0.22	0.20	0.29	0.29
Prices	**Avg. price**	18.12	16.86	49.4	49.4
	3P mark-up	0.54		0.23	
	WH mark-up	0.55	0.56	0.24	0.22
	1P mark-up	0.37	0.19	0.19	0.14
ΔSurplus	**Consumers**	NA	-4.38		0.05
	Sellers		-1.68		-2.66
	Wholesalers		1.76		2.50
	Amazon		0.10		-0.28
	Total		-4.20		-0.39

Table 9: *"Ban Reselling" Counterfactual: Gouda Cheese*

This table shows the equilibrium fees, prices and welfare when we ban reselling. We assume that all SbA products transition to FbA, and begin to be sold by wholesalers. When computing changes in seller/wholesaler surplus, we hold the identity of agents fixed even as the distribution methods change.

		$\gamma = 4$		$\gamma = 7$	
		Base	**Ban SbA**	**Base**	**Ban SbA**
Fees	**Avg. referral fee**	0.16	0.16	0.05	0.06
	Avg. FbA mark-up	0.95	0.96	0.94	0.94
	Total fee rate	0.22	0.22	0.10	0.11
Share	**AMZ share**	0.22	0.22	0.23	0.23
Prices	**Avg. price**	18.26	17.75	17.95	17.91
	3P mark-up	0.56	0.57	0.51	0.52
	WH mark-up	0.45		0.71	
	1P mark-up	0.48		0.19	
ΔSurplus	**Consumers**	NA	0.15		-0.17
	Sellers		0.00		0.03
	Wholesalers		-0.02		-0.03
	Amazon		-0.02		0.04
	Total		0.11		-0.12

Ban Reselling. Our next two counterfactuals aim to increase competition while preserving the gains from Prime and from product variety. We begin by banning reselling. This has been progressively implemented in India starting in 2013.[31] We assume that all SbA products transition to FbA, and begin to be sold by wholesalers. As above, this maintains only double – instead of triple – marginalization. The Prime program is preserved, but Amazon loses the gains from vertical integration. Table 9 shows the results.

The welfare consequences depend on the platform's dynamic incentives, γ. When γ is high, the platform uses reselling to lower prices. Banning reselling, then, leads to a decline in consumer surplus. Platform and seller welfare increases slightly, while consumer welfare drops. The opposite is true when γ is low: reselling is used to raise prices so banning it increases consumer surplus (with small losses in seller and platform surplus).

This yields our next main result: optimal regulation is product- (and platform-)specific. Banning the same instrument may increase welfare in mature categories but may decrease it in nascent ones. Regulation, then, should aim to preserve the gains from flexibility in fee instruments when competition is fierce, but prevent the costs of flexibility as competition decreases. This points towards either (i) separate regulation for dominant and nascent platforms or (ii) regulation that is robust to changes in market power.

Introduce Seller-Fulfilled-Prime. Last, we consider the (forced) introduction of a "Seller Fulfilled Prime" program, which requires Amazon to give the Prime checkmark to any seller that can consistently meet pre-specified performance metrics. This preserves the gains from Prime but enables competition in fulfillment services. Such a program previously existed, but has been progressively closed by Amazon over the past few years. We assume that all products continue to be distributed in their current methods, but the increased

[31] See here, for example.

Table 10: *"Seller-Fulfilled Prime" Counterfactual: Gouda Cheese*

This table shows the equilibrium fees, prices and welfare when we introduce a seller-fulfilled-Prime program, which pushes Amazon's FbA mark-up to zero.

		Gouda cheese		Baby Scale	
		Base	**SFP**	**Base**	**SFP**
Distn	**% SbA**	0.02	0.02	0.21	0.21
	% FbA	0.74	0.77	0.76	0.76
Fees	**Avg. referral fee**	0.11	0.12	0.15	0.15
	Avg. FbA mark-up	0.95	0.00	0.00	0.00
	Total fee rate	0.16	0.12	0.15	0.15
Share	**AMZ share**	0.22	0.23	0.29	0.29
Prices	**Avg. price**	18.1	17.1	49.4	49.4
	3P mark-up	0.54	0.59	0.23	0.23
	WH mark-up	0.55	0.60	0.24	0.24
	1P mark-up	0.37	0.31	0.19	0.19
△Surplus	**Consumers**		0.73		0.00
	Sellers		0.08		0.00
	Wholesalers		0.00		0.00
	Amazon		-0.19		0.00
	Total		0.62		0.00

competition pushes Amazon's mark-ups on fulfillment services to zero. Table 10 shows the results.

For Baby Scale and many other categories, the policy is non-binding since Amazon optimally sets a zero mark-up on FbA (see above for discussion). For selected categories such as Gouda, however, Amazon optimally charges a positive FbA tax. In that case, the intervention leads to a decline in total fees from 16% to 12%. Interestingly, referral fees increase in response to the regulation, offsetting some of the initial gains. Equilibrium prices fall, and consumer surplus rises, at the expense of platform welfare. Sellers and wholesalers benefit slightly from the transition.

This yields our last main result: fee instruments are substitutes from the perspective of the platform. Regulatory interventions that ban individual instruments may be offset by the endogenous response of (existing and new) instruments. This points regulators towards interventions that are robust to the endogenous response of fees.

8.3 Summary of Regulatory Counter-factuals

To conclude, table 11 shows the average counter-factual welfare implications across 10 product subcategories. Consistent with the results above, turning Amazon into a pure marketplace or a pure reseller leads to substantial declines in consumer and total welfare. Sellers benefit from the transition to a marketplace while wholesalers benefit from turning Amazon into a reseller. Given our estimated parameters, banning reselling also decreases consumer welfare. However, this depends on the estimated weight on consumer surplus and, therefore, in our market size assumption. Last, increasing competition in fulfilment increases consumer and total surplus, but only slightly since the intervention is rarely binding.

Table 11: *Average Counterfactual Results: Selected Categories*

This table shows the average welfare implications of regulatory interventions across 10 random subcategories. Changes in sellers/wholesaler surplus hold the identity of agents fixed even as the distribution method changes.

		Mktplace	Reseller	Ban SbA	SFP
ΔSurplus	**Consumers**	-5.39	-2.34	-1.78	0.10
	Sellers	0.23	-1.10	0.63	0.01
	Wholesalers	-0.74	0.92	-0.74	0.00
	Amazon	0.07	0.61	0.63	-0.03
	Total	-6.00	-2.30	-1.27	0.09

9 Conclusion

This paper develops and estimates a structural model of the Amazon platform, and uses it to study the "short-run" implications of regulatory interventions. Theoretically, it shows that Amazon's business model, fee policies and market power have important implications for the welfare consequences of regulation. This points regulators towards platform- or market-specific interventions that are robust to the endogenous response of platform fees and business models. Empirically, it shows that interventions that eliminate either the Prime program or product variety are likely to decrease welfare. This points regulators towards interventions that preserve Prime and product variety but increase competition.

These insights are based on selected categories and depend crucially on some estimation assumptions. In future versions of this paper, I hope to (a) scale up the analyses to hundreds of categories in order to perform richer cross-sectional analyses; (b) explore alternate fee structures such as advertising and steering; and (c) consider alternate estimation assumptions around the size of the market and Amazon's fee-setting behavior.

In the long run, there is further research needed to understand the determinants between business models, as well as the role of entry and competition for equilibrium outcomes. I pursue some of these in companion papers. Gutierrez (2022a) studies the determinants of reselling vs. marketplace intermediation. Gutierrez (2022b) extends the model to include buyer and seller entry decisions; and uses a reduced form Differences-in-Differences strategy to quantify the welfare consequences of private label product introduction.

References

Adachi, T. and M. Fabinger (2020). Multi-dimensional pass-through and welfare measures under imperfect competition. Working Paper.

Aguirre, I., S. Cowan, and J. Vickers (2010, sep). Monopoly price discrimination and demand curvature. *American Economic Review 100*(4), 1601–1615.

Almagro, M. and E. Manresa (2021). Data-driven nests in discrete choice models. Unpublished.

Amazon (2020). Fringe notions on antitrust would destroy small businesses and hurt consumers.

Anderson, S. and O. Bedre-Defolie (2017). Optimal variety and pricing of a trade platform.

Anderson, S. and O. Bedre-Defolie (2020). Hybrid platform model.

Anderson, S. P., A. de Palma, and B. Kreider (2001, aug). Tax incidence in differentiated product oligopoly. *Journal of Public Economics 81*(2), 173–192.

Andrews, D. W. K. and X. Shi (2013). Inference based on conditional moment inequalities. *Econometrica 81*(2), 609–666.

Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics 37*(3), 668–691.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*.

Berry, S., J. Levinsohn, and A. Pakes (2004, feb). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy 112*(1), 68–105.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics 25*(2), 242.

Bonhomme, S., T. Lamadon, and E. Manresa (2019). Discretizing unobserved heterogeneity. Unpublished.

Bonnet, C. and P. Dubois (2010, March). Inference on vertical contracts between manufacturers and retailers allowing for nonlinear pricing and resale price maintenance. *41*(1), 139–164.

Brynjolfsson, E., Y. J. Hu, and D. Simester (2011, aug). Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science 57*(8), 1373–1386.

Brynjolfsson, E., Y. J. Hu, and M. D. Smith (2003, nov). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science 49*(11), 1580–1596.

Brynjolfsson, E., Y. J. Hu, and M. D. Smith (2006). From niches to riches: Anatomy of the long tail. *Sloan Management Review*.

Brynjolfsson, E. and M. D. Smith (2000, apr). Frictionless commerce? a comparison of internet and conventional retailers. *Management Science 46*(4), 563–585.

Brynjolfsson, E. and M. D. Smith (2001). The great equalizer? consumer choice behavior at internet shopbots. *MIT Sloan Working Paper*.

Cabral, L. (2018). Standing on the shoulders of dwarfs: Dominant firms and innovation incentives. Working paper.

Cavallo, A. (2017, jan). Are online and offline prices similar? evidence from large multi-channel retailers. *American Economic Review 107*(1), 283–303.

Cavallo, A. (2018, mar). Scraped data and sticky prices. *The Review of Economics and Statistics 100*(1), 105–119.

Chen, N. and H.-T. Tsai (2021). Steering via algorithmic recommendations. Working paper.

Chevalier, J. and A. Goolsbee (2003). Measuring prices and price competition online: Amazon and barnes and noble. *Quantitative Marketing and Economics 1*(2), 203–222.

Chevalier, J. A. and D. Mayzlin (2006, aug). The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research 43*(3), 345–354.

Chintagunta, P. K., A. Bonfrer, and I. Song (2002, oct). Investigating the effects of store-brand introduction on retailer demand and pricing behavior. *48*(10), 1242–1267.

Congress (2020). Invetigation of competition in digital markets. Technical report, Antitrust Subcommittee. Congress,.

Conlon, C. T. and J. H. Mortimer (2021, dec). Efficiency and foreclosure effects of vertical rebates: Empirical evidence. *Journal of Political Economy 129*(12), 3357–3404.

Conlon, C. T. and N. Rao (2015). The price of liquor is too damn high: State facilitated collusion and the implications for taxes.

Cornière, A. and G. Taylor (2019, oct). A model of biased intermediation. *50*(4), 854–882.

Covarrubias, M., G. Gutiérrez, and S. Piton (2022). The amazon shock: Pass-through estimation through amazon annual fee changes.

Crawford, G. S., R. S. Lee, M. D. Whinston, and A. Yurukoglu (2018). The welfare effects of vertical integration in multichannel television markets. *86*(3), 891–954.

Cure, M., M. Hunold, R. Kesler, U. Laitenberger, and T. Larrieu (2021). Vertical integration of platforms and product prominence. Working paper.

Dolfen, P., L. Einav, P. Klenow, B. Klopack, J. Levin, L. Levin, and W. Best (2019, feb). Assessing the gains from e-commerce. Technical report.

Dubé, J.-P., A. Hortaçsu, and J. Joo (2020, feb). Random-coefficients logit demand estimation with zero-valued market shares. Technical report.

Ellickson, P. B., P. Kong, and M. J. Lovett (2018). Private labels and retailer protability: Bilateral bargaining in the grocery channel. Working Paper.

Etro, F. (2020). Product selection in online marketplaces.

Farrell, J. and M. L. Katz (2003, mar). Innovation, rent extraction, and integration in systems markets. *The Journal of Industrial Economics 48*(4), 413–432.

Forbes, K. F. (1988). Pricing of related products by a multiproduct monopolist. *Review of Industrial Organization*.

Gandhi, A., Z. Lu, and X. Shi (2020). Estimating demand for dierentiated products with zeroes in market share data. Unpublished.

Gomes, R. and A. Mantovani (2020). Regulating platform fees under price parity. CEPR Discussion Paper No. DP15048.

Gutierrez, G. (2022a). The determinants of amazon selling methods.

Gutierrez, G. (2022b). The seller welfare standard.

Hagiu, A. (2009, dec). Two-sided platforms: Product variety and pricing structures. *18*(4), 1011–1043.

Hagiu, A., T.-H. Teh, and J. Wright (2020). Should amazon be allowed to sell on its own marketplace? *SSRN Electronic Journal*.

Hagiu, A. and J. Wright (2015, jan). Marketplace or reseller? *Management Science 61*(1), 184–203.

Hamilton, S. F. (2009, feb). Excise taxes with multiproduct transactions. *99*(1), 458–471.

Hortacsu, A. and C. Syverson (2015, Fall). The ongoing evolution of us retail: A format tug-of-war. *Journal of Economic Perspectives 29*(4), 89–112.

Jiang, B., K. Jerath, and K. Srinivasan (2011, sep). Firm strategies in the mid tail of platform-based retailing. *Marketing Science 30*(5), 757–775.

Johnson, J., A. Rhodes, and M. R. Wildenbeest (2020). Platform design when sellers use pricing algorithms.

Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *The American Economic Review 75*(3), 424–440.

Katz, M. L. and C. Shapiro (1994, may). Systems competition and network effects. *Journal of Economic Perspectives 8*(2), 93–115.

Khan, L. M. (2016). Amazon's antitrust paradox. *The Yale Law Journal*.

Khan, L. M. (2018). Sources of tech platform power. *Georgetown Law Technology Review*.

Khan, L. M. (2019). The separation of platforms and commerce. *Columbia Law Review*.

Kirpalani, R. and T. Philippon (2021). Data sharing and market power with two-sided platforms.

Koijen, R. and M. Yogo (2014). Shadow insurance.

Li, J. (2019). Compatibility and investment in the u.s. electric vehicle market. *Unpublished*.

Liu, C., F. Niu, and A. White (2021). Optional intermediaries and pricing restraints.

Marshall, A. (1890). *Principles of Economics*. Macmillan.

Mitchell, S. and S. Sussman (2019, November). How amazon rigs its shopping algorithm. ProMarket.

Mulhern, F. J. and R. P. Leone (1991, oct). Implicit price bundling of retail products: A multiproduct approach to maximizing store profitability. *55*(4), 63–76.

Nevo, A. (2001, mar). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Peitz, M. and M. Reisinger (2014, dec). Indirect taxation in vertical oligopoly. *62*(4), 709–755.

Petrin, A. and K. Train (2010, feb). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research 47*(1), 3–13.

Quan, T. W. and K. R. Williams (2018, sep). Product variety, across-market demand heterogeneity, and the value of online retail. *The RAND Journal of Economics 49*(4), 877–913.

Rhodes, A. (2014, sep). Multiproduct retailing. *The Review of Economic Studies 82*(1), 360–390.

Rhodes, A., M. Watanabe, and J. Zhou (2021, feb). Multiproduct intermediaries. *Journal of Political Economy 129*(2), 421–464.

Rochet, J.-C. and J. Tirole (2003). Platform competition in two-sided markets. *Journal of the European Economic Association 1*(4), 990–1029.

Rochet, J.-C. and J. Tirole (2006). Two-sided markets: A progress report. *The RAND Journal of Economics 37*(3), 645–667.

Rohlfs, J. (1974). A theory of interdependent demand for a communications service. *The Bell Journal of Economics and Management Science 5*(1), 16.

Stone, B. (2013). *The Everything Store: Jeff Bezos and the Age of Amazon*. USA: Atlantic/Little, Brown.

Varian, H. R. (1985, sep). Price discrimination and social welfare. *American Economic Review 100*(4), 1601–1615.

Villas-Boas, S. B. (2007, apr). Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies 74*(2), 625–652.

Weyl, E. G. and M. Fabinger (2013, jun). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy 121*(3), 528–583.

Zhu, F. and Q. Liu (2018, aug). Competing with complementors: An empirical look at amazon.com. *Strategic Management Journal 39*(10), 2618–2642.