The Effects of Return Policies in E-commerce*

Peter Newberry
Department of Economics
Terry College of Business
The University of Georgia

Xiaolu Zhou
Amazon

September 4, 2024

Abstract

We quantify the welfare effects of return policies in a large online retail platform. Utilizing variation in return policies across sellers of tablets on Alibaba's Tmall, we estimate a model of demand and supply where the consumer has the option to return the product after purchasing it. We find that a more lenient return policy impacts demand through a reduction of the risk to the consumer, but does not serve to provide a strong signal of seller quality. On the supply side, estimates suggest that returns are costly for sellers, as they are approximately equal to the cost of proving a brand new tablet to the buyer. We use the model to estimate the effect of platform-level return policies, with a focus on the impact of the leniency of the uniform rule. As policies become more lenient, consumer surplus decreases and firm profit increases. The reduction in consumer surplus comes from the price response of firms, as they pass through increases in return cost to consumers. This suggests that the leniency of a platform-level return policy can work to undermine the goal of improving the customer experience.

Keywords: e-commerce, return policy, incomplete information, retail, market power
JEL Codes: L11, L15, L41, L81

*`pwnewberry@uga.edu` or `xzz109@psu.edu`

1 Introduction

Over the last decade, online shopping has become a prominent feature of the global economy. The US department of commerce reports that e-retail made up about 12% of all retail sales in the first quarter of 2020, which was up from 5% in the first quarter of 2011. Since the COVID-19 pandemic, this share has increased even further to between 14 and 15% of total retail sales each quarter.[1] A well-documented difficulty that e-retailers have had to overcome in order to compete with their brick-and-mortar counterparts and produce this growth is the fact that consumers do not have physical access to the product before they purchase it. One strategy to overcome this difficulty that has become ubiquitous over time has been to offer generous return policies, reducing the risk borne by consumers.[2] Anecdotal evidences suggests that return policies impact consumer decisions in online markets, but can come with significant costs to sellers.[3] The primary goal of this paper is to quantify the impact of return policies on consumers and sellers on a large online platform.

Specifically, we ask two questions. First, what is the role of return policies in determining demand? Previous theoretical models posit that return policies can impact demand through both an 'insurance effect' (see Che (1996)) and a 'signaling effect' (see Moorthy and Srinivasan (1995)). We separately identify these two effects and quantify their relative importance in driving demand and determining consumer surplus. Second, what is the impact of implementing a platform-level return policy? By answering this question we are able to assess the costs and benefits related to both the uniformity of the policy and the degree of leniency. The outcomes of this analysis are relevant to both platforms and policy makers as they weigh rules that restrict seller return policies. For example, Amazon.com puts rules restricting return policies in place for sellers who

[1]See `https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf`, accessed on 01/2023

[2]Other strategies previously studied include consumer reviews and ratings (e.g., Chevalier and Mayzlin (2006) and Newberry and Zhou (2019)), expert reviews (e.g., Hilger et al. (2011)), entry requirements (e.g., Newberry and Zhou (2022)), sales information lists (Sorensen (2007) and Newberry (2016)), required product information (e.g., Lewis (2011)), and seller certification (e.g., Farronato et al. (2020), Hui et al. (2016)).

[3]A 2014 survey by ComScore suggests that 66% of shoppers read return policies before purchasing a product and that 58-62% of shoppers consider return policies and procedures key factors influencing their decisions. See `https://www.comscore.com/Insights/Presentations-and-Whitepapers/2014/UPS-Pulse-of-the-Online-Shopper-A-Customer-Experience-Study`. Finally, estimates put the 'reverse-logistics costs' at $381 billion in 2017 and were estimated to reach $550 billion in 2020. See `https://www.logiwa.com/blog/when-ecommerce-return-management-costs-go-down-profit-goes-up`

use their Fulfilled by Amazon program as well as for sellers who don't.[4] Additionally, the Chinese government recently began enforcing a consumer protection law requiring a minimum leniency for returns on online purchases.[5]

To answer these questions, we collect data from sales of tablets on Tmall, which is Alibaba's primary business-to-consumer (B2C) platform. Tmall features thousands of *professional* sellers offering a wide variety of products. Tmall had 54 percent of the online B2C market share and total transactions reached $39 billion in Q3 2015.[6] Crucial for our analysis is that sellers on the platform can vary in their return policy. The platform sets a default policy regarding (1) the maximum number of days between the delivery of the product and the return request, (2) whether or not the buyer needs a reason for the return, and (3) whether or not the buyer pays for the return. The seller can then vary from the default policy on each of these dimensions.

Our data include all sellers who sold tablets on Tmall between September, 2014 and December, 2014. The data include the monthly sales and prices for each tablet, as well as product and seller characteristics. The seller characteristics include a consumer rating score, which is similar to the star rating on Amazon.com, as well as their return policy. We define a 'default seller' as one who has the default policy and a 'lenient seller' as one who has a more lenient policy in one or more of the three categories. Between 75% and 90% of the sellers are default sellers. In addition to the return policy, we also observe the percentage of products that are returned in each month for each seller, or the 'return rate'. On average, about 9% of tablets that are purchases are returned in any given month, which is lower than overall return rates reported in industry reports, which is likely to the nature of the tablet product category.[7]

Using these data, we provide a preliminary examination of the relationships between a lenient policy and demand and supply. We find that a lenient policy increases sales by almost 1% for sellers that have been on the platform for less than a year, suggesting that having a lenient return policy plays an important role for new sellers who are presumably perceived as more 'risky' by consumers

[4] See `https://www.fbamasterclass.io/post/understanding-amazon-return-policy-for-amazon-sellers`

[5] See `https://www.china-briefing.com/news/china-introduces-new-consumer-protection-law/`

[6] Information from `http://www.chinainternetwatch.com/15957/chinas-b2c-sales-q3-2015/`.

[7] Industry experts suggest that between 15 and 40% of online purchases are returned, depending on the product category, whereas only about 5-10% of brick and mortar purchases are returned. See `https://www.cnbc.com/2019/01/10/growing-online-sales-means-more-returns-and-trash-for-landfills.html`.

because, for example, there is not a long track record of rating scores. We find no systematic relationship between the leniency of a policy and return rates for a seller, which is likely due to the fact that the signaling and insurance effect work in opposite directions in determining return rates. Finally, we find that prices are positively associated with return rates, pointing to the fact that returns are costly for sellers. However the significance of the relationship depends on the controls and the data used, something that is likely due to the fact that the regression doesn't control for the interactions between return policies, demand, return rates, and price.

Together, the reduced form results highlight the need for a structural model in order to disentangle the insurance and signaling effects of a return policy on the demand side and estimate the cost of returns on the supply side. Therefore, we specify and estimate a discrete choice demand model and supply for differentiated products where the consumer has the option to return the product after they purchase it. We assume that consumers have homogenous preferences and know the prices and characteristics of each product *ex ante*, but only learn the true seller quality and their idiosyncratic preferences, or 'match quality', for a tablet if they purchase it. Before purchasing, the consumer forms expectations about seller quality and her idiosyncratic tastes, where the expected seller quality is formed from the seller ratings and the return policy. The fact that expectations of quality depend on the return policy is a reduced form way of modeling the signaling effect of the policy.

Once the consumer purchases the product, she realizes the true seller quality and the true match quality. If her realized utility is greater than the utility from returning the product, which includes a buyer return cost, then she keeps the product. Otherwise, she returns it. Thus, when deciding which product to purchase, the consumer considers the expected value of keeping versus returning each product. We assume that the cost of a return, which may be a monetary cost or other costs associated with a return, is homogenous across buyers and is a function of the return policy of the seller. This represents the insurance effect of the policy.

On the supply side, we assume that sellers play a static Bertrand-Nash pricing game in each month, where their marginal cost includes an expected return cost. There are two interesting features of sellers' optimal pricing decisions in this environment that differ from a standard model without returns. First, the prices are a function of the expected return costs incurred by the sellers. The fact that the consumers who are marginal in their purchase decision are more likely to return implies an increasing cost curve, which results in advantageous selection. Therefore, as we make the return policy more lenient, the sellers' return costs increase, resulting in higher prices for all consumers, even the ones who are very unlikely to return the product. Second, the mark-up term is a function of how sensitive demand, net of returns, is to prices. Intuitively, consumers will be less sensitive to prices at the purchase stage when they are protected by the return policy and because returns are costly, the overall sensitivity of demand will be lower compared to the situation where there are no returns. The mechanism is similar to moral hazard, as consumer are less sensitive to risk as the level of insurance against risk increases, and this moral hazard translates to more market power, all else equal.

The estimates of the preference parameters show that consumers are sensitive to prices, with the average elasticity of -3.9. The buyers place a value of about \$47 (19% of the average price) on a new seller who has a lenient policy *ex ante*, but this effect is not statistically significant. The statistical insignificance of this signaling effect is likely due to other quality signals on the platform, such as rating scores. Indeed, buyers place a statistically significant value of about \$43 (17%) on sellers with a relatively high rating. In contrast, the effect of a lenient policy on the value of a return is a statistically significant \$52 (21%), which can be interpreted as the return cost being \$52 lower for sellers with a lenient policy. The magnitude of the estimate is significant considering the average buyer return cost across all sellers is about \$22 (9%).

On the supply side, we find that the marginal return cost for a seller is about \$259 (105%) and is statistically significant, which implies an average expected return cost of about \$17 for each purchase. This estimate includes the cost of processing returns, the net loss in marginal cost, and the potential cost of a replacement and/or refund. To put this in perspective, the marginal cost of a product for the average seller is estimated to be about \$136, so the level of return cost is approximately the cost of sending an equivalent replacement tablet plus a recoup value of the old table that is about equal to the cost of processing the return. Overall, the estimates of the supply side suggest that the median seller has a mark-up of about 0.34.

We compare these estimates to a model that ignores returns and find an average elasticity of -2.6, which comes from the fact that a model without returns will incorrectly interpret consumers lack of response to price due to the insurance effect as price insensitivity. Similarly, most of the preferences for product characteristics are biased towards zero. On the supply side, the bias of the price coefficient leads us to underestimate marginal costs and overestimate the market power, with the median seller now having a mark-up of about 0.58. This comparison highlights the importance in modeling returns in retail markets, as any model that ignores them may overestimate the market power of the firms.

We use the estimated model to perform two exercises. First, we quantify the relative importance of the the insurance and signaling effects of a return policy on consumers by shutting down both effects together and then adding each effect individually, keeping prices fixed. Shutting down both effects reduces expected consumer welfare by 5%, as the value of buying from one of the sellers on Tmall decreases. If we add the signaling effect of the observed policies, but do not allow returns, consumer welfare increases by 1.0%, as demand increases for sellers who are able to signal their quality through a lenient policy. If, instead, we add the insurance effect of the observed policy, welfare increases by about 6%. Overall, the insurance effect accounts for about 80% of the value of observed return policies. On the seller side, we find that the signaling effect plays a pivotal role in determining demand for lenient sellers, whereas the insurance effect impacts both types of sellers. This exercise must be caveated by the fact that there are very few lenient sellers in the data, meaning there are some mechanics leading to signaling being less important. But the fact that the point estimate of the signaling effect is insignificant provides further evidence that the insurance effect of return policies dominate in this market.

Second, we perform a simulation where the platform sets a uniform return policy. To start, we form a baseline by shutting down returns completely. We then examine the impacts of moving from a no-return policy to one where all sellers set the default policy, which allows us to quantify the impact of the leniency of the uniform policy. If we perform this exercise assuming that firms cannot adjust prices, the more lenient policy results in an increase in consumer surplus of about 5%. This comes from the fact that the value of buying from a Tmall seller is now higher with the option to return. Firm profit is reduced by about 17% despite the increase in sales, which is due to the increase in return costs and the default sellers' profit is reduced by a greater amount because consumers are more likely to return their products. These results provide a quantification of the value of returns to consumers and the cost of returns to firms.

When we allow firms to adjust prices, things change. A move from a ban on returns to the default policy results in 3% price increase, on average, and over half of that price increase is due to the increase in market power. Another significant portion of the price increase comes from the increase in the return costs that are passed through to consumers. This results in an increase in the median mark-up of around 7% and the lower (in absolute value) price elasticity (-5.6 v -5.1 elasticity). Overall, consumer surplus falls by 11%. Seller profit increases modestly overall, as the increase in market power comes with increases in costs. Though, the move to a more lenient policy increases the profit for lenient sellers by more than default sellers because their returns are not increasing as much.

We then do the same exercise, but allow for the signaling effect of the policy to remain in place, which could be conceptualized as the sellers being able to signal in another way. It also serves to net out the signaling effects of a uniform policy. With signaling in place, prices increase by 3% and consumer surplus falls by 11% moving from a ban on returns to the default policy. Firms again benefit from moving to the default policy, but magnitude of the change is smaller than the magnitude without signaling. Overall, allowing signaling only slight dampens the impact of changing the leniency of the policy.

As a final quantification, we compare the outcomes in the observed market to ones where there is a uniform default policy rule. This type of rule means that the lenient sellers are forced to switch to a default policy, so return policies are becoming more strict overall. If this policy change eliminates signaling in addition to changing the return policy for lenient sellers, then it will result in lower consumer surplus and slightly lower profit. However, the loss in consumer surplus comes primarily through the signaling effect. If we allow for signaling under a uniform default policy, consumer surplus is higher than the observed market. This is inline with the previous results that more strict policies improve consumer outcomes.

In summary, the simulation results suggest that requiring sellers to set more lenient return policies may have had an adverse effect on consumer welfare. This is mainly working through the fact that more lenient policies result in higher costs to sellers, which are passed on to the consumer, and more market power through lower price elasticities. The fact that lenient return policies can decrease consumer surplus is something that comes out of a theoretical model developed in Vo (2022).

There are a couple caveats to our analysis. The first is that the market for tablets may not be representative of other markets which feature more uncertainty. In these markets, the consumer benefits of a lenient policy may outweigh the pricing effects. The second is that our analysis is in the short run. There could be long run effects of changing return policies in terms of market expansion, returning customers and entry/exit of firms. While these critiques limit the generalizability of our results, our analysis highlights mechanisms that are important to consider in any future work that examines return policies.

This paper contributes to a number of streams of literature. First, there is a large literature that studies factors that impact demand in online platforms. Quan and Williams (2018) examines the role of product variety, Chevalier and Mayzlin (2006) estimates the effect of consumer ratings, Ellison and Ellison (2009) analyzes the impact of price obfuscation, De los Santos et al. (2012) tests search models and estimates search cost, and Einav et al. (2014) quantifies the impact of taxes online, to name just a few. We contribute to this literature by examining the impact of return policies, which anecdotally, are critical factor in the functioning of an online retail platform. Second, there are an increasing number of empirical papers that analyze the impact of platform-level policies in online markets. Hui et al. (2021) examine a change in the thresholds to become a certified seller on eBay, Dinerstein et al. (2018) look at eBay's design of their searching mechanism, and Gutierrez Gallardo (2021) examines the impact of different vertical structures on Amazon that may arise due to regulation. See Baye and Prince (2020) for a good overview of issues related to online platforms and their regulation. In this paper, we focus on the a policy that restricts seller return policies, something that resembles policies implemented on Amazon and through a consumer protection law by the Chinese government. Third, there are a number of of theoretical and empirical papers in economics and marketing that examine product returns. These include papers by Vo (2022), Janssen and Williams (2022), Ibragimov (2022), Sahoo et al. (2018), Courty and Hao (2000), Anderson et al. (2009), Inderst and Ottaviani (2013), Inderst and Tirosh (2015) Petrikaitė (2018). To our knowledge, this is the first empirical paper that estimates a structural demand and supply model with the option to return, which allows us to quantify the effect of rules that restrict seller return policies. Finally, this is related to the industrial organization papers that estimate demand for differentiated products in a discrete choice framework, pioneered by Berry (1994). We demonstrate the degree to which estimated demand and supply parameters are biased if we ignore the consumer's option to return the product, highlighting the importance of considering returns in online and offline retail markets where returns are an important driver of demand.

The rest of the paper is organized as follows. Section 2 discusses the data and provides preliminary evidence of the different effects of the return policy. In section 3, we specify the model and discuss the implications of returns on supply and demand. In section 4 and 5, we present the estimates of the model and results of the simulations. Section 6 concludes.

2 Example

In this section, we specify a stylized theoretical example demonstrating the effects of changing the leniency of a return policy. We assume there exists one monopolist seller facing a demand curve given by $D(P, \gamma)$, where γ represents the leniency of the return policy. The slope, D_p, depends on the value of γ such that the second derivative of the demand function with respect to γ is positive, or $D_{p\gamma} < 0$. This is rationalized by an assumption that consumers who have high willingness to pay are less likely to return the product, meaning they are less affected by changes in the leniency of the policy.

The supply curve, $S(P, \gamma)$, is assumed to be upward sloping, $S_p > 0$. This comes from the fact that there is 'advantageous selection' in this market, as increases in sales are coming from the consumers that are more likely to return the product (the ones with lower willingness to pay). The slope of the supply curve again depends on the leniency of the policy such that $S_{p\gamma} < 0$.

and a continuum of buyers indexed by i. Each buyer has the following indirect utility of purchasing the product from the seller:

$$u_i = \delta_i - p + v(\delta_i) \tag{1}$$

and an indirect $u_i = 0$ for the outside option. The term δ_i represents consumer i's consumptions value for the product, which is assumed to be uniformly distributed between 0 and 1, p is the price, and the last term in the is the expected value the consumer gets from being able to return the product. The return value is a reduced form function encompassing the likelihood consumer i will return the product, the net value of consumption between the outside option and the product, and the cost of making a return. In this exercise, we are abstracting away from the price effects of returns, as it isn't essential to derive the comparatives statics we are interested in, but we will incorporate that into our empirical model in Section #. We parametrize this function as:

$$v(\delta_i) = \lambda\gamma(1 - \delta_i) \tag{2}$$

where the term in parentheses captures the fact that consumers with higher ex ante preferences are less likely to return the product and therefore, have lower value of returns. The term γ measures the leniency of the return policy, which directly impacts both the supply and the demand side, and the term λ is a demand-side shifter of the utility associated with a return such as the customer 'hassle cost'. The combined value of $\lambda\gamma$ represents the slope of the relationship between consumption value and overall expected return value, meaning it impacts the amount of heterogeneity in the value of returns across consumers.

Given the uniform assumption on δ, we can derive the inverse demand curve for the product as:

$$p = 1 - (1 - \gamma\lambda)s \tag{3}$$

where s is the share of customers who purchase the product.

On the supply side, we assume that the firm's marginal cost is made up of production cost c

and return costs function such that:

$$mc = c + z(\bar{\delta}) \tag{4}$$

where $\bar{\delta}$) represents the value of δ_i such that consumer consumer i is indifferent between purchasing the product and taking the outside option. We parameterize this function as:

$$z(\bar{\delta}) = \rho\gamma(1 - \bar{\delta}) \tag{5}$$

meaning the supply curve is given by:

$$mc = c + \rho\gamma(1 - \bar{\delta}) = c + \rho\gamma s \tag{6}$$

where the last inequality comes from the fact that $s = (1 - \bar{\delta})$. The term γ again shifts the leniency of the return policy and ρ shifts the cost of returns to the seller. Notice that the marginal cost is increasing in s, which comes from the fact that the people who are buying the product at lower prices are more likely to return the product. Therefore, this is a market that features advantageous selection. That is, as we expand demand, it becomes more costly to provide the product.

Finally, suppose there is a platform on which the seller operates that can adjust the leniency of the return policy through γ. The platform primary cares about revenue, as they earn royalties from sellers and consumer surplus, as they want to keep customers happy (i.e., they want them to return). When the platform increases γ this results in a rotation outward of demand and supply, as seen in figure #.

3 Data

The data used for this analysis come from the online platform Tmall, which is Alibaba's business-to-consumer marketplace. Tmall features mostly professional sellers and enterprises, making it is similar to Amazon Marketplace in the US. Tmall's sales reached $180 billion in 2015 and, which together with Alibaba's other platform Taobao, account for about 80% of China's e-commerce.[8,9] For comparison, Amazon's sales in 2015 were $107 billion.[10] Tmall's sellers span many many product categories, but our focus is on tablets.

Tmall has a few features that are intended to alleviate the problems associated with the information asymmetries that are common in online markets, most notably it has a system that allows consumers to rate their experience with a given seller/product across three different aspects of the transaction: description, customer services, and shipping. The seller's average rating (out of 5) in each category is displayed to consumers, where the average is across all the transactions for the seller that occurred in the previous six months. We combine the three scores into one rating by taking the simple average, as in Newberry and Zhou (2019), and refer to this as the sellers 'rating score'.

Sellers on Tmall choose their own return policy across three dimensions. The first dimension is the maximum acceptable number of days between the purchase of the product and the request for a return. The default for this dimension on Tmall is one week (seven days). The second dimension is whether or not the buyer needs a reason to return the product, such as 'it was broken' or 'it was the wrong product'. The default policy indicates that the buyer doesn't have to have a reason to return the product, so buyers can return a tablet because they 'didn't like the look of it', for example. The final dimension determines who pays the cost for shipping of the return, and the default is that the buyer pays for shipping. In spring of 2014, the Chinese government enacted a consumer protection policy that required all online sellers to accept returns for any reason within 7 days (see Guan (2020)).

We collected most of the data by scraping Tmall's website every two weeks during the period September, 2014 to December, 2014. For each scrape, we collect the total quantity sold in the previous month and the current price for every tablet that is available on Tmall at that time. To calculate a monthly price for each tablet, we take the average of the current scraped price and the

[8] Statistics can be found at `http://www.alibabagroup.com/en/ir/glance`.

[9] Information is from Alibaba's financial report, downloaded from `http://www.alibabagroup.com/en/ir/financial`.

[10] Information is from eBay's annual financial report downloaded from `https://investors.ebayinc.com/annuals.cfm`.

price from the previous scrape. Table 1 describes the market. After cleaning the data, we observe 203 total sellers in total and about 158 per month.[11] Nearly 50,000 tablets are sold each month, resulting in 9.5 million dollars in revenue. Almost 500 unique tablets were listed during our sample, based on tablet characteristics, and an average of about 350 unique tablets are listed per month. The first three rows and in the far right panel of Table 2 display averages of seller level performance measures. The average price is $250, the average seller sells about 280 tablets per month, which totals nearly $60,000 in revenue. Recall that these are professional sellers, so these performance numbers are much larger than we would see on the consumer-to-consumer marketplace Taobao. Additionally, the final column shows that there is a lot of variation in these measures.

Table 1: Description of the market

	All	
	Total	Average per month
No of sellers	203.00	158.50
Products	491.00	346.25
Quantity (000s)	182.33	45.58
Revenue (000s$)	38,338.75	9,584.69
Returns (000s)	20.85	5.21

Notes: Displayed are the aggregate statistics for sellers of tablets on Tmall both over the entire 4 month sample period and the average per month. Prices are converted to US dollars using an exchange rate equal to 6.33.

For each seller/month, we also observe a number of statistics about the quantity of returns. First, we observe the return rate, or the percentage of total products sold that are returned. Second, we observe the the total quantity of returns. Third, we observe the number of returns that fall under three different reasons: product quality, no product received, and no reason. There is a residual number of returns between the sum of these categories and the total, which we label as 'other'. One weakness of return data is that they are aggregated across all products that a seller offers. For example, if a store sells multiple products, then the return data will indicate the number of returns across all of those products. We assume that the return rate observed is equal to the

[11] We remove seller/months that have extreme values of return rates, which amounts to 5% of the original sample, and observations that appear to be phones mislabeled as tablets, which amounts to 2.3% of the original sample

return rate for the tablet category. Most of the sellers in our data sell only electronic goods, so this assumption is valid in the case that electronic goods in different categories have similar return rates.

The average return rate across sellers is about 9% (see row 4 in the right panel of Table 2), which is lower than estimates in industry reports.[12] This could be for many different reasons, but one that deserves discussion is the product category for which we have data. Presumably, there is not a lot of uncertainty when a consumer is choosing a tablet, or other electronic good. A product category like apparel, which features a lot of uncertainty in terms of how items will fit the consumers, likely has a higher return rate. So why do consumers return their tablet? While we do not observe the exact reason, we suspect that seller driven reasons like sending the wrong item or the item being broken due to poor packaging play a large role. Additionally, there may be some unknown brands which may or may not meet the consumers' expectations. Given that uncertainty doesn't play as big of role in this product category, our analysis may be considered a lower bound on the impact of returns.

In addition to the demand and return data, we collect a number of product and seller characteristics. The product characteristics are the brand, operating system, screen size, memory, storage, cellular internet capability, and the number of years since the product's introduction (i.e., product age) for each tablet. The second section of the right panel of Table 2 displays the average tablet characteristics across all sellers. The brand information indicates that there is not one dominant brand that is sold on Tmall.

The seller characteristics include the ratings scores for each seller at the time of each scrape, the tier of city the seller is located in, how long the seller has been on Tmall, and the seller 'type' as local or national in terms of their offline presence, as defined by Newberry and Zhou (2019). The final section of the right panel of Table 2 displays the average seller characteristics across all sellers. Around 75% of sellers are located in a large city, where a large city is defined as one of the 5 largest cities in China, and 21% are national retailers. Most sellers have been on Tmall for a long time, as the average age is about 30 months and 35% of sellers are highly rated, where a high

[12] See `https://www.cnbc.com/2019/01/10/growing-online-sales-means-more-returns-and-trash-for-landfills.html`.

rating is defined as having an average rating score (across the three ratings) equal to or above 4.8 out of 5. Finally, 27% of sellers are mid-rated, were a mid-rating is defined as having an average rating score between 4.7 and 4.8 out of 5.

Table 2: Characteristics by Policy

	Default Policy		Lenient Policy		All	
Variable	Mean	StDev	Mean	StDev	Mean	StDev
Seller Performance						
Revenue	63,737	173,510	33,966	57,739	60,433	164,915
Quantity	304	878	153	200	288	832
Price	252	200	204	192	247	197
Return Rate	0.09	0.07	0.07	0.04	0.09	0.07
% Return Quality	0.06	0.07	0.10	0.09	0.07	0.08
% Return Not Received	0.04	0.06	0.06	0.09	0.04	0.06
% Return No Reason	0.35	0.21	0.33	0.15	0.35	0.20
Product Characteristics						
Screen Size	8.48	1.09	8.30	0.91	8.46	1.07
Data	0.29	0.35	0.22	0.23	0.28	0.34
OS						
Android	0.74	0.39	0.84	0.31	0.75	0.39
Windows	0.23	0.37	0.15	0.29	0.22	0.37
Age	0.39	0.48	0.37	0.38	0.38	0.47
Storage	25.78	28.04	20.57	21.48	25.16	27.22
Memory	1.82	1.63	1.68	0.85	1.80	1.56
Brand						
Teclast	0.06	0.23	0.13	0.29	0.06	0.23
Odna	0.04	0.20	0.01	0.05	0.04	0.19
Miui	0.06	0.24	0.00	0.00	0.06	0.24
Microsoft	0.04	0.20	0.02	0.14	0.04	0.19
Lenovo	0.09	0.28	0.16	0.37	0.10	0.29
Samsung	0.08	0.27	0.08	0.27	0.08	0.27
Apple	0.03	0.16	-	-	0.02	0.15
Huawei	0.03	0.16	-	-	0.03	0.15
Other	0.57	0.49	0.59	0.46	0.57	0.49
Seller Characteristics						
Large City	0.75	0.43	0.68	0.47	0.75	0.43
National Retailer	0.20	0.40	0.20	0.40	0.21	0.41
Age	29.74	18.74	30.64	22.47	30.24	19.30
Highly Rated	0.35	0.48	0.46	0.50	0.35	0.48
Mid Rated	0.27	0.45	0.22	0.00	0.27	0.44

Notes:

The final piece of data we collect is the return policy for the each seller that appears in the data. We navigate to each seller's page on Tmall in January of 2018 and manually input the seller's return policy along the three dimensions. If there is no stated return policy for any dimension, we assume that the seller agrees to the default policy. There are a couple of issues due to the fact that these data were collected after the demand and returns data. First, there are a number of sellers who no longer appear on Tmall. For these sellers, we assume that they have the default return policy in the baseline and perform a robustness check in the reduced form analysis using only sellers who are still on the platform. Second, a seller could have changed their return policy in between the collection dates. We assume that the return policy we collect is the return policy during our sample period. To the extent that there are a lot of changes in the return policies, then this would result in measurement error in our policy variable.

Table 3 summarizes the return policy data. Slightly over half the sellers do not have a return policy available (second to last row of the table), with about half of these being because they are no longer on the platform when we collect these data. If we assume those missing data are the default policy, then about 90% of sellers have the default policy, and this number is 81% if we ignore the all the missing sellers and 86% if we ignore the sellers no longer on the platform. Most of the sellers who vary from the default policy make their policy more lenient, as we see 7% of sellers set a more lenient policy than the default. This number is almost 19% if we ignore the all missing data and 14% if we ignore the sellers who exited. This is due to the fact that we assume that sellers that do not have an explicit policy on their website use the default policy, meaning that a higher proportion sellers will have the default policy. Of those who vary from the default, 5% accept returns after more than 7 days, 4% offer to pay for return shipping, and 1% require a reason for the return.

Table 4 displays how the return policies vary by seller and brand. Sellers located in a large city, national retailers and those who have mid-level ratings are less likely to have a lenient policy than the average seller, while highly rated sellers are more likely to have a lenient policy than the average seller. There is some variation by brand, as sellers who sell Teclast, Lenovo, and Odna tablets are more likely to offer a lenient policy than the average seller.

Table 3: Return Policy Summary

Variable	Missing=Default	Excluding All Missing	Exclude Exit
Days to Return			
7	0.94	0.90	0.88
15	0.05	0.07	0.08
30	0.00	0.01	0.01
Seller Pays Shipping	0.04	0.23	0.19
Reason Required	0.01	0.03	0.04
Default Policy	0.90	0.81	0.86
Lenient Policy	0.07	0.19	0.14
Missing Policy Total		0.51	
Missing Policy Exit		0.22	

Notes:

We now move to the left and middle panels of Table 2, which separates the average statistics across sellers who vary by the policy. Looking at the bottom section, sellers with a lenient policy are less likely to be located in a large city, equally likely to be a national retailer, and have been on Tmall longer than sellers with the default policy. Sellers who have a lenient policy are more likely to be highly rated, which is preliminary evidence of a possible signaling effect. In line with the statistics from Table 4, the middle section of Table 2 shows that sellers who have a lenient policy are more likely to sell Teclast and Lenovo compared to sellers that have the default policy. A higher percentage of sellers who have a lenient policy sell tablets that have the Android operating system and these sellers also sell tablets that have more storage. Other than brand, operating system, and storage, there are not large differences in the average tablet characteristics across sellers who vary in their policy. However, there are some significant differences across policies in terms of prices, quantities, and return rates, which are shown in the top section of Table 4. Sellers who have a lenient policy offer cheaper products and sell fewer tablets, which results in significantly less revenue per month. Sellers with a lenient policy have a lower return rate than sellers with the default, which is another piece of evidence in favor of a potential signaling effect. Finally, lenient sellers have a higher share of returns due to the quality of the product than default sellers.

Table 4: Policy by Brand and Seller

Variable	Days to Return 7	15	30	Seller Pays Shipping	Reason Required	Lenient	Default
Seller Characteristics							
Large City	0.94	0.06	0.00	0.05	0.03	0.07	0.90
National Retailer	0.87	0.10	0.01	0.03	0.06	0.07	0.84
Highly Rated	0.93	0.05	0.00	0.07	0.00	0.10	0.90
Mid Rated	0.94	0.04	0.01	0.03	0.00	0.06	0.91
Brand							
Teclast	0.85	0.08	0.00	0.31	0.00	0.23	0.77
Odna	0.88	0.12	0.00	0.12	0.00	0.12	0.88
Miui	0.92	0.08	0.00	0.00	0.00	0.00	0.92
Microsoft	0.97	0.00	0.03	0.00	0.00	0.03	0.97
Lenovo	0.88	0.12	0.00	0.06	0.00	0.12	0.82
Samsung	0.92	0.08	0.00	0.00	0.00	0.08	0.92
Apple	1.00	0.00	0.00	0.00	0.00	0.00	1.00
Huawei	1.00	0.00	0.00	0.00	0.00	0.00	1.00
Other	0.93	0.05	0.00	0.07	0.03	0.09	0.88

Notes:

3.1 Preliminary Evidence

In this section, we provide a descriptive analysis of the effect of return policies. To be consistent, we use the same that data we use to estimate the structural model, which is aggregated to the seller-month level, rather than the product-month level. The primary reason for the aggregation is because the return rates and return policies are at the seller level, but it is also because zeros are far less common at the seller level. We calculate the seller level demand by aggregating the quantity of sales across all products that a seller offers in a given month. The price and the product characteristics for a seller-month are formed by calculating the within-seller-month share weighted average of these variables. Therefore, each seller sells their weighted average product each month at the weighted average price, meaning that the terms ‘seller’ and ‘product’ are interchangeable hereafter.

We begin the preliminary analysis by providing evidence of the effect of the return policy on demand. First, we run the following regression model:

$$Log(Q_{st}) = x_{st}\beta^d + y_s\mu^d + \alpha^d Log(P_{st}) + \gamma^d \text{Lenient}_s + \epsilon_{st} \tag{7}$$

where Q_{st} is the quantity of tablets sold for seller s in month t. The product and seller characteristics, denoted x_{st} and y_s, respectively, are the same as those found in Table 2. The product characteristics for a seller vary across time due to new products being introduced, products being dropped, and changes in the within-seller-month relative sales, while the seller characteristics are time-invariant. While there is variation in the rating score across time, it is relatively small over our sample period, so we define the high and mid rating dummy variables based on the average rating score across months. Additionally, instead of using the continuous, time-varying seller age variable, we define a seller as new if they have been on Tmall less than 1 year as of the beginning of our sample, and old otherwise. The sales weighted average price is given by P_{st} and Lenient_s is a dummy variable indicating whether or not the seller has a lenient return policy.

The results of seven different specifications are presented in Table 6. The first four specifications include all sellers under the assumption that seller who left the platform have the default policy and the last two only include sellers that had not exited by the time we collected the policy data. We omit the seller and product characteristics in the interest of space. Results for an OLS regression are in column (1) and show that consumers are sensitive to prices and the magnitude suggests that the price elasticity is -1.42. Sellers with a lenient return policy have higher demand, but the coefficient is not significant. In column (2), we instrument for price with a variable measuring the relative share of listings that are tablets versus sell phones for rival sellers. This is similar to a BLP style instrument, but we are using a variable that is not included as a demand shifter. The results indicate that the endogeneity of prices resulted in a price coefficient biased towards zero, as the price elasticity is nearly twice as big at -3.11, and the lenient policy is still statistically equal to 0.

Table 5: Demand Regressions

Variable name	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Price	-1.42***	-3.11***	-3.09***	-3.00***	-2.57***	-2.57***	-2.47***
	(0.19)	(0.55)	(0.55)	(0.54)	(0.64)	(0.63)	(0.60)
Lenient Policy	0.09	-0.13	0.92**		0.03	1.00**	
	(0.28)	(0.30)	(0.45)		(0.34)	(0.50)	
Lenient Policy*Old			-1.84***			-1.83**	
			(0.60)			(0.77)	
Lenient Policy*New				0.91**			0.98**
				(0.44)			(0.49)
Constant	6.31***	12.59***	12.56***	12.13***	10.12***	10.24***	9.75***
	(1.17)	(2.29)	(2.27)	(2.22)	(2.61)	(2.61)	(2.47)
Obs	634	634	634	634	435	435	435
Prod Chars	Y	Y	Y	Y	Y	Y	Y
Sell Chars	Y	Y	Y	Y	Y	Y	Y
Sample	All	All	All	All	Policy	Policy	Policy
IV	N	Y	Y	Y	Y	Y	Y

Notes:

In column (3) we add interactions between a lenient return policy and the experience of the seller. The coefficient on the policy variable is now positive and significant and the magnitude implies that sellers with a lenient return policy who are new to the platform have about 1% higher demand. The significance disappears and even becomes negative for sellers who are more experienced. This suggests that the leniency of the policy is more important for new sellers than old sellers, which is likely due to the uncertainty related to the quality of these sellers. Column (4) presents a regression with only the lenient policy variable for new sellers included and the coefficients on price and the policy are similar to specification (3). For reasons we will describe below, this is how we define the lenient variable in our structural model. The final three columns are equivalent to columns (2)-(4), but here we exclude sellers who exited the platform. The results are similar, which can serve as justification for the assumption that sellers without an observable policy have the default policy.

In the next exercise, we examine the relationship between the return policy and the return rates. To do this, we run the following regression:

$$r_{st} = x_{st}\beta^r + y_s\mu^r + \alpha^r Log(P_{st}) + \gamma^r \text{Policy}_s + \epsilon_{st} \tag{8}$$

where r_{st} is the share of products returned for seller s in month t, and the other covariates are the same as in Equation 7. Again, we present results with ((1)-(4)) and without ((5)-(8)) the sellers who exited the market. Specifications (1) and (5) do not include any seller or product characteristics, specifications (2) and (6) include only seller characteristics, specifications (3), (4), (7) and (8) include both. Specifications (4) and (8) consider the effect of leniency for young sellers. The estimates of the leniency variable is always negative, indicating that sellers with a lenient policy have fewer returns, an implication of a possible signaling effect. However, the estimates are not significant in any of the specifications. This is in line with the idea that the signaling effect and the insurance effect of a lenient policy work in opposing directions in terms of their impact on returns. Prices have a positive relationship with returns, although it is only significant when we do not include seller characteristics.

Table 6: Return Frequency Regressions

Variable name	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Price	3.26**	3.49*	1.85	1.93	2.37	3.59	3.06	3.01
	(1.49)	(1.85)	(1.67)	(1.64)	(1.53)	(2.23)	(2.17)	(2.06)
Lenient Policy	-0.98	-1.04	-0.96		-1.88	-1.58	-1.07	
	(1.05)	(0.98)	(0.91)		(1.24)	(1.22)	(1.17)	
Lenient Policy*New				-1.16				-2.67
				(1.36)				(1.69)
Constant	-5.95	2.08	6.76	6.38	0.20	-1.59	-0.99	-0.98
	(8.18)	(8.10)	(6.93)	(6.79)	(8.43)	(9.75)	(8.84)	(8.44)
Obs	634	634	634	634	435	435	435	435
Prod Chars	N	Y	Y	Y	N	Y	Y	Y
Sell Chars	N	N	Y	Y	N	N	Y	Y
Sample	All	All	All	All	Policy	Policy	Policy	Policy
IV	Y	Y	Y	Y	Y	Y	Y	Y

Notes:

Finally, we examine the relationship between prices and returns with the following specification and present the result in Table 14. Again, specifications (1)-(3) use all the data and (4)-(6) use sellers who have not exited, Specifications (1) and (4) do not include any seller or product characteristics, specifications (2) and (5) include only seller characteristics, specifications (3) and (6) include both. Generally, there is a positive relationship between prices and return rates, which is suggestive evidence that returns are costly to sellers. However, the coefficient is not consistently significant when including product and seller characteristics, which is likely due to the fact that this reduce form representation doesn't account for the interaction between demand, return rates, return policies, and pricing decisions.

$$Log(P_{st}) = x_{st}\beta^p + y_s\mu^p + \iota r_{st} + \epsilon_{st} \tag{9}$$

Table 7: Price Regressions

Variable name	(1)	(2)	(3)	(4)	(5)	(6)
Return Rate	0.009**	0.001	0.004	0.013***	0.004	0.006**
	(0.004)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)
Constant	5.216***	3.861***	3.534***	5.184***	3.784***	3.465***
	(0.041)	(0.209)	(0.207)	(0.048)	(0.250)	(0.238)
Obs	634	634	634	435	435	435
Prod Chars	N	Y	Y	N	Y	Y
Sell Chars	N	N	Y	N	N	Y
Sample	All	All	All	Policy	Policy	Policy

Notes:

In summary, the data indicate that there is a positive relationship between a lenient policy and demand for new sellers, but we cannot separately identify the signaling and insurance effects. Additionally, there is evidence that returns are costly to sellers, but without a model of demand and supply, we cannot identify that cost. Therefore, we specify and estimate a structural model in the following sections, which also allows us to perform simulations quantifying the effects of different return policies.

4 Model

We follow the literature that estimates discrete choice models of demand in the spirit of Berry (1994), but we add the option for the buyer to return the product after receiving it. We assume that the return policy can impact demand through a signaling effect and an insurance effect.

The lifetime indirect utility for consumer i who buys and keeps (k) a tablet from seller s in month t is given by:

$$u_{ist}^{k} = x_{st}\boldsymbol{\beta}^{d} - \alpha p_{st} + \xi_{st} + \epsilon_{ist} \tag{10}$$

where x_{st} is a vector of product characteristics that are observable to the consumer before they make their purchase decision and p_{st} is the price. The term ξ_{st} represents a vertical component of utility, or 'seller quality', and ϵ_{ist} is horizontal tastes for the seller, or the 'match quality'. These two components are unknown to the consumer before they purchase the product. We assume that the consumer uses the information on the website to form their expectation of ξ_{st}, which we denote $\tilde{\xi}_{st}$. Therefore, once she purchases the product, she learns $\xi_{st} = \tilde{\xi}_{st} + \nu_{st}^{k}$, where ν_{st}^{k} represents the difference between the expectation of seller quality and true seller quality.[13] We assume that the match quality is made up of a component known to the consumer ex ante, η_{ist}, and an unknown stochastic shock learned after purchase, ε_{ist}^{k}, such that $\epsilon_{ist} = \eta_{ist} + \varepsilon_{ist}^{k}$. Therefore, the *ex post* indirect utility can be re-written as:

$$u_{ist}^{k} = \underbrace{\underbrace{x_{st}\boldsymbol{\beta}^{d} - \alpha p_{st} + \tilde{\xi}_{st}}_{\tilde{\delta}_{st}^{k}} + \nu_{st}^{k}}_{\delta_{st}^{k}} + \eta_{ist} + \varepsilon_{ist}^{k} \tag{11}$$

We assume that the consumer makes their purchase decision based on their belief about seller quality, so that the *ex ante* expected indirect utility is:

$$\tilde{u}_{ist}^{k} = E_{\varepsilon^{k}}\left[\tilde{\delta}_{st}^{k} + \eta_{ist} + \varepsilon_{ist}^{k}\right] \tag{12}$$

[13] We do not make assumptions about the distribution of ν^{k}, but do assume that consumers know how observables impact ξ_{st}.

where the belief the consumer forms is given by the following linear functional:

$$\tilde{\xi}_{st} = \underbrace{\mu_1^d \text{High Rating}_s + \mu_2^d \text{Med Rating}_s + \mu_3^d \text{National}_s + \mu_4^d \text{Large City}_s + \mu_5^d \text{New}_s}_{y_s \boldsymbol{\mu}^d} + \gamma^d Lenient_s + \phi_{st} \tag{13}$$

All of the variables in this function are dummy variables that are defined earlier. We note that the $Lenient_s$ variable is equal to one if the seller has a lenient policy **and** is new to the platform. This is because the reduce form evidence suggests that the policy is impactful for new sellers, but for old sellers the point estimate implies a negative relationship. This goes against the theory that the leniency of a policy should positively impact demand through both the signaling and insurance effect. We therefore assume that old sellers that have a lenient policy are equivalent to sellers who have the default policy. The parameter γ^d represents the signaling effect of the return policy.

The term ϕ_{st} is made up of two components. The first is a time-invariant component of seller quality that is unobserved to the econometrician but observed to the consumer before she makes her purchase decision, ϕ_s, and the second is an unobserved (to the econometrician) demand shock for seller s in month t, $\Delta\phi_{st}$.

If the consumer chooses the product, she learns u_{ist}^k and then can return the product return the product at a cost $C_{is}(x_{st}, y_s, Lenient_s)$ and receive a benefit $B_{ist}(x_{st}, y_s)$. The benefit of the return represents the value of the next best option, which could be buying again from the same seller, buying from a new seller, or taking the outside option. We allow this to be a function of product and seller characteristics of seller s, which serve as a proxies for shifters of these options. We assume that the cost of returning a product to seller s is also a function of seller and product characteristics. Additionally, the return costs vary across sellers based on their return policy. So the *ex post* utility of a return is given by:

$$u_{ist}^r = B_{ist}(x_{st}, Y_s) - C_{is}(x_{st}, Y_s, Lenient_s) \tag{14}$$

We cannot identify the impact of x_{st} and y_s on $B_{ist}(\cdot)$ and $C_{is}(\cdot)$, so we parameterize the *ex* *post* net return utility as:

$$u^r_{ist} = \underbrace{x_{st}\boldsymbol{\beta}^r + \gamma^r Lenient_s + y_s\boldsymbol{\mu}^r}_{\tilde{\delta}^r_{st}} + \nu^r_{st} + \varepsilon^r_{ist} \tag{15}$$

where ν^r_{st} is a seller/month level shifter to the return value or return cost and ε^r_{ist} is a return shock, both of which are realized after receiving the product. These shocks can represent variation in the return costs and/or return value across consumers and sellers. Therefore, the *ex ante* utility of a return can be expressed as:

$$\tilde{u}^r_{ist} = E_{\varepsilon^r}\left[\tilde{\delta}^r_{st} + \varepsilon^r_{ist}\right] \tag{16}$$

The impact of the return policy on the return cost, γ^r represents the level of insurance offered by different policies. That is, the lower the costs of returning, the higher the level of insurance. Note that we do not include the price in the return value but, in theory, we could include it to account for refunds and/or receiving a replacement when the consumer returns the product.[14]

The consumer will choose to return the product if:

$$\begin{aligned} u^r_{ist} &> u^k_{ist} \\ \tilde{\delta}^r_{st} + \nu^r_{st} + \varepsilon^r_{ist} &> \tilde{\delta}^k_{st} + \nu^k_{st} + \eta_{ist} + \varepsilon^k_{ist} \\ \underbrace{\tilde{\delta}^r_{st} + \underbrace{\nu^r_{st} - \nu^k_{st}}_{\nu_{st}}}_{\delta^r_{st}} + \varepsilon^r_{ist} &> \tilde{\delta}^k_{st} + \eta_{ist} + \varepsilon^k_{ist} \end{aligned} \tag{17}$$

We cannot separate the seller quality learned after purchase, ν^k_{st}, and the shifter of the return value, ν^r_{st}, so we combine it into one seller/month unobservable ν_{st} that is realized after purchase.

Before purchasing, the consumer does not know the realizations of ε^k_{ist}, or ε^r_{ist}, so when she is choosing which product to purchase, she must take expectations over these objects. The expected

[14] That is, we could identify a different price coefficient for returns and purchasing.

utility of purchasing from seller s from the consumer's perspective is:

$$E[u_{ist}] = E[\max\{u^k_{ist}, u^r_{ist}\}] \tag{18}$$

We assume that the return shock, ε^r_{ist}, and the realization of match quality shock, ε^k_{ist}, are i.i.d. extreme value random variables, so that the expected utility is the inclusive value of purchasing the product:

$$E[u_{ist}] = \log(\exp(\tilde{\delta}^k_{st} + \eta_{ist}) + \exp(\tilde{\delta}^r_{st})) \tag{19}$$

Finally, we add a pre-purchase shock that doesn't carry over to the return decision, meaning that it represents randomness in the purchase decision that is not related to the true realize utility.

$$\begin{aligned} \bar{u}_{ist} &= E[u_{ist}] + \upsilon_{ist} \\ &= \log(\exp(\tilde{\delta}^k_{st} + \eta_{ist}) + \exp(\tilde{\delta}^r_{st})) + \upsilon_{ist} \end{aligned} \tag{20}$$

We normalize the value of the outside option to 0, such that:

$$\bar{u}_{i0t} = 0 + \upsilon_{i0t} \tag{21}$$

The outside option can buying a tablet from a seller who is not on Tmall. We use the total number of tablets sold in China during our time period to measure the total market size and then subtract out our observed sales to get the sales of the outside option.[15] Assuming that υ_{ist} is also iid extreme value results in the probability that a consumer purchases from seller s, given by:

$$\begin{aligned} P_{ist} &= \frac{\exp(E[u_{ist}])}{1 + \sum_{s' \in S} \exp(E[u_{is't}])} \\ &= \frac{\exp(\tilde{\delta}^k_{st} + \eta_{ist}) + \exp(\tilde{\delta}^r_{st})}{1 + \sum_{s' \in S} \left(\exp(\tilde{\delta}^k_{s't} + \eta_{is't}) + \exp(\tilde{\delta}^r_{s't}) \right)} \end{aligned} \tag{22}$$

where S is the set of sellers. Once she receives the product, the probability that she returns the

[15] Data source.

product is given by:

$$R_{ist} = \frac{\exp(\delta^r_{st})}{\exp(\tilde{\delta}^k_{st} + \eta_{ist}) + \exp(\delta^r_{st})} \tag{23}$$

Therefore, the market share for seller s in month t is:

$$s_{st} = \int_{i \in \mathcal{I}} P_{ist} dF_i \tag{24}$$

where $\mathcal{I}$ is the set of all consumers. The share of returns is for seller s is:

$$r_{st} = \int_{i' \in \mathcal{I}_s} R_{ist} dF_{i'} \tag{25}$$

The integration in both equations is over consumers' unobserved (to the econometrician) utility that is realized before the purchase decision, η_{ist}. Here, $\mathcal{I}_s$ denotes the set of consumers who purchase product s, which demonstrates the importance of accounting for the selection of consumers who make a return decision for seller s. That is, the consumers who purchase are likely to have a high value of η_{ist}, meaning they are less likely to return the product. This is similar to the selection issues faced by search models (e.g., De los Santos et al. (2012)).

It is useful to define two additional share terms before moving to the supply side. First, is the the share of the market that purchases and keeps the good:

$$s^k_{st} = \int_{i \in \mathcal{I}} P_{ist}(1 - R_{ist}) dF_i \tag{26}$$

and second is the share of the market purchases and returns the good:

$$s^r_{st} = \int_{i \in \mathcal{I}} P_{ist} R_{ist} dF_i \tag{27}$$

The sum of these two represent the share of the market that purchases from seller s:

$$s_{st} = s^r_{st} + s^k_{st} \tag{28}$$

4.1 Supply Model

The profit of seller s is given by:

$$\pi_{st} = M_t \bigg(s^k_{st}(\mathbf{p}_t) \left(p_{st} - c^p_{st}\right) + s^r_{st}(\mathbf{p}_t) c^r_{st} \bigg) \tag{29}$$

where M_t is the market size, c^p_{st} is the marginal cost of the product for seller s in period t, and c^r_{st} is the net value to the seller for each returned product, which includes all possible costs and benefits of a return. $\mathbf{p}_t$ is a vector of prices for all sellers in month t. In estimation, we can make different assumptions about what is contained c^r_{st}, which impacts the interpretation of this value. In our baseline, we assume that the seller keeps p_{st} and doesn't pay c^p_{st} when a product is returned, such that $c^r_{st} = \kappa_{st} + p_{st}$. Under this scenario, κ_{st} represents the processing costs of a return, the non-recoupable marginal cost, the potential cost of replacement, and the potential lost revenue from a refund. We do not observe whether the consumer got a refund or was issued a replacement, so we are not able to separately estimate the costs associate with each scenario. The profit then can be re-written as:

$$\pi_{st} = M_t \bigg(\Big(s^k_{st}(\mathbf{p}_t) + s^r_{st}(\mathbf{p}_t) \Big) p_{st} + s^r_{st}(\mathbf{p}_t) \kappa_{st} - s^k_{st}(\mathbf{p}_t) c^p_{st} \bigg) \tag{30}$$

Taking the first order condition:

$$\begin{aligned} \frac{\delta \pi_{st}}{\delta p_{st}} &= 0 = M_t \left(\left(\frac{\delta s^k_{st}}{\delta p_{st}} + \frac{\delta s^r_{st}}{\delta p_{st}} \right) p_{st} + \left(s^k_{st} + s^r_{st} \right) + \frac{\delta s^r_{st}}{\delta p_{st}} \kappa_{st} - \frac{\delta s^k_{st}}{\delta p_{st}} c^p_{st} \right) \\ &\rightarrow p_{st} = \lambda^p_{st} c^p_{st} - \lambda^r_{st} \kappa_{st} - \chi_{st} \end{aligned} \tag{31}$$

where

$$
\begin{aligned}
\lambda^c_{st} &= \frac{\frac{\delta s^k_{st}}{\delta p_{st}}}{\frac{\delta s^k_{st}}{\delta p_{st}} + \frac{\delta s^r_{st}}{\delta p_{st}}} > 1 \quad MC \text{ mark-up} \\
\lambda^r_{st} &= \frac{\frac{\delta s^r_{st}}{\delta p_{st}}}{\frac{\delta s^k_{st}}{\delta p_{st}} + \frac{\delta s^r_{st}}{\delta p_{st}}} < 0 \quad RC \text{ mark-up} \\
\chi_{st} &= \frac{s^k_{st} + s^r_{st}}{\frac{\delta s^k_{st}}{\delta p_{st}} + \frac{\delta s^r_{st}}{\delta p_{st}}} < 0 \quad \text{Market power mark-up}
\end{aligned}
\tag{32}
$$

In a standard model without returns the first order condition is given by:

$$
p_{st} = c^p_{st} - \tilde{\chi}_{st} \tag{33}
$$

where $\tilde{\chi}_{st} = \frac{s_{st}}{\frac{\delta s_{st}}{\delta p_{st}}}$.

The pricing first order condition differs from a standard model without returns in a few important ways. First, there is a mark-up term on the marginal cost, which we denote λ^p_{st}. This comes from the assumptions about what happens with marginal cost and prices when there is a return. Specifically, because the model assumes that the seller does not pay c^p_{st} when the product is returned but still receives p_{st}, the mark-up on marginal cost is higher than 1. We note that this would be true even if only a fraction of the time they get to keep p_{st} or pay c^p_{st}.

Second, there is a mark-up on the return cost, λ^r_{st}. An interesting feature of this term is that this will get higher the more lenient the return policy due to a mechanism similar to advantageous selection in insurance markets. That is, the consumers who purchase the product after a marginal change in the return policy are the ones that are most likely to return it, meaning the expected cost of returns is going to increase for sellers. This means that the price will increase for all consumers, even the ones who are very unlikely to return the product.

Finally, the 'market power' mark-up term, χ_{st}, is similar to the standard model, as it is the total share of purchases for seller j divided by the derivative of the share of purchases with respect to p. However, in a model with returns, there is a higher share of purchases, all else equal, because consumers are protected against risk. More importantly, consumers will not be as price sensitive in their purchase decision because they are insured, so the denominator of χ_{st} is closer to 0 than than the denominator of $\tilde{\chi}_{st}$. This resembles a moral hazard mechanism, as consumers behave more risky (i.e., at higher prices) when they are protected against risk with the ability to return. Therefore, in a model with returns, the 'market power' mark up term is higher than it would be in a standard model.

The intuition of this is also similar to that of the market power that arises in a search model. That is, as the expected value of searching a product increases, then consumers are less sensitive to changes in search costs. But once they search, it is costly to choose another option. Together this gives firms additional market power. Here, the search costs are equivalent to the cost of purchasing the good and the value of searching is equivalent to the value of returns. As the value of returns grows, consumers are less price sensitive, but the consumers don't fully realize the increase in the value of the returns because returns are costly. We again note that this would be true even if the seller gets to keep p_{st} from only a portion of the returns.

In estimation, we parameterize marginal product cost as a function of product and seller characteristics:

$$c^p_{st} = x_{st}\boldsymbol{\beta}^c + y_s\boldsymbol{\mu}^c + \omega_{st}$$

In theory, we could also parameterize κ_{st} as a function of observables, but it has been difficult to identify these parameters in practice. Therefore, we assume a homogenous return cost for all sellers/products such that $\kappa_{st} = \kappa \quad \forall s$.

5 Estimation

We utilize the data on purchase shares and the return rates in order to identify the model. The parameters to estimate are separated into the preference/learning parameters, $\theta^d = \{\boldsymbol{\beta}^d, \alpha, \boldsymbol{\mu}^d, \gamma^d\}$, the return parameters, $\theta^r = \{\boldsymbol{\beta}^r, \boldsymbol{\mu}^r, \gamma^r\}$, and the cost parameters, $\theta^c = \{\boldsymbol{\beta}^c, \boldsymbol{\mu}^c, \kappa\}$.

We separate the estimation of the demand and supply parameters. On the demand side, the estimation follows that of Berry (1994) and Berry et al. (1995), in that we 'invert' the shares in order to solve for the mean purchase and return values, $\tilde{\delta}^k_{st}$ and δ^r_{st}. While our model is homogenous, the inversion requires a contraction as in Berry et al. (1995) because we need to integrate over the consumer unobservable η and because we have to invert both the purchase share and the return share.

Specifically, we have an inner and outer contraction. Conditional on a vector of $\tilde{\delta}^k_{st}$, we can use a seller month level contraction to solve for the δ^r_{st} such that the return rates predicted by the model equal the observed return rates. Then the outer loop iterates of $\tilde{\delta}^k_{st}$ until the purchase shares predicted by the model equal the predicted shares. Within the loops, we integrate over the distribution of η using quadrature integration in order to calculate the shares.

Once we have the values of $\tilde{\delta}^k_{st}$ and δ^r_{st}, we can run simple regression models in order to estimate the parameters. The consumption preference regression model is

$$\tilde{\delta}^k_{st} = x_{st}\boldsymbol{\beta}^d - \alpha p_{st} + \underbrace{y_s\boldsymbol{\mu}^d + \gamma^s Lenient_s + \phi_s}_{\text{Seller FE}} + \Delta\phi_{st} \tag{34}$$

which we estimate via 2SLS using the measure of tablet versus cell phone sales of other sellers as the instrument for price. In practice, we estimate seller fixed effects and then regress the fixed effects on y_s and $Lenient_s$ in order to estimate the learning parameters. The return preference regression model is:

$$\delta^r_{st} = x_{st}\boldsymbol{\beta}^r + \underbrace{y_s\boldsymbol{\mu}^r + \gamma^r Lenient_s + \nu_s}_{\text{Seller FE}} + \Delta\nu_{st} \tag{35}$$

which we estimate via OLS with seller fixed effects and, again, we regress the fixed effects on y_s and $Lenient_s$ in order to estimate the effect of static seller characteristics.

Once we have the demand parameters, we estimate the supply side via OLS of the first order condition equations:

$$\begin{aligned} \frac{1}{\lambda^p_{st}} p_{st} + \frac{1}{\lambda^p_{st}} \chi_{st} &= c^p_{st} - \frac{\lambda^r_{st}}{\lambda^p_{st}}\kappa \\ &= x_{st}\boldsymbol{\beta}^c + y_s\boldsymbol{\mu}^c - \frac{\lambda^r_{st}}{\lambda^p_{st}}\kappa + \omega_{st} \end{aligned} \tag{36}$$

5.1 Identification

The complication in identifying the demand parameters is the fact that many of the same covariates enter both the utility of purchasing and the utility of making a return. The key to identifying these parameters is that we observe both the purchase share and the return rate. Intuitively, using these two different pieces of data, we are able to solve for δ^r_{st} separately from $\tilde{\delta}^k_{st}$, so we can use the covariation in δ^r_{st} and $Lenient_s$, for example, to identify γ^r and the covariation in $\tilde{\delta}^k_{st}$ and $Lenient_s$ to identify γ^d. This demonstrates how we separately identify the signaling effect from the insurance effect of the return policy.

The primary identifying assumptions are that the covariates (besides price) are conditionally independent of the time varying seller-level shocks in $\{\Delta\phi_{st}, \Delta\nu_{st}, \omega_{st}\}$ and the consumer-level shocks in $\{\varepsilon^k_{ist}, \varepsilon^r_{ist}, \eta_{ist}, \upsilon_{ist}\}$. In addition, in order to identify the impacts of the seller-level, time-invariant characteristics, including the return policy, we make the additional assumptions that y_s and $Lenient_s$ are conditionally independent of $\{\phi_s, \nu_s\}$.

Recall that ϕ_s is the unobserved (to the econometrician) variation in purchase shares for seller s. While the exogeneity of the seller's location, age, and national presence is straightforward, one may argue that a seller's ratings and return policy are a function of ϕ_s. To provide justification for this assumption in terms of ratings, we note that we have estimated the model using the time-varying ratings and the results do not change significantly. Regarding the return policy, the assumption is violated if sellers set their return policy with knowledge of ϕ_s. We posit that the return policy is likely set before the realization of the seller level variation in purchases that is not accounted for through the product characteristics and the other seller characteristics. We also point out that ϕ_s is known to the consumer at the purchase stage, so sellers do not set the policy in order to provide a signal about the value of ϕ_s in our model.

The seller level unobservable in the return decision, ν_s, represents both the realized seller quality that is not learned through the information signals and any seller-level variation in the benefits/cost of returns. Similar to the purchase stage, the exogeneity of seller characteristics besides the return policy is easily argued and using time varying ratings does not change the results significantly. The assumption regarding the policy implies that the seller sets their return policy without considering

ν_s, which is comprised of ν_s^k and ν_s^r. The independence of ν_s^k comes from the assumption that consumers are rational in terms of how the policy impacts their belief about seller quality. In other words, the policy does not provide any new information when the consumer reaches the return stage, meaning there is no signaling motive to set the policy based on the value of ν_s^k. Regarding possible correlation between the policy and ν_s^r, we again argue that sellers are likely unaware of the variation in returns not accounted for through product and seller characteristics when setting their return policy.

5.2 Estimates

The first column of the left panel and the right panel of Table 8 displays the estimates of the demand model, with standard errors in parentheses. Due to linearity of the model, the calculation of the standard errors is standard, though we do account for the distribution of the seller fixed effect estimates when calculating the standard errors for the seller characteristics.

The price coefficient is -0.024, resulting in an average price elasticity of demand (i.e., purchasing and keeping the product) of nearly -4. The result is not statistically significant, but is similar in magnitude to the reduced form significant estimate from Section 3.1. As indicated in Table 9b, less than 1% of sellers have inelastic demand under this price coefficient. The preferences for product characteristics all have their predicted sign, but most of them are imprecisely estimated. The coefficients on the seller characteristics, which represent the impact of information signals, also have signs that are in line with theory, with a few of them being statistically significant. Specifically, a consumer's belief is higher for sellers with a national offline presence and high or medium ratings. Using the price coefficient, we calculate that a consumer's belief about the quality of a seller/product is between \$43 and \$48 higher than a seller with low ratings. The impact of the return policy on consumer beliefs is positive, but it is not significant. The magnitude of the coefficient suggests a similar impact as ratings, as a consumer's belief is about \$47 higher for new sellers with a lenient return policy. The fact that it is not significant suggests that consumers use the other signals available to infer quality of a seller.

Table 8: Demand Model Estimates

(a) Main Estimates

	Purchase Preferences			Return Preferences
	Full Model	No Returns		Full Model
Price Sensitivity			Product Characteristics	
Price	-0.024 (0.016)	-0.013 (0.011)	Intercept	-9.504 (0.965)
			Screen Size	-0.023 (0.094)
Product Characteristics			Cellular Network	-0.254 (0.226)
Intercept	-10.814 (1.843)	-6.950 (1.276)	Android	-0.983 (0.653)
Screen Size	0.858 (0.666)	0.403 (0.461)	Windows	-0.485 (0.679)
Cellular Network	0.684 (0.610)	0.251 (0.423)	Product Age	0.130 (0.192)
Android	-3.382 (2.438)	-2.637 (1.688)	Storage	0.001 (0.004)
Windows	-2.446 (2.097)	-2.098 (1.452)	Memory	-0.047 (0.074)
Product Age	-0.237 (0.372)	-0.200 (0.258)	Top Brand	0.187 (0.353)
Storage	0.111 (0.072)	0.061 (0.050)	September	1.225 (0.116)
Memory	0.244 (0.194)	0.124 (0.135)	October	-0.567 (0.113)
Top Brand	1.080 (0.699)	0.699 (0.484)	November	-0.580 (0.111)
September	1.162 (0.396)	1.305 (0.274)		
October	-0.984 (0.538)	-0.710 (0.373)	Seller Characteristics	
November	-0.405 (0.328)	-0.276 (0.227)	Lenient	1.243 (0.518)
			Large City	0.813 (0.207)
Seller Characteristics			National	1.310 (0.214)
Lenient	1.121 (1.088)	1.119 (0.539)	New Seller	-0.203 (0.208)
Large City	0.666 (0.435)	0.630 (0.215)	High Rating	-0.974 (0.219)
National	1.494 (0.450)	1.195 (0.223)	Medium Rating	0.106 (0.210)
New Seller	0.021 (0.437)	0.029 (0.217)		
High Rating	1.038 (0.460)	0.401 (0.228)		
Medium Rating	1.150 (0.440)	0.746 (0.218)		

(b) Summary

	Full Model	No Returns
Keep Elasticity	-3.902	-2.526
Share Inelastic	0.008	0.050
Return Elasticity	0.030	-
Conditional Return Elasticity	0.021	-
Return Cost	22.182	-

Notes: .

The estimates of the net return value appear in the right panel of Table 8. We interpret the coefficients as the extent to which the net value of returning a product is impacted by the given variable, where a positive estimate means that the covariate increases the benefit and/or reduces the cost of returning. Similar to the purchase stage, many of the estimates of the impact of product characteristics are imprecisely estimated. Many of the seller characteristics, however, have a significant impact on the net return value. Sellers located in a large city ($34) and who have an national offline presence ($55) have a higher return value, which suggests that consumer return costs might be lower for these types of sellers. Sellers with a high rating have a net-return value that is $41 lower than than sellers with a low rating, which suggests that products purchased from theses sellers are likely better than what the consumers could get if they request a return.

We find that a leniency policy increases a consumer's return value by a statistically significant 1.243, which suggests that consumer return costs are $52 lower for sellers with a lenient policy. This represents variation in the monetary costs of dealing with a return and the net value of receiving a replacement/refund from the seller. The magnitude of the estimate suggests that the policy is an important determinant of returns, as the average consumer return costs across all sellers is about $22.[16] Finally, Table 9b shows that the average unconditional return elasticity is small (0.03) and that the conditional (on purchase) elasticity is 0.021. The significant difference between the conditional and unconditional return elasticities highlight the importance of accounting for unobserved preferences that carry over from the purchase stage to the return stage (i.e., η).

The supply estimates are presented in the first column 9, with standard errors in parentheses. The signs of the effects of product characteristics on marginal costs mostly line up with priors, with the only possible exception being the age of the product. However, this coefficient is not statistically significant. The seller characteristics we include have positive and significant effects on costs, except for the dummy variable indicating that the seller is in a large city. Given these estimates we calculate the product cost for the median seller, which is about $136.

[16]We calculate this by taking the difference between the value of purchasing and the value of returning the product for each seller, assuming that this difference represents the cost of a return.

Table 9: Supply Model Estimates

(a) Main Estimates

	Marginal Cost	
	Full Model	No Returns
Returns		
Return Cost	259.09 (46.77)	-
Prooduct Characteristics		
Intercept	-392.87 (59.97)	-363.02 (53.09)
Screen Size	46.42 (5.37)	43.01 (4.75)
Cellular Network	102.66 (13.75)	93.40 (12.12)
Android	-154.25 (32.07)	-144.95 (28.30)
Windows	-216.04 (34.94)	-189.06 (30.93)
Product Age	17.40 (10.00)	7.76 (8.83)
Storage	5.50 (0.26)	4.83 (0.23)
Memory	15.28 (3.61)	14.27 (3.20)
Top Brand	45.37 (9.66)	43.82 (8.55)
Seller Characteristics		
IV: input	1.23 (0.40)	1.17 (0.35)
Large City	4.56 (10.56)	6.83 (9.32)
National	25.39 (11.00)	24.71 (9.63)
New Seller	19.91 (10.25)	13.56 (9.07)
High Rating	79.83 (11.31)	70.94 (9.40)
Medium Rating	28.79 (10.74)	25.41 (9.07)

(b) Summary

	Full Model	No Returns
Product Cost	136.51	118.09
Expected Return Cost	17.41	-
Mark Up	0.34	0.58

Notes: .

The estimate of the cost of a return to a seller is about \$260, which implies an expected return cost for each transaction of over \$17 for the median seller. Recall that the return costs includes the processing cost, the non-recouped marginal cost, the cost of a replacement, and the lost revenue if there is a refund issued. Therefore, one could interpret the magnitude of this cost as the seller giving the consumer a replacement (\$136) and the non-recouped marginal cost of the old tablet plus the processing cost totaling \$124. Another interpretation is that the cost represents the lost revenue of a refund (\$195), and processing cost plus the non-recouped marginal cost totaling \$65. Under either interpretation, the cost of returns to sellers is large.

The second column of the left panel of Table 8 and the second column of 9 display the estimates of the model if we were to ignore returns. In this case, it becomes a homogenous logit model that has a linear estimating equation. The price coefficient is estimated to be about half the coefficient of the full model, resulting the the elasticity for the median seller to be about -2.5. The parameters on most of the other covariates are also smaller (in absolute value) than the full model. This comes from the fact that the model without returns interprets the lack of response to a given covariate as low preferences for them, when it is atually due to the fact that consumers are not tied to the product once they buy it. For a given price change, for example, the consumer will not be as responsive in a world with returns because they know they can always return the product afterwards. This bias in the price coefficient leads to 5% of sellers/months having inelastic overall demand. Another interesting estimate of the no-returns model is that of the lenient variable. If we include this as a utility shifter in the no-returns model, we get a significant and positive effect. However, this the equivalent to what we did in Section 3.1, as it is confounding the signaling and insurance effects.

On the supply side, the bias in the price coefficient leads to lower estimates of marginal cost and higher mark-ups (0.58 for the median seller). This demonstrates the importance of accounting for returns when analyzing competition in both online and offline retail markets.

6 Simulations

In this section, we examine the effects of return policies in two ways. First, we quantify the relative role of the signaling and insurance effects of the return policy in this market, holding prices fixed. Second, we analyze the impact of the platform making a uniform return policy rule, with a focus on the role of the leniency of the rule.

6.1 Signaling versus Insurance

We start by setting a baseline where neither of the signaling nor the insurance effects are in the market. In practice, we do this by setting $Lenient_s = 0$ and $\delta^r_{st} = -\infty$ for all sellers. By setting the policy variable to zero, we make all sellers on the platform identical, *ex ante* and setting the value of a return to $-\infty$ is equivalent to making the cost of a return infinite. We present outcomes under this scenario in the second column of Table 10, where the observed outcomes are in the first column. We focus on consumer welfare and quantity sold by sellers, but also include the return rates and consumer return costs. Consumer welfare is calculated as the expected indirect utility of purchasing in this market, which in this case is given by:

$$CS = \int_i \log \left(\sum_{s,t} \exp(\delta_{st}^k - \gamma^k Lenient_s + \eta_{ist}) dF_i \right)$$

The results show that consumer surplus decreases by about 5% when the policy is removed as demand moves to the outside option and the average seller sells 7 fewer tablets. The lenient sellers are the ones that are hit the hardest, as sales drop from 264 to 105 tablets per month, demonstrating that these sellers rely on the policy. That is, they no longer are separated from default sellers with the policy signal

Table 10: Signaling versus Insurance

		Observed	Ins=N, Sig=N	Ins=N, Sig=Y	Ins=Y, Sig=N
Consumer Surplus ($M)		206.84	195.58	197.42	205.21
Surplus Per Consumer ($)		96.21	90.99	91.86	95.44
Consumer Return Cost ($K)		76.692	-	-	76.494
	Lenient	264.04	105.32	272.81	114.69
Quantity	Default	288.44	286.67	281.83	292.92
	All	287.59	280.38	281.52	286.74
	Lenient	0.081	-	-	0.105
Return Rate	Default	0.115	-	-	0.116
	All	0.114	-	-	0.116

Notes:

In the third column of the table, we add the signal back into the consumer's decision to purchase the product and present the outcomes. Moving from column two to column three features an increase in consumer surplus of just below 1%. Sales for lenient sellers increase to a level that is even higher than the observed outcomes, while default sellers see a decrease in sales because the lenient sellers enjoy the benefit of signaling without the competition from the insurance effect. When adding the insurance effect (column 4), the consumer surplus increases by about 6% and all sellers have an increase number of sales. However, this comes with an increase in the return rates, specifically for lenient sellers. The return rate for lenient sellers under this scenario is 10.5%, while in the observed scenario is 8.1%, highlighting the role of signaling in reducing returns.

Overall, these results suggest that the insurance value of the return policy makes up around 80% of the value of return policies to consumers. One caveat to this is that signaling directly impacts only lenient sellers, whereas insurance affects all sellers. Therefore, the relative impact of signaling is going to be mechanically lower. This can be seen through the relative impact of the two effects on sales, as signaling plays a significant role in determining the *distribution* of demand across different seller types, but does not impact overall demand a great degree.

6.2 Uniform Policy

Next, we examine the impacts of the platform making a uniform policy return policy rule, paying particular attention to how the leniency of the rule affects outcomes. Amazon requires sellers who subscribe to their Fulfilled by Amazon (FBA) program to have a uniform policy rule. Even sellers who do not use the FBA program are required to have a certain level of leniency.[17] Additionally, in the spring of 2014, the Chinese government implemented a law that required a certain level of leniency. We view our simulations an input into debates about making such policies.

To do this, we set $Lenient_s = 0$ in the purchase utility to remove signaling and adjust the return utility, δ^r_{st}, such that the consumer cost of a return is the same for all sellers. We start by forming a baseline where there are no returns, or $\delta^r_{st} = -\infty$, and then show how outcomes change when the policy is set to the default policy for all sellers.

[17] See `https://www.repricer.com/blog/amazons-new-returns-policy/`

Table 11: Outcomes Under Uniform Ban and Default Return Policy without Signaling

			Optimal Prices		Fixed Prices	
		Observed	Ins=N, Sig=N	Ins=Def, Sig=N	Ins=N, Sig=N	Ins=Def, Sig=N
Profit ($M)		2.605	2.587	2.600	3.132	2.601
Consumer Surplus ($M)		206.84	232.79	206.18	195.58	205.07
Surplus Per Consumer ($)		96.21	108.09	95.89	90.99	95.37
Total ($M)		209.45	235.38	208.78	198.71	207.67
Firm Return Cost ($K)		616.556	-	612.000	-	619.277
Consumer Return Cost ($K)		76.692	-	75.667	-	75.772
Total ($K)		693.247	-	687.667		695.049
Elasticity		-5.059	-5.572	-5.073	-5.525	-5.063
Mark up		0.336	0.307	0.328	0.382	0.338
MC Pass Through		1.098	1	1.091	1	1.097
Market Power		50.172	45.667	49.863	45.585	50.129
RC Pass Through		0.084	-	0.079	-	0.082
	Lenient	13,906	4,811	6,059	6,364	5,734
Profit ($)	Default	16,526	16,735	16,776	20,245	16,791
	All	16,435	16,321	16,404	19,763	16,407
	Lenient	227.79	221.27	219.36	227.79	227.79
Prices ($)	Default	247.87	238.10	246.67	247.87	247.87
	All	247.18	237.51	245.72	247.18	247.18
	Lenient	0.081	-	0.021	-	0.027
Return Rate	Default	0.115	-	0.114	-	0.116
	All	0.114	-	0.113	-	0.115

Notes:

The fourth and firth column of Table 11 shows outcomes under a return ban and a uniform default policy assuming that sellers cannot adjust prices. We expand the outcomes from the previous exercise to now include more supply side information such as profit (total, means), prices (means), elasticities (medians), and mark-ups (medians) and we break down mark-ups into the three categories highlighted in Equation 32. The results indicate that, as we make the uniform policy more lenient, consumer surplus increases by 4.8%, which can represent the value consumers place on a more lenient return policy, all else equal. Total firm profit decreases by 17% overall, with lenient sellers losing 10% of their profit and default sellers losing more than 17%, which can represent the costs associated with more leniet return policies. The overall loss in firm profit is due to the increase in costs (from returns) outweighing any gains from the increase in demand and default sellers are impacted more by the return costs as they have an 11.6% return rate versus 2.7%

for lenient sellers.

When we allow firms to adjust prices in response to the change in leniency, the story changes. Moving from column 2 to column 3 shows that consumer surplus falls by 11.3% when the policy becomes more lenient and firm profit increase by less than 1%. The reduction in consumer surplus from the fact that sellers increase prices by 3.4%. Firms are able to recoup the costs of returns by charging higher prices. The increase in prices comes from sellers passing through return costs onto consumers and from an increase in mark-ups due to market power. The average return rate with the lenient policy is 11.3%, suggesting that an increase in the return seller cost by $1 increases the expected return cost by 11 cents. The results indicate that the median seller marks up return costs by 0.08, translating to a (0.08/.011) 73% mark-up on the expected return cost. The mark-up due to market power increases by 9% with the more lenient policy. In total, the mark-up for the median seller increases by 6.8%. Interestingly, the average lenient seller actually decreases their price, which is likely due to the increased competition from the default sellers. Though, lenient sellers benefit more from a more lenient uniform policy, as they get more sales, but don't increase their return cost as much as default sellers.

While Table 11 focused on the return leniency, there is also the fact that implementing a uniform policy removes any signaling effect. In order to examine this, we repeat the same exercise in columns 2 and 3 of Table 11, but we now allow for signal related to the return policy to still be visible. The results are in columns four and five of Table 12, with the results when the signal is removed repeated in columns 2 and 3. Most of the results are qualitatively similar, but dampened to a degree because the signal relieves some of the uncertainty. This suggests that signaling and insurance effects of a uniform policy interact, but only to a small degree. Another interpretation is that, if firms are able to signal their quality in other ways in response the platform forcing a uniform more lenient policy, then the consumers (sellers) would be slightly better (worse) off.

Table 12: Outcomes Under Uniform Ban and Default Return Policy with Signaling

			Optimal Prices		Optimal Prices	
		Observed	Ins=N, Sig=N	Ins=Def, Sig=N	Ins=N, Sig=Y	Ins=Def, Sig=Y
Profit ($M)		2.605	2.587	2.600	2.585	2.599
Consumer Surplus ($M)		206.84	232.79	206.18	234.62	208.56
Surplus Per Consumer ($)		96.21	108.09	95.89	108.96	97.00
Total ($M)		209.45	235.38	208.78	237.21	211.16
Firm Return Cost ($K)		616.556	-	612.000	-	598.060
Consumer Return Cost ($K)		76.692	-	75.667	-	75.024
Total ($K)		693.247	-	687.667	-	673.084
Elasticity		-5.059	-5.572	-5.073	-5.579	-5.082
Mark up		0.336	0.307	0.328	0.307	0.329
MC Pass Through		1.098	1	1.091	1	1.090
Market Power		50.172	45.667	49.863	45.460	49.797
RC Pass Through		0.084	-	0.079	-	0.078
	Lenient	13,906	4,811	6,059	12,061	14,839
Profit ($)	Default	16,526	16,735	16,776	16,460	16,456
	All	16,435	16,321	16,404	16,308	16,400
	Lenient	227.79	221.27	219.36	216.08	218.61
Prices ($)	Default	247.87	238.10	246.67	238.07	246.59
	All	247.18	237.51	245.72	237.31	245.62
	Lenient	0.081	-	0.021	-	0.020
Return Rate	Default	0.115	-	0.114	-	0.113
	All	0.114	-	0.113	-	0.110

Notes:

Finally, we can use these exercises to quantify the impact of implementing a uniform default policy rule in our market. To do so, we compare the outcomes under the default policy in column three of Table 11 to the observed outcomes in column one in in Table 11. Moving to a uniform default policy removes the signaling effect and makes the policy more strict, as lenient sellers are now forced to have the default policy. This policy change results in consumers being worse off despite the fact that they are facing lower prices. Sellers are also worse off, as the more strict policy results in lower mark-ups and that outweighs the reduction in return costs. The lower mark-up comes from both lower mark-ups of return costs and a loss in market power.

If we allow for signaling to remain despite the uniform policy, consumer surplus now increases from the policy, as consumers face lower prices without losing the signal. Overall, sellers are still losing from the policy, but lenient sellers are actually better off. This is because they are still able to signal their quality, increasing their demand, but they have lower return costs with the more strict policy.

6.3 Discussion

Our results suggest that lenient return polices may have adverse effects on consumers through firm price responses and benefit sellers. Given anecdotal evidence of the value of easy returns to consumers and the cost of such returns to sellers, it is counterintuitive that we find consumers are worse off with lenient return policies and sellers are better off. However, it is important to know that our findings do not imply that consumer do not value easy returns. In fact, we do find that consumers place a significant value on sellers with a more lenient policy. But, those lenient policies come with a cost in the form of higher prices, and our results suggest that these price effects outweigh the benefits. We suspect that, when consumers are asked about return policies in surveys, they are not fully considering the possible price effects. Our results suggest that if consumers were asked "would you be willing to pay \$X for a lenient return policy", where \$X is our estimated increase in prices, the consumers would answer "no". We are not aware of any survey or other evidence that quantifies the WTP for returns.

On the firm side, we find that, while returns are very costly, sellers can benefit from a lenient policy rule because it gives them more pricing power. This counters anecdotal evidence that returns can be detrimental to sellers. We note that our results suggest that lenient policies help the average or median seller, but that may not be true for all sellers. Indeed, if we attempt to move to an even more lenient policy than the default policy, there are some sellers that want to set an infinite price, as the returns are too costly for them to make a positive profit from participating in the market.

It is also important to keep in mind that uncertainty doesn't play a huge role in this market, with less than 10% of tablets being returned. Therefore, there is not a ton of risk on the consumers' part. This can be likened to other insurance markets, where it doesn't make sense for many people to be insured, given their risk and the cost of insurance. In other retail markets that feature more uncertainty, for example apparel, it might be the case that the value of returns outweighs the cost, in terms of price increases. But even in this case, our analysis suggests that the benefit of a lenient return policy to consumers is dampened by the market power that results.

Something that our model does not capture is any dynamic effects of lenient return policies, on both the demand and supply side. Specifically, it could be the case that lenient return policies have some market expansion effects by attracting new customers and/or retaining current customers. Additionally, setting more lenient policies may impact the market structure by firms entering/exiting. In this sense, our analysis serves as a 'short-run' examination of return policies.

7 Conclusion

Our analysis examines the impact of return policies on the market for tablets on Alibaba's Tmall. On the demand side, return policies serve more as a form of insurance than as a signal of seller quality. On the supply side, returns are costly to sellers and provide them with some market power. We demonstrate that policies which feature uniform return rules can have adverse effects on consumers as the rules become more lenient, something that is due to this market power. While this must be caveated by the fact that this is a single product category and a single platform, it suggests that the pricing effects of loosened return policies should be considered when designing or regulating platforms. We believe that our model could be extended to include endogenous return policies, firm entry and exit, and consumer dynamics in order to further study the impact of return polices that we see in practice. Further, with a more extensive data set, one could study the heterogeneity in the effects of policies across product categories and, therefore, analyze how platforms can design polices that may or may not vary by category.

References

Anderson, E. T., K. Hansen, and D. Simester (2009). The option value of returns: Theory and empirical evidence. *Marketing Science 28*(3), 405–423.

Baye, M. R. and J. Prince (2020). The economics of digital platforms: A guide for regulators. *The Global Antitrust Institute Report on the Digital Economy 34*.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841–890.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.

Che, Y.-K. (1996). Customer return policies for experience goods. *The Journal of Industrial Economics*, 17–24.

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

Courty, P. and L. Hao (2000). Sequential screening. *The Review of Economic Studies 67*(4), 697–717.

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American economic review 102*(6), 2955–80.

Dinerstein, M., L. Einav, J. Levin, and N. Sundaresan (2018). Consumer price search and platform design in internet commerce. *American Economic Review 108*(7), 1820–59.

Einav, L., D. Knoepfle, J. Levin, and N. Sundaresan (2014). Sales taxes and internet commerce. *American Economic Review 104*(1), 1–26.

Ellison, G. and S. F. Ellison (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica 77*(2), 427–452.

Farronato, C., A. Fradkin, B. Larsen, and E. Brynjolfsson (2020). Consumer protection in an online world: An analysis of occupational licensing. Technical report, National Bureau of Economic Research.

Guan, Y. (2020). Right to return goods without reason in chinese e-commerce as the references for indonesian laws. *International Journal of Research and Analytical Reviews 7*(4), 653–658.

Gutierrez Gallardo, G. (2021). The welfare consequences of regulating amazon. *Available at SSRN 3965566*.

Hilger, J., G. Rafert, and S. Villas-Boas (2011). Expert opinion and the demand for experience goods: an experimental approach in the retail wine market. *Review of Economics and Statistics 93*(4), 1289–1296.

Hui, X., M. Saeedi, Z. Shen, and N. Sundaresan (2016). Reputation and regulations: Evidence from ebay. *Management Science 62*(12), 3604–3616.

Hui, X., M. Saeedi, G. Spagnolo, and S. Tadelis (2021). Raising the bar: Certification thresholds and market outcomes. *American Economic Journal: Microeconomics*.

Inderst, R. and M. Ottaviani (2013). Sales talk, cancellation terms and the role of consumer protection. *Review of Economic Studies 80*(3), 1002–1026.

Inderst, R. and G. Tirosh (2015). Refunds and returns in a vertically differentiated industry. *International Journal of Industrial Organization 38*, 44–51.

Lewis, G. (2011). Asymmetric information, adverse selection and online disclosure: The case of ebay motors. *American Economic Review 101*(4), 1535–46.

Moorthy, S. and K. Srinivasan (1995). Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Science 14*(4), 442–466.

Newberry, P. and X. Zhou (2019). Heterogeneous effects of online reputation for local and national retailers. *International Economic Review 60*(4), 1565–1587.

Newberry, P. W. (2016). An empirical study of observational learning. *The RAND Journal of Economics 47*(2), 394–432.

Petrikaitė, V. (2018). A search model of costly product returns. *International Journal of Industrial Organization 58*, 236–251.

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

Sahoo, N., C. Dellarocas, and S. Srinivasan (2018). The impact of online product reviews on product returns. *Information Systems Research 29*(3), 723–738.

Sorensen, A. T. (2007). Bestseller lists and product variety. *The journal of industrial economics 55*(4), 715–738.

Vo, M. P. (2022). Product return policies, pricing, and consumer welfare. *Pricing, and Consumer Welfare (May 8, 2022)*.

8 Additional Tables and Figures

Table 13: Outcomes Under Uniform Strict & Default Return Policy without Signaling

			Optimal Prices		Fixed Prices	
		Observed	Ins=Strict, Sig=N	Ins=Def, Sig=N	Ins=N, Sig=N	Ins=Def, Sig=N
Profit ($M)		2.605	2.590	2.600	2.965	2.601
Consumer Surplus ($M)		206.84	227.07	206.18	198.46	205.07
Surplus Per Consumer ($)		96.21	105.46	95.89	92.32	95.37
Total ($M)		209.45	229.66	208.78	201.43	207.67
Firm Return Cost ($K)		616.556	143.667	612.000	195.597	619.277
Consumer Return Cost ($K)		76.692	18.258	75.667	23.429	75.772
Total ($K)		693.247	161.926	687.667	219.026	695.049
Elasticity		-5.059	-5.391	-5.073	-5.372	-5.063
Mark up		0.336	0.317	0.328	0.376	0.338
MC Pass Through		1.098	1.022	1.091	1.029	1.097
Market Power		50.172	46.738	49.863	46.911	50.129
RC Pass Through		0.084	0.021	0.079	0.027	0.082
	Lenient	13,906	5,067	6,059	6,177	5,734
Profit ($)	Default	16,526	16,749	16,776	19,156	16,791
	All	16,435	16,343	16,404	18,706	16,407
	Lenient	227.79	222.42	219.36	227.79	227.79
Prices ($)	Default	247.87	238.23	246.67	247.87	247.87
	All	247.18	237.69	245.72	247.18	247.18
	Lenient	0.081	0.005	0.021	0.007	0.027
Return Rate	Default	0.115	0.026	0.114	0.038	0.116
	All	0.114	0.026	0.113	0.037	0.115

Notes:

Table 14: Outcomes Under Uniform Strict & Default Return Policy with Signaling

		Observed	Optimal Prices Ins=Strict, Sig=N	Optimal Prices Ins=Def, Sig=N	Optimal Prices Ins=Strict, Sig=Y	Optimal Prices Ins=Def, Sig=Y
Profit ($M)		2.605	2.590	2.600	2.588	2.599
Consumer Surplus ($M)		206.84	227.07	206.18	229.02	208.56
Surplus Per Consumer ($)		96.21	105.46	95.89	106.39	97.00
Total ($M)		209.45	229.66	208.78	231.60	211.16
Firm Return Cost ($K)		616.556	143.667	612.000	141.131	598.060
Consumer Return Cost ($K)		76.692	18.258	75.667	18.100	75.024
Total ($K)		693.247	161.926	687.667	159.231	673.084
Elasticity		-5.059	-5.391	-5.073	-5.401	-5.082
Mark up		0.336	0.317	0.328	0.313	0.329
MC Pass Through		1.098	1.022	1.091	1.022	1.090
Market Power		50.172	46.738	49.863	46.506	49.797
RC Pass Through		0.084	0.021	0.079	0.021	0.078
	Lenient	13,906	5,067	6,059	12,574	14,839
Profit ($)	Default	16,526	16,749	16,776	16,466	16,456
	All	16,435	16,343	16,404	16,331	16,400
	Lenient	227.79	222.42	219.36	216.73	218.61
Prices ($)	Default	247.87	238.23	246.67	238.20	246.59
	All	247.18	237.69	245.72	237.46	245.62
	Lenient	0.081	0.005	0.021	0.005	0.020
Return Rate	Default	0.115	0.026	0.114	0.026	0.113
	All	0.114	0.026	0.113	0.025	0.110

Notes: