Optimal Pricing and Product Returns*

Kadircan Çakmak[†] Bart Bronnenberg[‡] Tobias J. Klein[§]

December 13, 2022

Abstract

E-commerce has provided consumers with the opportunity to evaluate a product's search attributes prior to purchase. At the same time, its experience attributes, e.g., personal feel and fit, can often not be fully evaluated online without physical inspection and are only revealed after purchase upon delivery. To offset this difficulty, e-commerce retailers and online marketplaces offer lenient return policies to consumers who feel that the goods they bought disappoint. Using a novel data set of transactions and customer returns at a large online marketplace in the apparel industry, we document that product return rates in the apparel industry are high, ranging from about 5% to about 45% of purchases for a given product. In addition, return rates causally depend on the prices paid. We conceptualize how pricing decisions are affected by returns through (1) return rates, and (2) their dependence on price. Estimating a choice-based model of demand that includes product returns, we show that ignoring product returns severely biases estimates of demand and may lead to incorrect pricing decisions. Incorporating product preferences, (price-dependent) return rates, and return costs in the profit function, we show how to solve for optimal prices. While return costs push prices up by increasing the marginal costs, larger price elasticities of sales net of returns pull prices down. Finally, we show how prices, return costs and profits change in equilibrium.

Keywords: product return, demand estimation, optimal pricing, online shopping

JEL Codes: D83, D81, D18, D12, D11, C01, C51, L81, M31

This is a draft that was prepared for the job market. Please contact the authors before citing it.

*We would like to thank participants of the Structural Econometrics Group at Tilburg University for their helpful comments. This paper is the job market paper of the first author.

[†]Tilburg University, contact address k.cakmak@tilburguniversity.edu

[‡]C.E.P.R. and Tilburg University

[§]C.E.P.R. and Tilburg University

1 Introduction

Online shopping has become an integral part of modern consumers' lives. E-commerce currently constitutes 15% of all consumer transactions (USA-Census-Bureau, 2021). The distinguishing feature of online shopping is that it enables consumers to visit stores and view a wide range of products without incurring travel costs. Online shopping also makes it easy to compare products on search features, like price, style, color, or size, and many websites provide these comparisons automatically (see, e.g., Amazon.com, Zalando.com). Another advantage is that e-commerce retailers and online marketplaces can personalize what they offer to consumers, and assist the latter sort through the vast array of goods available from their stores.

However, the advantages of online shopping notwithstanding, fully experiencing attributes like feel and fit is not possible online. This leads to post-purchase preference updates –possibly even reversals– and regret. Competing platforms and online retailers realize this and offer lenient product return policies to attract consumers.

In practice, firms' return policies lead to a large number of product returns. In 2016, a CNBC news item entitled "A $260 billion ticking time bomb" described the ever-increasing costs of product returns to e-retailers (Reagan, 2016). The report called for the management of returns, especially for apparel products. Recent estimates show that the total value of returned items in the United States of America (USA) reached $370 billion in 2018, an amount equal to almost 10% of all retail purchases (APPRISS, 2018; OPTORO, 2018), and continues increasing dramatically after the coronavirus pandemic to $760 billion in 2021 (OPTORO, 2022). Historically, return rates remain at 5-8%; however, a typical return rate for a modern e-commerce retailer currently stands at a much higher level, 15-20%. Moreover, only 20-40% of the return costs can be recovered by e-retailers

(OPTORO, 2016). The return rates are especially high in the apparel category and may reach 40% (OPTORO, 2018). Aside from this being wasteful in direct monetary terms, it is also an environmental disaster. For instance, it is estimated that 10% of all product returns are discarded or destroyed and end up in landfills (Reagan, 2019).

The problems of high return rates may be exacerbated by the fact that firms do not take product returns into account when they make pricing decisions. For two reasons, this may be the case in practice. First, data on orders are usually easier to obtain than data on final sales, as the latter need to first be constructed by linking data on orders and returns. Second, standard pricing formulas do not take returns into account. If, as a consequence, the prices firms set are too low, then the equilibrium number of returned products may be too high.

These considerations shape the three main goals of this paper. Our first goal is to document the extent and nature of product returns in the apparel industry in more detail. This includes providing evidence on heterogeneity in return rates across consumers, products, and brands. Our second goal is to investigate whether there is empirical support for the hypothesis that returns causally depend on the price. This hypothesis is derived from the idea that everything else being equal, higher prices leave consumers with less surplus; therefore, consumers may be more likely to return a product when the price is high. Our third goal is to propose a way to derive optimal prices in the presence of return costs that also takes a possible dependence of the return rate on the price into account.

Using data from a large Turkish online marketplace, we document that product returns vary strongly across products and brands and average between about 5% and 46% for the 100 best-selling female jeans for the observed period in 2019. We show that the return rate indeed increases in price. A direct consequence of this finding is that ignoring product returns makes markets look less price sensitive than they really are. The reason is that final sales depend on price not only because initial sales depend on price, but also because the return rate depends on price. As a consequence, an increase in price leads to fewer initial sales and more returns. In our case, the price elasticities of final sales, when returns are taken into account, are on average 2.1% higher than the elasticities for initial sales, or orders.

We assume that firms do not take returns into account. This allows us to use the inferred demand system to recover marginal cost and to compute prices that maximize profit in the presence of returns. We find that firms would need to lower their prices by 1.5% to maximize profits even if return costs were 0. This effect is solely due to the positive dependence of the return rate on price. When we additionally incorporate return costs, final equilibrium prices go up by 2% on average.

The remainder of this paper is structured as follows: We review the related literature in Section 2. In Section 3, describe our data. We present the main empirical facts about product returns in Section 4. We study optimal pricing in Section 5. Section 6 explains our empirical strategy, including identification, model specification, and estimation methods. Section 7 presents the empirical results for the estimation of demand and return models. Section 8 quantifies the changes in equilibrium prices and profits when product returns are incorporated into the analysis. The last section concludes.

2 Literature

This paper contributes to the literature on customer product returns and optimal pricing in management and economics in three ways. First, we contribute to these fields by documenting that product return rates are heterogeneous (across products and brands) and depend causally on price. Second, we study the effect of product returns on estimates of demand primitives, in particular the price elasticity of final sales, using a rich transaction-level e-commerce data set that includes product returns and consideration sets. Third, we reformulate the firm's problem to include product returns and return costs and derive optimal pricing policies in equilibrium.

Research on product returns can be categorized into three groups. The first group develops applied microeconomic theory for optimal pricing and the return policy. For instance, Che (1996) and Shulman et al. (2011) model monopoly and oligopoly sellers of experience goods and suggest that it is optimal to offer free returns (or small restocking fees) and higher prices. However, when multiple sellers compete, they find equilibria where consumers pay for returns. Ketzenberg and Zuidwijk (2009) suggests an optimal pricing and return policy for a monopoly seller based on a two-period theoretical model where firms charge recovery costs for returned items and a reselling option in a later period. Finally, Ofek et al. (2011) builds a theoretical model to study when brick-and-mortar stores decide to enter online markets in the presence of product returns. Our study contributes to these theoretical studies by deriving an optimal pricing rule for multi-product firms that can be used to numerically solve for optimal prices at the firm level and in equilibrium. Our pricing rule is general and can be used in markets with multiple firms. Inputs into this are estimates of demand for all products as a function of all prices, estimates of return rates and their dependence on price, return costs, and marginal costs. It is also useful from an academic perspective, as it can be used to infer marginal costs when return costs are known.

The second set of studies on product returns seeks to determine the effect of offering pre-purchase information and the effect of return policy design on return rates and post-purchase consumer behavior. Anderson et al. (2009) develop a structural model of demand in which customers decide whether to purchase an item first and then whether to keep or return it. Shulman et al. (2015) also builds a structural model with reference-dependent utilities and estimates it. Using experimental data, they show that the provision of pre-purchase information can reduce the number of product returns. However, these studies do not analyze differential return rates at the product level and their impact on demand. Sahoo et al. (2018) and Minnema et al. (2016) study the effect of customer reviews on product returns. They find that a high number of reviews leads to a lower probability that a product is returned. They also find that a period with unusually many positive reviews (for instance, compared to long-run averages) results in an increase in sales but also an increase in the likelihood that products are returned. Shehu et al. (2020) estimates the effect of free shipping promotions on returns when the firm otherwise charges a fixed amount for shipping. They find that although free shipping increases final sales (initial sales minus returns), the costs associated with returns more than compensate for the additional profit and create an overall loss. In contrast to that, Petersen and Kumar (2009) shows that introducing product returns can increase profitability, by increasing the likelihood of a purchase, despite the cost of returns. Bower and Maxham III (2012) analyzes post-return behavior. They find that spending is significantly higher in the long run if shipping is free for returns. We contribute to this literature by documenting that the return rate causally depends on the price. In our case, this dependence is positive. This finding is in line with the hypothesis that the surplus associated with keeping the product is an important input into the decision to return the product. Everything else equal, a higher price leads to a lower surplus, and therefore the likelihood of returning the product is higher.

The third group of papers zooms in further and focuses on predicting returns at the consumer and product level. Ketzenberg et al. (2020) explores predictive methods to identify customers who abuse return options using transaction-level data from a US-based retailer. Kedia et al. (2019) constructs a predictive model based on hidden customer tastes and product features and estimates the model using deep neural networks. We add to this strand of the literature by documenting that variation in the price can be used to explain variation in return rates.

To the best of our knowledge, there is no literature on how product returns affect pricing when the return rate depends on the price. This includes the absence of work on structural modeling of consumer behavior in an e-commerce context with product returns and measuring the extent to which estimates of price elasticities are affected by product returns. Whereas Dzyabura et al. (2018) focuses on differential return rates across products to make recommendations about which products to sell online, it takes the products' prices, the restocking costs, and the return rates as exogenous inputs. Our contribution is to develop a general pricing rule in the presence of product returns and show how one can use it in practice to determine optimal prices.

3 Data

3.1 Purchase and return data

Our data set is provided by a large Turkish online marketplace specializing in apparel. It covers all purchases and returns at the individual level in the following five categories: "jeans", "dresses", "pullovers", "trousers," and "coats." Our data are at the customer-product-day level and were collected during 105 consecutive days between 29-06-2019 and

11-10-2019. They record purchases and returns for a sample of 1.4 million users and cover 1.3 million transactions. A customer is included in our sample if she made at least one search of an item on the website of the marketplace during one of five randomly selected days.[1] A total of 530 thousand out of the 1.4 million customers in the sample made at least one transaction.

The variables in the transaction data set are identifier codes for consumer, product, time, and transaction. The consumer code is an identifier for the user (user ID) and remains the same if the user logged into the website on a PC or accessed the marketplace through a mobile application. Users visit the website through either the iPhone (50%) or the Android (45%) application in an overwhelming majority of cases. Product attributes are a product code, category, gender, price, and brand code. Product codes are the same for the different sizes of the exact same product, but they are specific to other characteristics including color. The categorical variable "gender" indicates the gender of the intended user of the product, using the following levels: "female," "male," and "unisex." The variable price is a rescaled version of the original price (for confidentiality). The Brand code is an indicator of the brand of the product. Additional attributes are the average rating of the product, the number of reviews, and the number of customers who marked the product as a favorite.

The transaction data also cover a time-stamped return dummy and a classifier for why the product was returned (e.g., wrong size, product damaged, etc.). We can match these data to the click-stream data of the section below by user IDs and time stamps. If the customer orders multiple items in one session, we track this through the transaction code, which is an identifier for the order decision. The average rating is the average number of stars received by the product over the entire observation period (not the average rating

[1] This was done to keep the volume of data manageable. See Section 3.2.

Table 1: Products

sales rank	cumulative market share	brand	rating	average price	std. price	average buying price	std. buying price	return rate
1	4.9	1042	4.4	0.95	0.0	0.95	0.0	8.8
6	15.6	40	4.5	1.9	0.0	1.9	0.0	11.1
11	20.5	18	3.8	1.24	0.25	1.2	0.2	37.2
16	24.4	1042	4.3	1.14	0.0	1.14	0.0	12.4
21	27.6	1042	4.3	0.95	0.0	0.95	0.0	10.5
26	30.3	1042	3.8	1.14	0.0	1.14	0.0	28.3
31	32.8	40	3.9	1.69	0.05	1.7	0.04	24.1
36	35.0	1042	4.4	1.13	0.09	0.99	0.25	19.2
41	37.1	40	3.8	1.9	0.0	1.9	0.0	20.8
46	39.0	40	4.4	1.39	0.16	1.26	0.2	12.8
51	40.8	40	4.6	1.9	0.0	1.9	0.0	6.6
56	42.5	40	4.1	1.9	0.0	1.9	0.0	17.1
61	44.2	1042	4.5	1.06	0.0	1.06	0.0	6.4
66	45.6	40	4.5	1.27	0.13	1.25	0.06	14.4
71	47.0	40	3.8	1.89	0.05	1.89	0.05	14.6
76	48.3	1042	3.9	0.95	0.0	0.95	0.0	7.7
81	49.5	1042	4.0	1.06	0.0	1.06	0.0	22.8
86	50.7	1042	4.4	0.57	0.17	0.45	0.11	13.5
91	51.8	40	4.1	1.79	0.2	1.7	0.22	20.1
96	52.9	4951	4.1	2.47	0.0	2.47	0.0	32.7
90% of market (average)			4.2	1.5	0.1	1.48	0.1	20.8
90% of market (weighted av.)			4.2	1.38	0.07	1.37	0.07	17.1

This table summarizes the product level data for the female jeans category. The statistics described at the bottom rows (denoted as all) cover 90% of the sales in the market. We report the average across products and the average weighted by market share.

at the time of purchase). We also observe the total number of reviews for each product.

Table 1 shows summary statistics for a selection of products from the female jeans category. The selection is the same in this table and in some of the tables and figures that will follow: we show data for every fifth product within the first 100 ordered by sales, from 1 to 96. The second column of the table shows that the top 100 products account for more than 50% of the sales in the category. The third column shows brand identifiers. There are two dominant brands and the market can be thought of as concentrated. The product ratings have a mean of 4.2 out of 5 and vary between 3.8 and 4.6 across the top-selling products. The average price is the average of the prices observed over the period of 105 days our data covers, and the average purchase price is weighted by the number of purchase instances. The standard deviations in price are computed per product over the entire period. If a product is always sold at the same price, the purchase price and average price are by construction the same, and their standard deviation is 0. For products with price variation, we observe lower average buying prices. The last column shows that there is considerable heterogeneity in return rates. Moreover, the return rate is generally lower for products with higher rankings.

Table 2 shows summary statistics for purchase and return decisions for customers who purchased at least one item from the same category, female jeans. At the 90th percentile, customers purchase 3 items. On average customers buy once every 22.5 days and the Fraction of products returned is about 18%.

3.2 Click-stream data

In addition to the purchase and return data, we also use click-stream (search) data for the same customers and categories. These data were collected for five randomly selected days between 29-06-2019 and 11-10-2019. The unit of observation in the click-stream data set is a click by a consumer (identified by an individual code) on a product's detail page (identified by a unique product code) through a device (identified by a unique device code) in a search session.[2] There are approximately 21 million observations (clicks) by 1.4 million customers in 3.4 million search sessions for 28,821 distinct products in the click-stream data. Each click is time-stamped. Time is measured in seconds. We capture all activity in any of the 5 apparel categories mentioned above, including session activity

[2] A search session is a unique identifier determined by the (web) app that is associated with the visitor when she starts her session. It remains the same until the user closes the app or browser.

Table 2: Purchases and returns

	average	std.	min.	10th pctl.	med.	90th pctl.	max.
purchases							
purchased overall	1.54	1.36	1.0	1.00	1.00	3.0	48.0
value all purchases	1.41	0.52	0.3	0.84	1.33	1.9	13.3
number of days with purchase	1.17	0.48	1.0	1.00	1.00	2.0	16.0
average days between days with purchase	22.5	20.4	1.0	3.00	16.0	53.0	104.0
average number of products bought on same day	1.16	0.44	1.0	1.00	1.00	2.0	16.0
returns							
fraction number products returned	0.18	0.36	0.00	0.0	0.0	1.0	1.0
number purchase days with at least one return	0.27	0.55	0.00	0.0	0.0	1.0	10.0
fraction number products returned if at least one returned	0.96	0.14	0.17	1.0	1.0	1.0	1.0
... as previous, and multiple were bought on same day	0.68	0.25	0.17	0.5	0.5	1.0	1.0
fraction value all purchases returned	0.18	0.36	0.00	0.0	0.0	1.0	1.0

This table summarizes the transaction level data for the female jeans category. Only customers who bought at least one item from this category are used in the analysis. Customers who are likely to be reselling (those who buy more than 50 items) are excluded from the analysis.

Table 3: Choice sets

sales rank product bought	fraction sessions	cumulative market share	size choice set	average buying price	average price other products	std. price other products
1	0.39	2.4	5.1	0.95	0.84	0.14
6	0.07	10.1	5.1	1.90	1.20	0.14
11	0.07	17.0	3.6	1.18	1.24	0.17
16	0.04	21.7	6.2	1.14	0.90	0.12
21	0.03	25.5	6.1	0.95	0.91	0.15
26	0.01	27.9	4.9	1.14	0.99	0.14
31	0.01	30.6	5.3	1.64	1.61	0.17
36	0.02	32.7	4.9	1.14	1.06	0.11
41	0.01	34.5	8.6	1.84	1.26	0.26
46	0.01	35.3	8.8	1.43	1.12	0.27
51	0.04	37.9	6.8	1.90	1.34	0.21
56	0.02	39.5	6.3	1.90	1.37	0.22
61	0.03	41.2	4.7	1.06	0.95	0.13
66	0.00	42.8	3.0	1.14	0.95	0.19
71	0.01	44.4	3.6	1.90	1.25	0.10
76	0.02	45.0	6.7	0.95	0.95	0.17
81	0.01	46.4	6.0	1.06	0.96	0.18
86	0.02	47.3	2.9	0.38	0.67	0.09
91	0.01	49.2	5.1	1.74	1.18	0.09
96	0.02	50.3	4.8	2.47	1.56	0.20
all (bought one of 100 product)	3.33	51.5	5.3	1.27	1.05	0.15
bought one of the other products	2.4		5.3	1.6	1.2	0.18
bought no product	94.28		5.5		1.54	0.28

This table summarizes the characteristics of the individual choice sets based on 5 days of click-stream data set. We consider a sequence of clicks in a day as a search session and they are included in this table's analysis if the customer clicked at least two product in the female jeans among the products in Table 1. If there are multiple purchases in the same search session, the same search session is used multiple times.

across more than one category.[3]

Table 3 shows summary statistics on choice sets, which are defined as the sets of products customers look at (these are sometimes also referred to as search sets). The structure of the table is similar to the structure of Table 1. Each row is for a product that was eventually bought. The table shows the fraction of sessions in which the product is searched. Next, we show the market shares. They are numerically different from the ones in Table 1 because here we report them for the five days for which we have click-stream data. As expected, the differences between the ones reported here and the ones in Table

[3]Thus, when a consumer clicks first on the detail page of a dress and then continues with a pullover, we observe search for both categories.

1 are minor. The average size of the choice set, conditional on buying a product, is about 5 products. However, surprisingly, there is considerable heterogeneity in the size of these conditional choice sets across different products. The table also shows that choice sets contain products of similar prices. That is, there is a strong correlation between the observed purchase price and the price of other products in the choice set.

In the structural estimation part, we restrict our sample to the 100 best selling products in the "female jeans" category.[4]

4 Two stylized facts on product returns

In this section, we first document that return rates differ across products. Then, we show that the likelihood to return a product depends positively on the price.

4.1 Heterogeneity of return rates across products

Accounting for product return decisions in demand estimates is essential for several reasons. Chiefly among them, if the return rates vary across products or brands, product return decisions are likely related to product features that cannot be observed online. To illustrate how different return rates are, even among the most popular products, we plot in Figure 1 the observed return rates for 20 out of the 100 best-selling products. As before, in Table 1 and 3, we choose every fifth product in the female jeans category. Figure 1 shows that there is substantial variation across products, even within brands. Figure A.2 in Appendix A shows that the variation in return rates is also large across brands.

We also analyze return rates across gender and categories. Table A.1 in Appendix A

[4] The reason for this is that we only observe prices for products bought. By restricting the sample to considering the most popular products, we minimize the requirement to impute missing prices. The missing prices are 30% of all prices. The analysis presented here imputes missing prices using the last price observed. Prices in our online setting are all national.

Figure 1: Return rates for female jeans

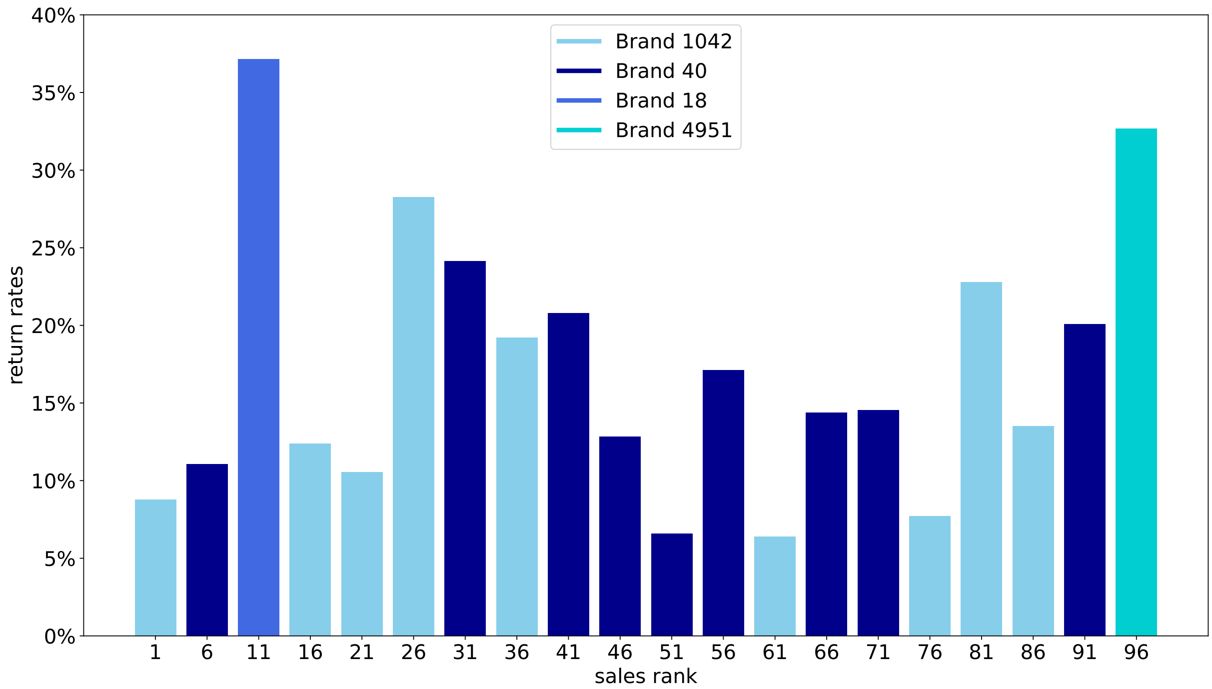

In this plot, we have the average return rates of a subset of best seller products. The colors of the bars represent the brands.

shows that return rates are higher for females. This is driven by return decisions in the dress category. In other categories, return rates across genders are similar. As one may expect, return rates are larger when the uncertainty about the fit is higher. For instance, the pullover category has smaller return rates than the female dress category.

Together with the return decisions, we also observe stated reasons for returning an item. Table A.2 in Appendix A) shows that a large fraction of individuals names a size mismatch as the main motivation to return an item. Among the other reasons, taste mismatch, delivery issues, and payment problems were common other reasons.

Finally, we observe product ratings in our data set. We will use this information as one of the control variables in Section 7.2 where we run several logistic regressions to predict product returns.[5] For completeness, Appendix A Figure A.3 shows that return

[5]The rating data set is collected in October 2020 and the ratings are the ratings recorded at the time of sampling, not at the time of purchase.

rates and ratings are inversely related.

4.2 Return rates and prices

There is no clear behavioral or economic foundation that foreshadows or dictates the direction of the relationship between prices and return rates. On the one hand, customers may have confirmation bias and keep the expensive products, which leads to a lower return rate when the price increases. On the other hand, the monetary cost of a taste mismatch increases with price and customers may be more motivated to return the item and replace it. In the latter scenario, we would expect to see a larger return for more expensive items. The direction and empirical magnitude of the relation between prices and return rates are important because they will have an impact on the optimal pricing decision. To this end, we present evidence about this relation using several descriptive analyses.

Figure 2 shows the results of a kernel regression of return rates on products' price deviations from the category mean. We find the return rates and weighted averages of the prices for each product together with their purchase frequency. Next, we compute the deviation of average prices from their mean category prices. We implement the kernel regression by taking the purchase frequencies (as the weights of the return rate-price deviation pair observations) into account. The figure shows results for all products together and separately for 3 categories.[6] The figure shows a clear positive dependence of return rates on prices. Interpreting the relation when pooling across categories, return rates are about 30% (4.7 percentage points) higher over the inter-quartile range (-25% to 35%) of price deviations. suggesting that when the monetary cost of a mismatch is high,

[6] We don't show the pullover and coat categories because they are out of season in August and September which covers almost all our data collection period.

Figure 2: Dependence of return rate on price (using only between-product price variation)

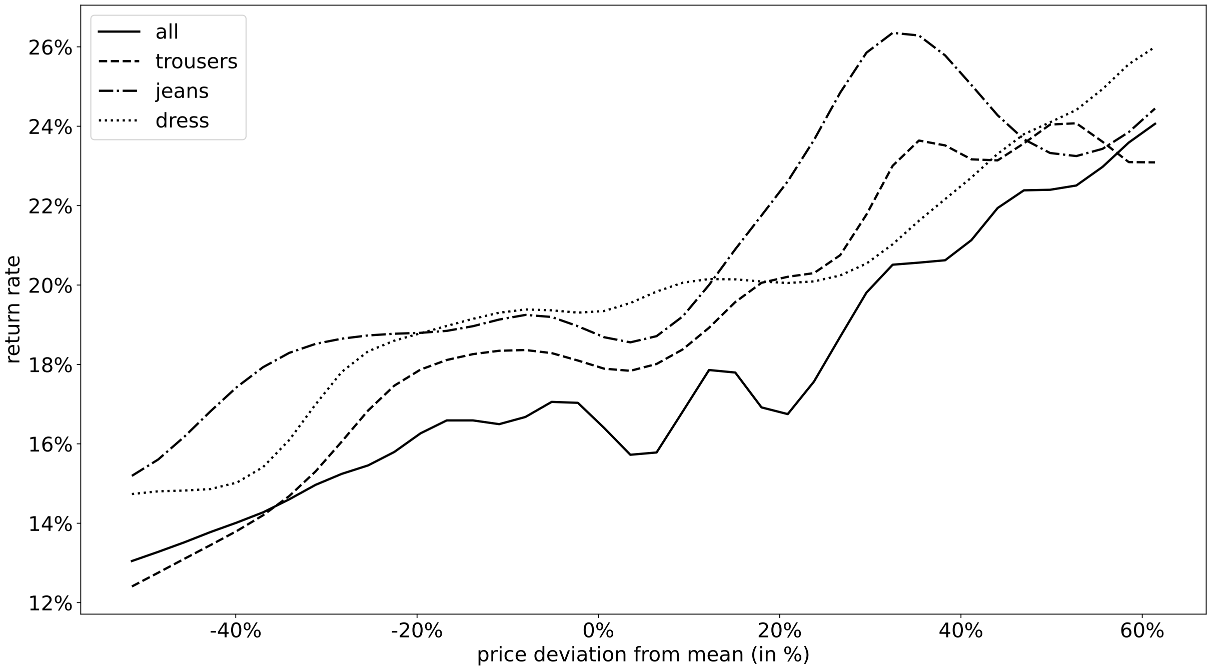

This figure shows the results of kernel regression for return rates on price deviations of a product from its mean category price. We don't show pullovers and coats since they are purchased less frequently due to the summer season, though they are included in all. The bandwidths are set according to the rule of thumb bandwidth.

customers have a stronger tendency to return the item, i.e., that the second argument above dominates.

Figure 2 is based on exploiting variation in prices across products. We next look at the relation between price and returns using only the variation of prices within products.[7] For this, we take the subset of products exhibiting price variation and calculate the weighted average price at which each product is sold. The weight is given by the quantity sold. Then, we create an observation for each observed value of the price a product is sold at and calculate the deviation of that price from its weighted average price and the return rate when the product was sold for that price. Finally, we perform a kernel regression of return rates on the within-transformed prices, again using sales as the weight.

[7] Table 1 shows that such variation exists for some products. Such price variation exists because of price experiments conducted by the sellers on the marketplace and occasional sales campaigns for some products.

Figure 3: Dependence of return rate on price (using only within-product price variation)

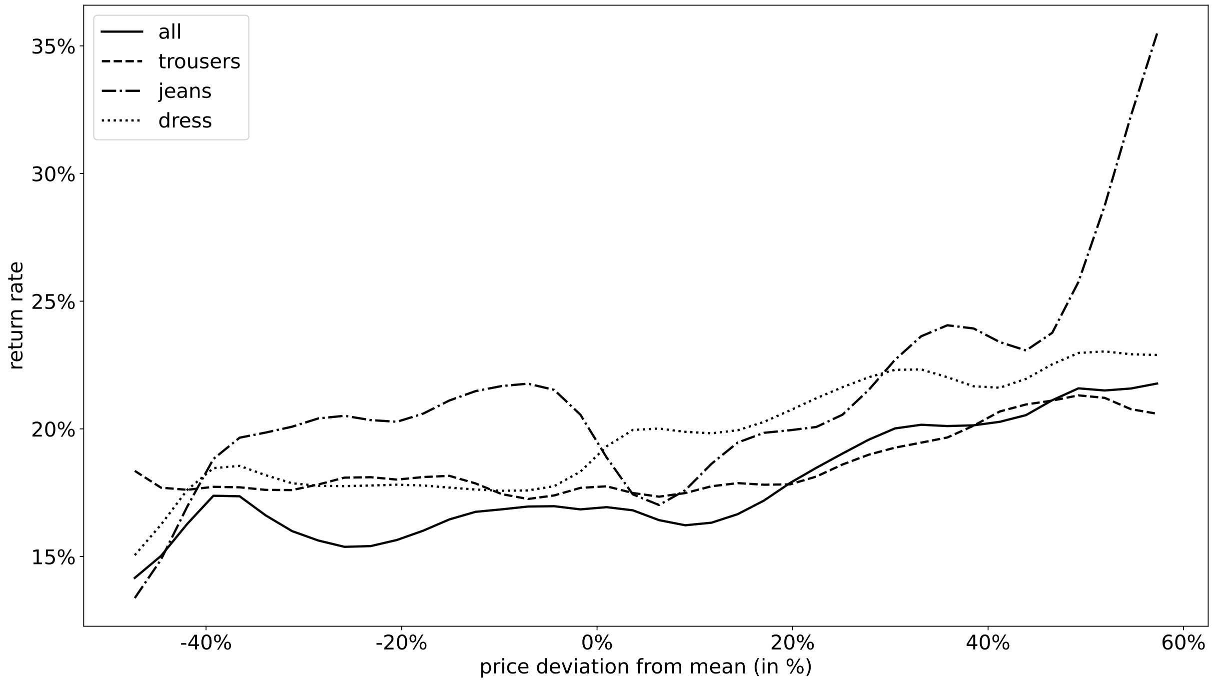

This figure shows the results of kernel regression for return rates on price deviations of a product from its mean price. Products showing price variation across time are included to be able to identify the effect of price change on return rate at the product level. The bandwidths are set according to the rule of thumb bandwidth.

Figure 3 shows the result. When again pooling across categories, return rates are once more 30% higher over the inter-quartile range (-25% to 30%). So, if a product is bought at a lower price than its average sales price, it is less likely to be returned while it is more likely to be returned if it is bought for a higher price than its mean price.

In Section 7, we complement these results with estimates from logistic regressions. The results later serve as input into our counter-factual pricing scenarios in Section 8.

5 Optimal pricing

In this section, we study pricing in the presence of product returns. We start with the standard firm problem and then modify the demand model and cost structure to incorporate product returns. Then, we derive optimal prices and discuss how they depend on product returns and the costs of these returns to the firm. Finally, we discuss how the dependence of the return rate on the price and the existence of return costs push optimal prices in opposite directions.

5.1 Incorporating product returns into the firm problem

Consider that firms have the following gross profit function, π_f:

$$\pi_f = \sum_{j \in f} s_j(p, \theta) \cdot (p_j - mc_j) \tag{1}$$

where $s_j(p, \theta)$ is the demand for the product j given a price vector p and preference parameters θ. $s_j(p, \theta)$ is multiplied by the markup, i.e., the difference between the respective price p_j and marginal cost mc_j, and summed up across the products of the firm f.

In the standard interpretation, choices are represented as being final and product returns are not modeled. However, product returns affect profits in at least two ways: (i) they reduce revenue, and (ii) they raise costs. To represent (i), we introduce a *return rate function* $r_j(p,\gamma)$ that interacts with $s_j(p,\theta)$ and determines the *returned* and *kept* quantities as $r_j(p,\gamma)\cdot s_j(p,\theta)$ and $(1-r_j(p,\gamma))\cdot s_j(p,\theta)$, respectively. Notice that –in line with the evidence we presented in Section 4 above– the return rate is product-specific and is allowed to depend on the price.

To account for (ii), we include the costs of returns when the product is returned, in combination with the function describing the return rate. The return cost is born by the firms, not consumers. Not only is this consistent with business practice in our empirical application, but it is also becoming the dominant arrangement in e-commerce, possibly due to competition over customers.

Restating the profit function to incorporate reduced demand from returns and their costs, we thus propose

$$\pi_f = \sum_{j\in f} \underbrace{s_j(p,\theta)\cdot(1-r_j(p,\gamma))}_{\text{actual transactions}}\cdot(p_j - mc_j) - \underbrace{s_j(p,\theta)\cdot r_j(p,\gamma)}_{\text{returns}}\cdot rc_j, \tag{2}$$

where rc_j is the constant marginal return cost of the product j. The markup that a firm obtains is scaled down by $1-r_j(p,\gamma)$ while the cost it covers has the additional part $s_j(p,\theta)\cdot r_j(p,\gamma)\cdot rc_j$.

We now start with an analysis of $s_j(p,\theta)\cdot(1-r_j(p,\gamma)$. We are interested in the demand elasticity, i.e., how price affects the quantity sold versus the quantity sold and kept (i.e., net of returns), and the relation between these elasticities. This relation can be made explicit using the following derivation of own price elasticities:[8]

$$\frac{\frac{\partial(s_j(p)\cdot(1-r(p)))}{s_j(p)\cdot(1-r(p))}}{\frac{\partial p}{p}} = \underbrace{\frac{\partial s_j(p)}{\partial p_j}\cdot\frac{p_j}{s_j(p)}}_{\text{Conventional Elasticity}} - \underbrace{\frac{\partial r_j(p)}{\partial p_j}\cdot\frac{p}{1-r_j(p)}}_{\text{Adjustment from Returns}} \tag{3}$$

The deviation from the baseline elasticity estimate will depend on the second term on the right-hand side of equation (3). Theoretically, we know that $s_j(p)$ is a decreasing function of p_j but not much is known about the sign and size of the marginal impact of prices on consumers' tendencies to return items $\partial r_j(p)/\partial p_j$ in general. In Section 4, we have argued that this dependence could be positive because higher prices leave consumers with less surplus, everything else equal, or negative, when there is confirmation bias. Moreover, we have documented that the dependence on the price is positive for our data, suggesting that at least in our case, the first effect dominates. Consequently, the second term in (3) that is subtracted is positive, which means in turn that final sales depend more on price than initial orders. That is, conventional estimates of elasticities that are based on order data and that do not take product returns into account are biased toward zero.

5.2 Pricing rule

We now show that the effect of price on the tendency of consumers to return items also impacts pricing. The derivation of the optimal pricing rule given the standard profit function in equation (1) is well-known (see e.g. Nevo, 2001)

$$\frac{\partial \pi_f}{\partial p_j} = s_j(p) + \sum_{l\in f}(p_l - mc_l)\frac{\partial s_l(p)}{\partial p_j} = 0 \Rightarrow p = mc + \left(\Omega \circ -\frac{\partial s(p)}{\partial p}\right)^{-1}\cdot s(p),$$

[8]To simplify notation, we omit parameters θ and γ but keep p to make it clear that both s and r are functions of the price vector.

where Ω is the ownership matrix, and $\circ$ is element-wise multiplication. Entry ij of Ω is an indicator for products i and j being owned by the same firm.

For the modified profit function in equation (2), we can derive the optimal pricing equation in a similar way:[9]

$$p = mc + \left(\Omega \circ - \left(\frac{\partial s(p)}{\partial p} + \underbrace{\frac{\partial(-r(p) \circ s(p))}{\partial p}}_{A}\right)\right)^{-1} \cdot \left(s(p) \underbrace{\circ (1 - r(p))}_{A} - \underbrace{\Omega \circ \frac{\partial(s(p) \circ r(p))}{\partial p} \cdot rc}_{B}\right) \tag{4}$$

where the additional terms with respect to the original rule are annotated with curly brackets. The terms denoted A are related to the changes in purchased quantities. The term denoted B is incorporating return costs associated with returned products into the pricing rule. The direction of the effect of these additional terms on the price is ambiguous. Therefore, determining the direction of the price adjustments from returns requires empirical measurement in a real-world scenario. In the next section, we will some related intuition by studying the special case of a single-product monopoly seller.

5.3 Illustrative example

Constructing a single product monopoly example is useful for understanding how returns affect pricing in equation (4). If we take the first order condition of equation (2) and

[9] See Appendix C for the derivation.

rearrange terms in a convenient way, we get[10]

$$p = \left(mc + \frac{\tau'(p)}{s'(p) - \tau'(p)} \cdot rc \right) + \frac{-s(p)}{s'(p) - \tau'(p)} \cdot (1 - r(p)) \tag{5}$$

where $\tau'(p)$ is the partial derivative of $s(p) \cdot r(p)$ with respect to p, i.e. the total share of returned goods given price. Suppose for now that the return rate does not depend on the price, i.e. set $r'(p) = 0$. Then

$$p = \underbrace{\left(mc + \frac{r(p)}{1 - r(p)} \cdot rc \right)}_{\text{marginal cost component}} + \underbrace{\frac{-s(p)}{s'(p)}}_{\text{markup component}} .$$

When prices do not affect the return rates, the markup component is as usual. However, we allocate the return cost coming from the returned items on the kept items, leading to the term $r(p)/(1-r(p))$ preceding the return cost, rc. This shows that one of the effects is that product returns have an effect similar to an increase in marginal cost. The intuition here is that for each item that is eventually sold, a marginal cost has to be paid when the item is sold, and a return cost has to be paid for $r(p)/(1 - r(p))$ items that are returned.

In addition, as we have documented in Section 4, the return rate depends positively on the price. To see the effect thereof, assume that $r(p) = \bar{r} + \alpha_r \cdot p$. Substituting this into $\frac{\tau'(p)}{s'(p)-\tau'(p)}$ and $\frac{-s(p)}{s'(p)-\tau'(p)}$ in equation (5) and taking the derivative with respect to α_r shows how the marginal cost and markup components change due to this. Appendix C shows that the derivative is positive for the first term and negative for the second term. The intuition is as follows: since α_r increases the return rate at a given price, there will be an increase in the transfer of return costs coming from a returned share on the kept share. This increases the total marginal cost of a final sale. For the markup term, actual

[10] To ease the notation: $f'(x)$ stands for the derivative of $f(x)$ with respect to x.

sales will decrease at a higher speed when the price is elevated so firms will have smaller room for price increases. This decreases margins.

Put differently, product returns have two effects on prices. They lead to upward pressure on the price because they have a positive effect on the marginal cost; they lead to downward pressure because they make demand more price elastic (when the return rate depends positively on the price, as it does in our case).

Figure 4 illustrates this for a single product monopolist facing a logistic demand and return function, respectively. To construct the three plots, we set $\{\bar{u}_d = -2,\ \alpha_d = 2\}$, and $\gamma = \{\bar{r} = -2,\ \alpha_r = 0.5\}$ where α_r and α_d are taste parameters of price in return rate and demand functions, respectively.

Figure 4(a) plots profits against prices for different levels of the return costs. We see that profits decrease in the return cost for any price. The optimal price when return costs are zero is lower than the price when return costs are ignored; and the higher the return costs, the higher the optimal price.

In Figure 4(b), we plot the optimal price against the return rate. We do so for 4 scenarios: incorporating the effect of returns on the price sensitivity of demand, or not incorporating it; combined with incorporating return costs, or not incorporating return costs. We see that the presence (absence) of return costs makes the optimal price rise (fall) as a function of the return rate.

Finally, Figure 4(c) plots optimal prices against the price sensitivity of returns. The marginal cost- and markup terms are shown by dotted and dashed lines (right axis) and the baseline price and optimal price (sum of the marginal cost and markup terms) are shown by solid and dash-dotted lines (left axis), respectively. The figure shows that the stronger the effect of price on the return rate, the larger the marginal cost effect on prices

Figure 4: Illustration

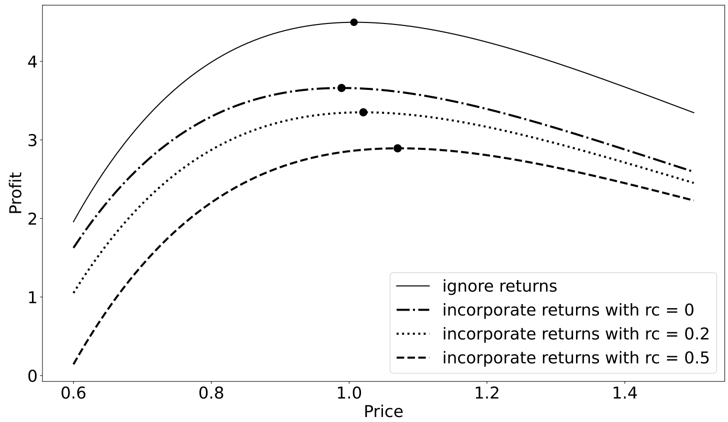

(a) Dependence of optimal price on return costs

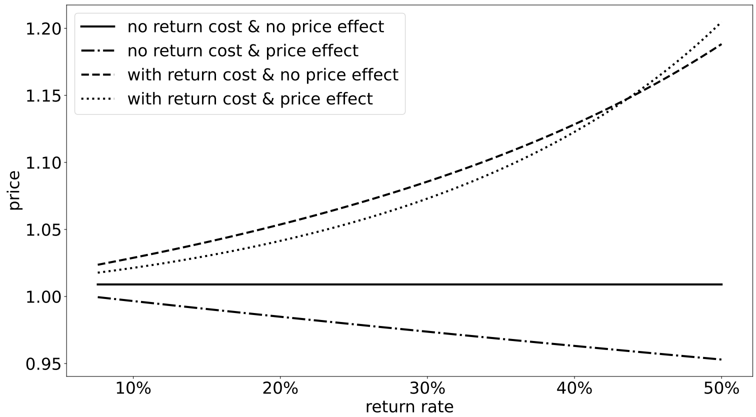

(b) Dependence of optimal price on return rate

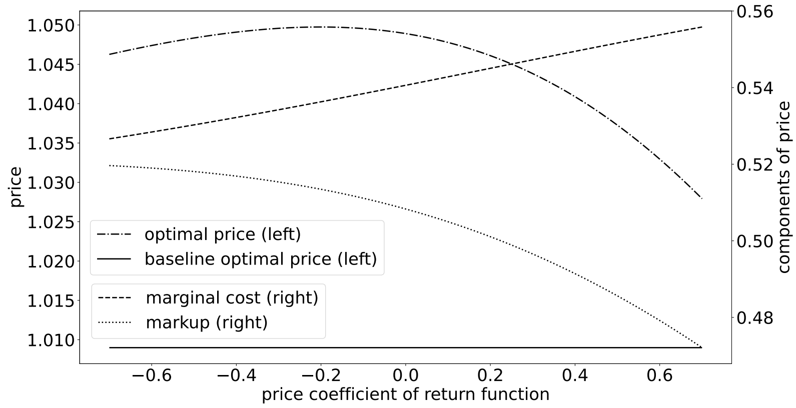

(c) Dependence of optimal price on the price sensitivity of returns and the smaller the markup effect on prices. The combined effect (dash-dotted line) is non-monotone.

6 Empirical approach

The primary goal of this paper is to investigate how the pricing decision is affected by product returns. In Section 5, we have shown that for this, we need to estimate on the one hand how initial orders for each product depend on all prices; and on the other hand how the return rate varies across products and depends on the price.

One way to proceed would be to take a fully structural approach. For this, one could for instance spell out a sequential search model in which consumers have the option to return the product at some (possibly non-monetary) cost. Then, one could estimate the model and obtain demand for given prices, cross-price effects, and return rates. The advantage of this approach is that it is fully micro-founded and internally consistent. Therefore, one can meaningfully characterize welfare effects, which is however not the focus of this paper.

Pursuing a fully structural approach also comes at a cost. Such a model is very complex, which means that many modeling choices need to be made. For instance, a parametric functional form for the search cost distribution has to be specified; and whether and how search costs vary across products. Moreover, identification is less direct and it is substantially harder to estimate the model parameters.

Therefore, we instead use empirical models that are also used by firms when they make pricing decisions. In particular, we use a multinomial logit model for initial orders and a binary logit model for the return decision. The option value of being able to return the product is captured by the intercepts in the multinomial logit model.

At the same time, we allow for rich heterogeneity across consumers by using observed consideration sets when we construct the multinomial logit probabilities. Thereby, we can capture substitution patterns that are much richer than the ones that are traditionally implied by the multinomial logit model. This is of key importance for obtaining realistic results when we conduct our counterfactual simulations.

6.1 Initial orders

We use a multinomial logit model with individual-specific consideration sets to model initial orders. We directly observe the choice sets for a subset of individuals for which we have the 5 days of click stream data. These data are also informative about the distribution of consideration sets for a given choice that was made. For the other individuals –those for which we only have transaction-level data– for which we do not observe consideration sets, we use the information on the distribution of consideration sets from the click-stream data to obtain choice probabilities using simulated consideration sets. Details on this can be found in Appendix B. Estimation is done by maximum likelihood using individual-level data.

The choice probability of product j at time t by individual i with a known consideration set $\mathcal{S}_i$ is given by the usual multinomial logit formula:

$$Pr(y_{it} = j \mid \mathcal{S}_i) = \frac{exp(u_{jt})}{\sum_{k \in \mathcal{S}_i} exp(u_{kt})} \tag{6}$$

where $u_{jt} = X_j \cdot \beta - \alpha_d \cdot p_{jt}$ and X_j is a vector of characteristics of product j, β are associated taste parameters, α_d is the price coefficient and p_{jt} is the price of the product j at time t.

If we have simulated consideration sets, we have to take an average over the simulated probabilities corresponding to simulated consideration sets. Let's denote the l^{th} simulated consideration set of the product j as CS_{jl} :

$$Pr(y_{it} = j) = \frac{1}{L}\sum_{l=1}^{L} Pr(y_{it} = j \mid \mathcal{CS}_{jl}) = \frac{1}{L}\sum_{l=1}^{L} \frac{exp(u_{jt})}{\sum_{k \in \mathcal{CS}_{jl}} exp(u_{kt})}$$

Consumers for whom we have click stream data choose between the outside option and all products they have inspected. In line with this, the sets $\mathcal{S}_i$ for consumers for whom we have click-stream data contain an outside option. The consumers for whom we have transaction data buy one product for sure. Therefore, the draws of sets $\mathcal{CS}_{jl}$ for those consumers for whom we only observe what they have bought do not contain an outside option. That is, we model product choice conditional on buying any inside good.

The parameters for the demand model are identified from the variation in choice sets, the variation in average ratings within the brands' products, and price variation within products across days. Identification of brand dummies and the coefficient for the average rating depends on the relative number of purchases for these products after normalizing the outside option's value to zero.

We estimate the parameters by maximum simulated likelihood. For this, we simulate consideration sets as described in Appendix B. We use them to simulate choice probabilities for individuals for which we do not observe the consideration sets. For those who have observed consideration sets, the computation of multinomial logit probability is straightforward.

The log-likelihood function is

$$L(\beta, \alpha_d) = \sum_{j=1}^{J}\sum_{t=1}^{T} n_{jt} \cdot log\left(\frac{1}{L}\sum_{l=1}^{L} Pr(y_{it} = j \mid \mathcal{CS}_{jl})\right) + \sum_{i \in \mathcal{I}} log(Pr(y_{it} = j \mid \mathcal{S}_i)), \quad (7)$$

where n_{jt} is the frequency of choosing product j at time t.

6.2 Returns

We use a binary logit specification to model the return rate $r(p, \gamma)$. We capture heterogeneity across products by including product fixed effects. We then estimate the dependence of the return rate on the price.

We can identify the coefficients in the logistic regression of return dummies on product fixed effects and price from the relative frequencies of product returns and price variation in transactions within products.

When estimating the parameters, we are not restricted to the sample of products (100 best sellers) that we use for estimating demand. The reason is that there is no missing price problem in the transaction level data, i.e. we observe the characteristics of a product for every consumer. This provides us with the opportunity to study the dependence of the return rate more broadly.

7 Results

7.1 Initial orders

Table 4 shows our demand estimates. Column (1) reports estimates when we only use the transaction data; column (2) reports estimates when we also use the click-stream data.[11]

We control for brand fixed effects and for differences across products within the same brand using the average rating. Table D.3 in Appendix D shows that brand coefficients are distributed in a relatively large range. This captures the taste heterogeneity of customers

[11]Using only the click-stream data is not feasible, as there is not enough variation in prices within brands across the 5 days for which we have click-stream data to identify the price coefficient.

Table 4: Demand parameters

	(1) only data 100 days	(2) both data sets
price	2.17	1.914
	(0.018)	(0.015)
average rating	0.705	0.692
	(0.018)	(0.018)
brand fixed effects	yes	yes
purchase week fixed effects	yes	yes
log likelihood	-74452	- 87606
nr. observations	60544	118706[a]

This table summarizes the estimation results for three models: (1) demand estimation for 100 days of purchase data without observed consideration sets and (2) our custom model described by the likelihood function in Equation 7. The characteristics include brand fixed effects, average rating and price. The standard errors are robust standard errors of MLE estimation.
a Number of transactions in 100 days and search & transaction (including no purchase) in 5 days.

controlling for price and ratings and determines the level of initial product orders. As expected, demand depends positively on the rating. In addition, we control for seasonal patterns using week fixed effects.

The average implied elasticity of demand with respect to the price is around -2.

7.2 Product returns

Our estimates for the return model are reported in Table 5. They complement the evidence that we have already presented in Figures 2 and 3 and that we have discussed in Section 4.

We estimate six specifications of the model using data for all category-gender pairs and one model on the demand estimation sample, i.e. the 100 top-selling female jeans. The first three columns use price as a continuous variable and control for (1) product and time FE, (2) brand and category-gender FE along with the average rating, and (3) time FE in addition to (2). Columns 4-6, repeat the same regressions in 1-3 but use different variables capturing the dependence on the price, namely dummies for intervals of price deviation from the respective mean prices of the product. The last column is the same model as column 1, but here we restrict the sample to the one we use for estimating demand. The coefficients are significantly positive and range between 0.2-0.4. This means return rates depend positively on the price, even after controlling for product and time FE. We observe similar values in columns 1-3 and column 7. This exhibits that in our showcase category, female jeans, there is no systematic difference from the overall pattern. We see a monotonically increasing trend in the coefficients of dummies as the deviation from the mean price increases. We will use the results in column 7 in our counterfactual simulations when we need $r(p)$ of (4).

Table 5: Parameters of the return model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
price	0.258	0.393	0.413	-	-	-	0.206
	(0.018)	(0.005)	(0.005)				(0.119)
(-40%< $\tilde{p}$ <-20%)[a]	-	-	-	0.372	0.303	0.321	-
				(0.049)	(0.044)	(0.045)	
(-20%< $\tilde{p}$ <0%)	-	-	-	0.444	0.356	0.4	-
				(0.048)	(0.043)	(0.043)	
(0%< $\tilde{p}$ <20%)	-	-	-	0.511	0.414	0.466	-
				(0.048)	(0.043)	(0.043)	
(20%< $\tilde{p}$ <40%)	-	-	-	0.62	0.545	0.576	-
				(0.049)	(0.045)	(0.045)	
(40%< $\tilde{p}$ <60%)	-	-	-	0.688	0.625	0.619	-
				(0.056)	(0.05)	(0.051)	
average rating	-	-0.417	-0.388	-	-0.355	-0.349	-
		(0.006)	(0.006)		(0.006)	(0.006)	-
product fixed effects	yes	no	no	yes	no	no	yes
brand fixed effects	no	yes	yes	no	yes	yes	no
purchase week fixed effects	yes	no	yes	yes	no	yes	yes
category-gender fixed effects	no	yes	yes	no	yes	yes	no
log likelihood	-332569	-344917	-342743	-332500	-348332	-345335	-26589
number observations	801524[b]	801524	801524	801524	801524	801524	65929

This table summarizes 7 models of logistic regression of product returns on various control variables using the transactions from all categories & estimation sample. First three columns are logistic regressions of returns on price and (1) product and week, (2) category-gender, and week and (3) category-gender, brand and week fixed effects. The next three columns is the regression of returns on bins of the deviations of prices from their mean at the product level and the fixed effects in the order of first three columns. The last column is the logistic regression of returns on price and product fixed effects for the demand estimation sample of female jeans.

a The intervals represent the range of dummies for the deviation of a product's price from its mean.
b The number of the transactions for products sold more than 20 times and show price variation.

We see similar values for the coefficient on the average rating across these models. They are negative and significant. This means that products are less likely to be returned when they are better rated.

7.3 Elasticities for gross and net sales

We derived the price elasticity of return adjusted demand in (3) for a general multiplicative structure. Relying on our choice for the functional forms, we can continue our analysis from where we left off and calculate the percentage deviation in return adjusted elasticity with respect to the original one. The second term in (3) under the functional form assumptions we have made is

$$\begin{aligned}\frac{\partial(1-r_j(p))}{\partial p_j}\cdot\frac{p}{1-r_j(p)} &= \frac{-(-\alpha_r\cdot exp(\bar{r}-\alpha_r\cdot p))}{(1+exp(\bar{r}-\alpha_r\cdot p))^2}\cdot\frac{p}{1-r(p)}\\ &= \alpha_r\cdot r(p)\cdot(1-r(p))\cdot\left(\frac{p}{1-r(p)}\right)=\alpha_r\cdot r(p)\cdot p\end{aligned}$$

This means the rate of deviation from the original own price elasticity is[12]

$$\frac{\alpha_r\cdot r(p)\cdot p}{\alpha_d\cdot p\cdot(1-s(p))}=\frac{\alpha_r\cdot r(p)}{\alpha_d\cdot(1-s(p))} \tag{8}$$

where the denominator is the absolute value of the original own price elasticity.

Figure 5 shows demand elasticities for the same sample of products as in Figure 1. These are all products that are among the top 100 selling products in the category. The elasticities are higher on average when we use data on sales net of returns. Furthermore, the divergence between the two elasticities varies by product. In Figure D.5 in the Online Appendix, we show the distribution of the ratio between the net and the gross price

[12] In absolute value terms.

Figure 5: Elasticities of sales with and without returns

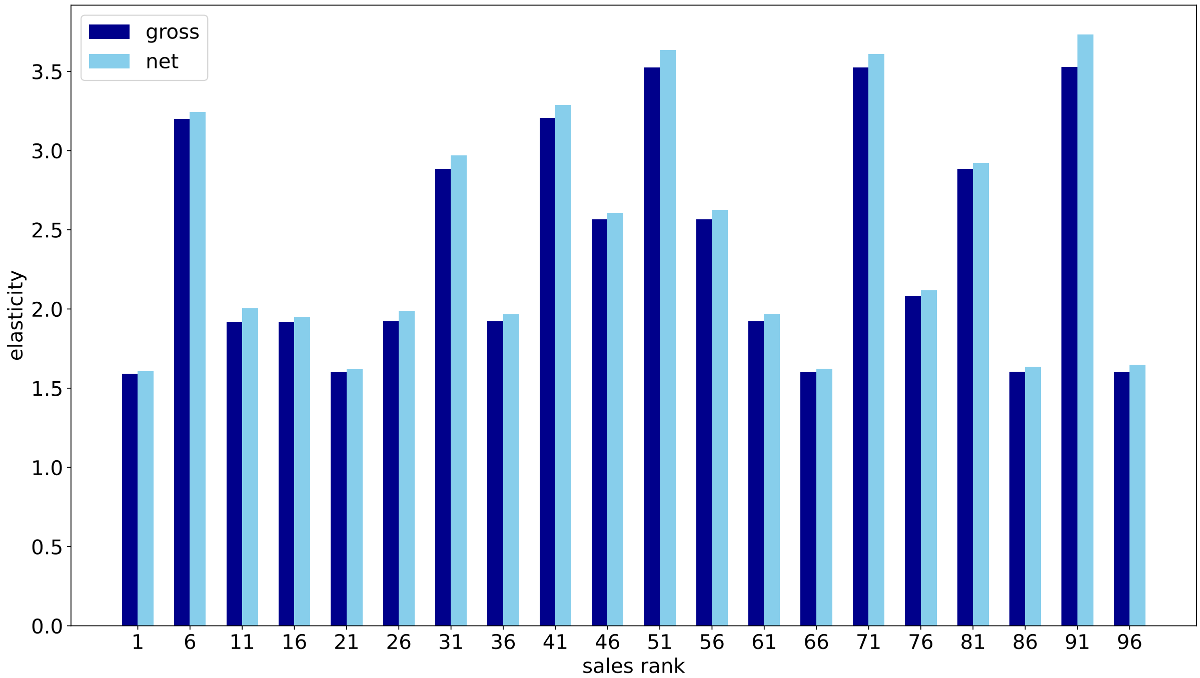

This figure shows the price elasticity of demand for all 100 products. Elasticities based on the demand model without returns are shown with light gray color. The elasticities increase to the levels shown by the dark gray color. These values are computed based on the estimates in Table 5 and 4.

elasticities. Net elasticities are 1-6% larger compared to gross elasticity, on average around 2%.

8 Counterfactual analysis

The empirical estimation results along with the return cost data provided by the marketplace allow us to run counterfactual pricing scenarios where firms set the prices taking into account the effects of product returns. We use the theoretical results derived in Equation (4) to run two counter-factual scenarios: (i) firms use the new pricing rule as their best response while other firms keep their prices constant and (ii) all firms adopt the new pricing rule and they converge to the equilibrium of mutual optimal responses. We summarize the results in Table 6 and illustrate how the prices change in equilibrium

Figure 6: Dependence of price adjustment on return rates

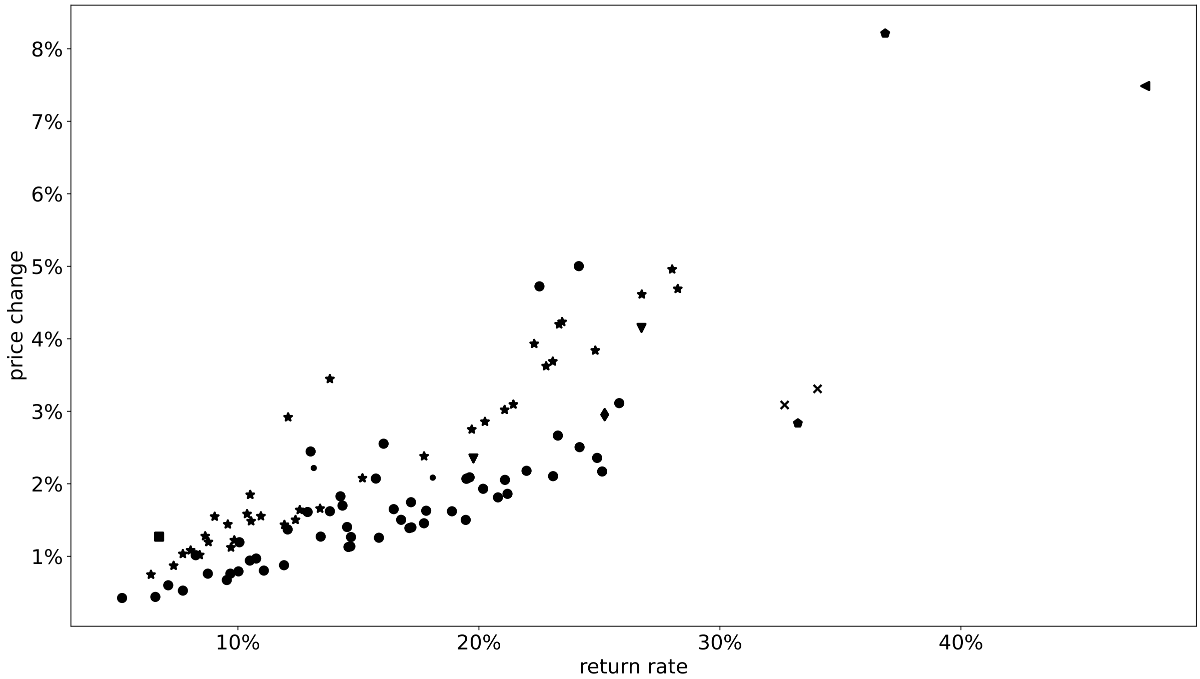

This scatter plot shows the original return rates of the products (in x-axis) and the price changes (in percentage) in the counter-factual equilibrium pricing scenario (in y-axis). Each marker represents a different brand.

in Figures 6 & 7.

In best response pricing, we assume that only one firm adopts the pricing rule in (4) and other firms keep their prices fixed. We do this exercise for each firm separately and obtain the new set of prices for each firm. The counterfactual results based on each firm's modified price (while keeping other firms' prices constant) will be recorded. In equilibrium pricing, all firms respond to the prices they observe with the new pricing formula and they converge to a price equilibrium. We obtain similar results for the best response and equilibrium prices. Therefore, to be concise, we only discuss the results of equilibrium pricing counterfactual.

Figure 6 shows how equilibrium prices shift from existing prices with respect to return rates. Each marker is a product and the marker types represent brands. We see a clear positive link between return rate level and price increase. This is caused by the return

Figure 7: Price changes for selected products

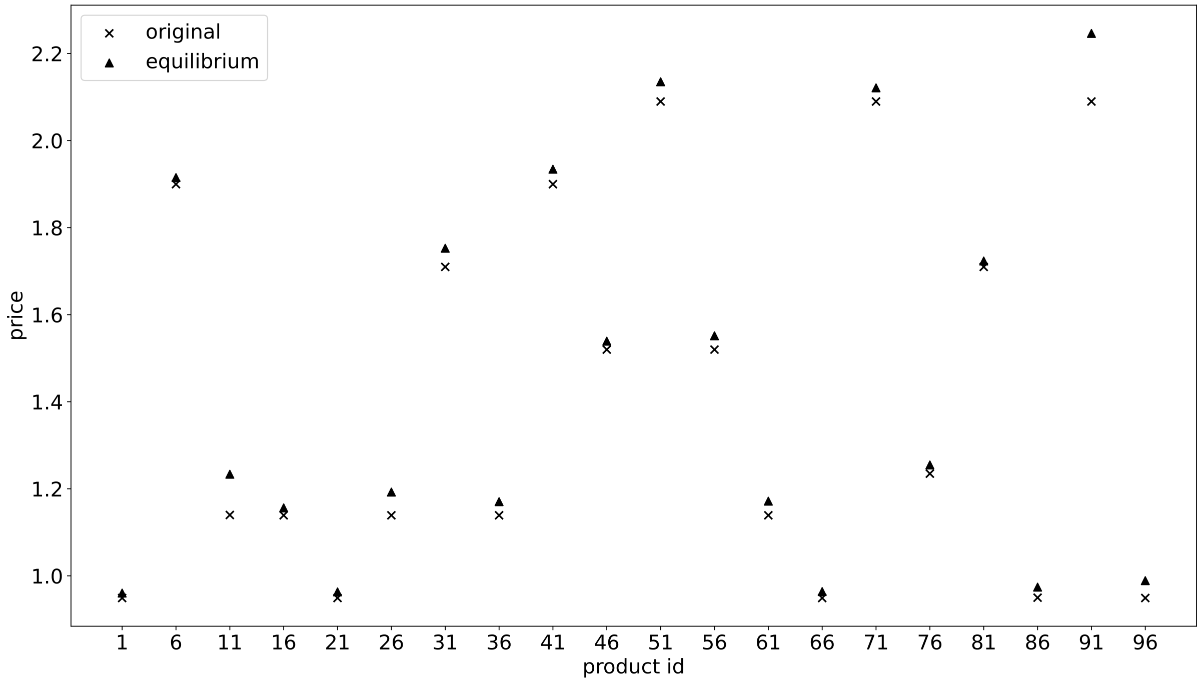

This plot shows the original and counter-factual equilibrium prices for the subset of products presented in Data Section based on the structural estimates when product returns are considered.

cost which imposes a stronger effect on prices as return levels rise. One can also see the changes for the subset of products that we used in the descriptive statistics (Section 3). For a given original price level, we see variation across equilibrium prices, e.g. products 11 and 16. This is mainly due to the variation in product return rates.

The reflection of these price changes can be seen in Table 6. Each row in this table summarizes the changes in the new pricing equilibrium for a firm. The first column is the number of products sold by each firm. We have two dominant firms and several small firms. The following columns show the change in the weighted average of the respective column names. As it is clear from Figure 6, we have an increase in prices for all firms. This leads to a decrease in the initial online orders (note that the outside good's price is fixed at 0). This first-order effect created a decrease in total return costs paid, although the weighted average of return rates may be larger for some firms. The decrease in the

Table 6: Effect of incorporating returns on market equilibrium

# products	price	return rate	initial orders	return cost	profit
53	1.34%	-0.15%	-3.18%	-3.32%	0.42%
36	1.76%	-0.39%	-2.40%	-2.78%	0.40%
2	3.17%	1.07%	-11.57%	-10.63%	1.76%
2	1.95%	0.30%	-3.19%	-2.89%	0.75%
2	3.42%	0.55%	-6.05%	-5.54%	1.89%
2	7.89%	1.20%	-12.96%	-11.92%	2.85%
1	7.42%	1.67%	-21.55%	-20.23%	6.33%
1	2.92%	0.77%	-6.64%	-5.92%	1.82%
1	1.25%	0.18%	-0.09%	0.09%	1.51%

This table provides the percentage changes for the specified set of variables in new equilibrium prices with respect to original prices at the firm level. The return rates and prices are weighted according to the weights of the products within the firm.

weighted averages of return rates for the top two firms shows they may lead the customer to the products with lower expected return rates by adjusting their prices. The decrease in initial orders also leads to a decrease in finalized transactions but at the optimal point, we see profit increases ranging between 0.4% to 6% thanks to the decrease in return costs.[13]

Recall from the discussion in Section 5 that the positive dependence of the return rate on prices leads to downward pressure on prices while taking the additional cost into account leads to upward pressure on prices. Our results show that the second effect dominates. As a result, the number of initial orders decreases and therefore the total number of returns, even though the return rate slightly increases. Fewer resources are spent on product returns, and firm profits increase.

9 Conclusion

In this paper, we study the optimal pricing problem of sellers in the presence of customer product returns. Using a novel transaction and clickstream data set from a large Turkish

[13] The larger values are for the smaller firms. A weighted average would yield 0.6%.

online marketplace, we first document two key facts about return rates: (i) return rates are heterogeneous across products, and (ii) return rates depend positively on price.

We then develop an optimal pricing rule incorporating product returns and show how marginal cost and markup channels change when returns are accounted for. Qualitatively, our theory and the empirical facts suggest that the markup term should shrink while the marginal cost term expands.

We estimate a demand and a return model to quantify the sizes of these two effects and determine which one is dominant. Based on the estimated taste primitives, we solve for the market equilibrium with the new pricing rule. The results suggest that firms decrease their prices up to 5%, with an average of 1.5%, if there were no return costs. However, when the data provider's return cost estimate is incorporated, prices increase up to 6%, with an average of 2.1%. So, on average, 1.5% shrinkage in markup terms is dominated by a 3.6% increase in marginal costs. In the equilibrium, the total number of product returns decreases, although return rates increase, because of the reduction in online orders.

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.

APPRISS (2018). 2018 consumer returns in the retail industry.

Bower, A. B. and J. G. Maxham III (2012). Return shipping policies of online retailers: Normative assumptions and the long-term consequences of fee and free returns. Journal of Marketing 76(5), 110–124.

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

Dzyabura, D., S. El Kihal, and M. Ibragimov (2018). Leveraging the power of images in predicting product return rates. SSRN Electronic Journal, 1–33.

Kedia, S., M. Madan, and S. Borar (2019). Early bird catches the worm: Predicting returns even before purchase in fashion e-commerce. arXiv preprint arXiv:1906.12128.

Ketzenberg, M. E., J. D. Abbey, G. R. Heim, and S. Kumar (2020). Assessing customer return behaviors through data analytics. Journal of Operations Management 66(6), 622–645.

Ketzenberg, M. E. and R. A. Zuidwijk (2009). Optimal pricing, ordering, and return policies for consumer goods. Production and Operations Management 18(3), 344–360.

Minnema, A., T. H. Bijmolt, S. Gensler, and T. Wiesel (2016). To keep or not to keep: Effects of online customer reviews on product returns. Journal of retailing 92(3), 253–267.

Nevo, A. (2001). Measuring market power in the readyâtoâeat cereal industry. Econometrica: Journal of the Econometric Society 69 (2), 307–342.

Ofek, E., Z. Katona, and M. Sarvary (2011). âbricks and clicksâ: The impact of product returns on the strategies of multichannel retailers. Marketing Science 30(1), 42–60.

OPTORO (2016). Consumer returns report: How flexible return policies can drive customer loyalty.

OPTORO (2018). State of retail returns 2018: Insights on the returns industry.

OPTORO (2022). 2022 state of returns: How good returns processes are driving customer loyaltyâand revenue.

Petersen, J. A. and V. Kumar (2009). Are product returns a necessary evil? antecedents and consequences. Journal of Marketing 73(3), 35–51.

Reagan, C. (2016, Dec). A $260 billion 'ticking time bomb': The costly business of retail returns.

Reagan, C. (2019, Jan). That sweater you donât like is a trillion-dollar problem for retailers. these companies want to fix it.

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

Shehu, E., D. Papies, and S. A. Neslin (2020). Free shipping promotions and product returns. Journal of Marketing Research 57(4), 640–658.

Shulman, J. D., A. T. Coughlan, and R. C. Savaskan (2011). Managing consumer returns in a competitive environment. Management Science 57(2), 347–362.

Shulman, J. D., M. Cunha Jr, and J. K. Saint Clair (2015). Consumer uncertainty and purchase decision reversals: Theory and evidence. Marketing Science 34(4), 590–605.

USA-Census-Bureau (2021, May). Quarterly retail e-commerce sales 1st quarter 2021.

Online Appendix

A Data

In this section, we share additional descriptive statistics of the data set.

Figure A.1: Total and net sales for a sample of female jeans

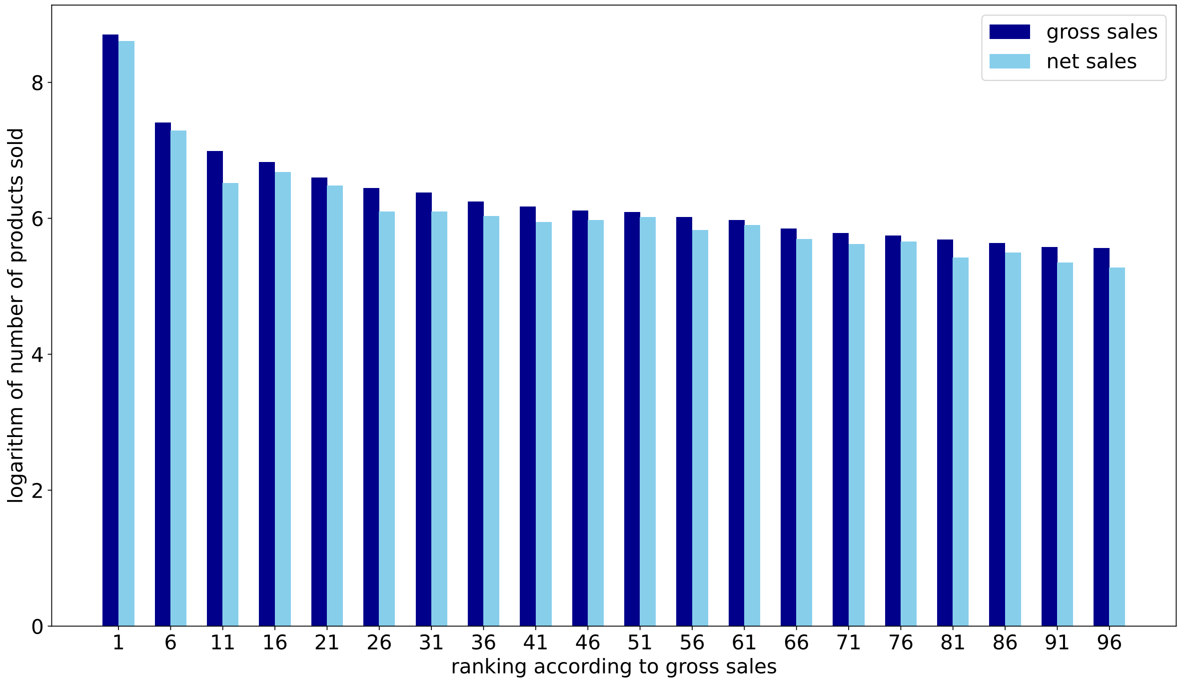

This plot shows the logarithm of the number of sales in y-axis and the product id (ranked according to the number of gross sales) in x-axis. One can notice that the rankings will change if we sort according to net sales even by 10 ranks.

Figure A.2: Return rates by brand

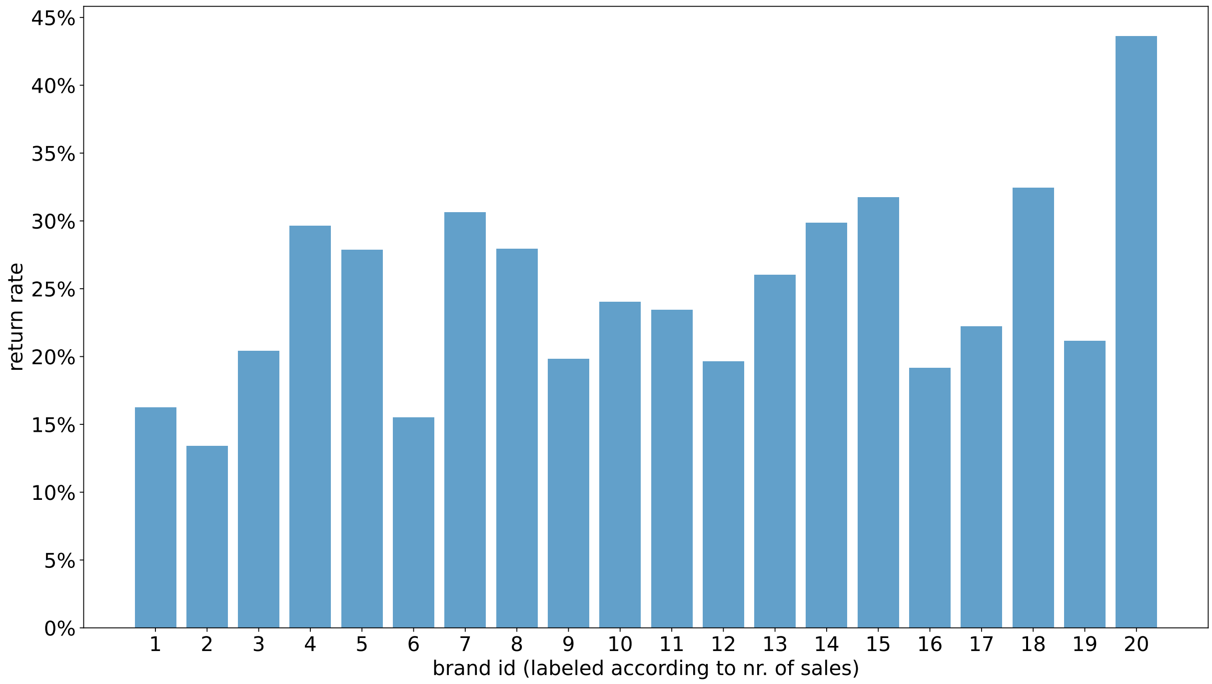

In this plot, we have the average return rates of brands for the 20 best seller brands in female-jeans category.

Table A.1: Return rates by category and gender

Gender	Category	Return Rate
Female	Trousers	17.36%
Female	Coat	19.24%
Female	Pullover	7.83%
Female	Jeans	18.44%
Female	Dress	18.74%
Male	Trousers	14.25%
Male	Coat	20.29%
Male	Pullover	7.3%
Male	Jeans	18.09%

This table is shows the average return rate in gender-category pairs, which will define a market (set of substitute products) in our analysis.

Table A.2: Return reasons

return reason	share
Size mismatch	49.99 %
Dislike product model	9.21 %
Dislike product quality	8.8 %
Unliked product	4.71 %
Abandon	4.6 %
Wrong order	4.46 %
Various reasons of size and quality	3.2 %
Defective product	2.56 %
Order another product	2.25 %
Reason not specified	2.2 %
Delay	1.95 %
No delivery	1.05 %
Choose wrong payment method	1.04 %
Different product	0.89 %
Cancel address change	0.75 %
Cancel economic reason	0.74 %
Various reasons	0.73 %

This table shows the return reasons that are stated more than 500 times. Some common reasons are clickable options in the website and if the customer cannot find a satisfying option, she can write her own explanation. Reasons related to cancellations without receiving the product (wrong payment method, no delivery, wrong order -if less than 3 days- are dropped in the estimation.)

Figure A.3: Dependence of return rates on product ratings by category

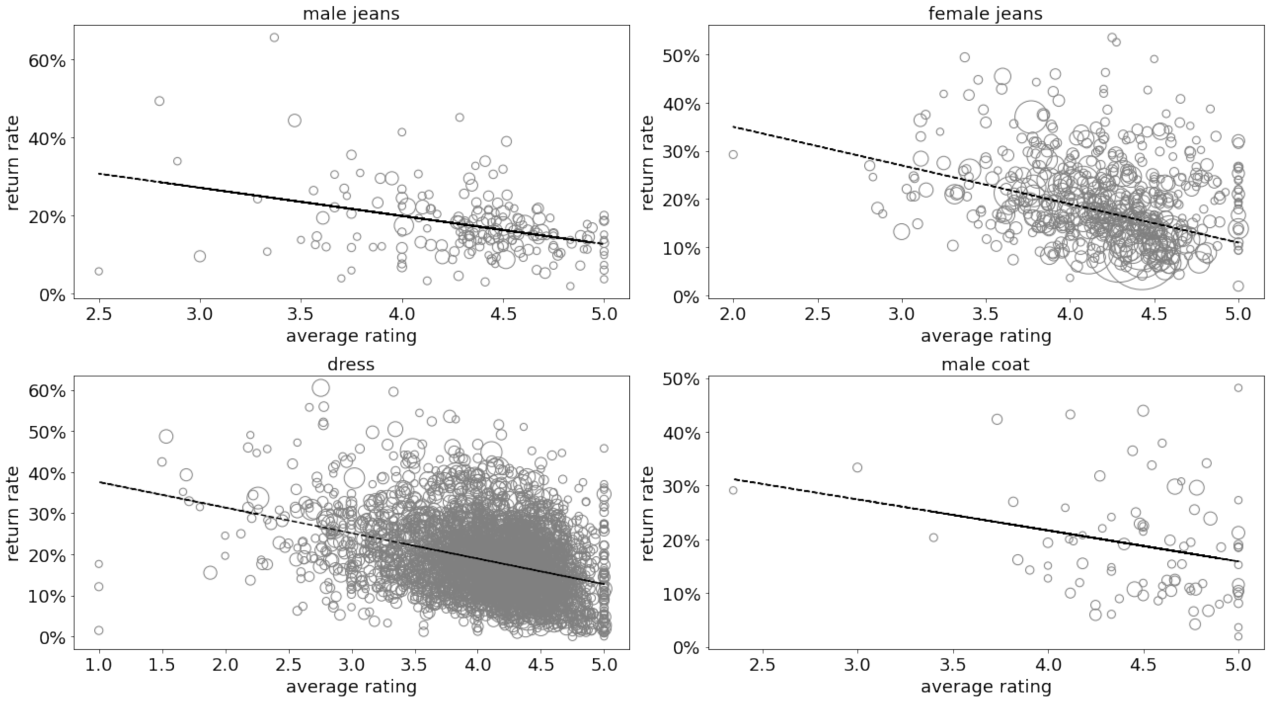

This scatter plot shows the ratings (in $x-$axis) and return rates in percentages (in $y-axis$) with the observation numbers represented by the radius of circles for the markets labeled in their titles. The linear relationship (estimated using weighted least squares) is shown by the dashed line.

Figure A.4: Dependence of return rates on product ratings

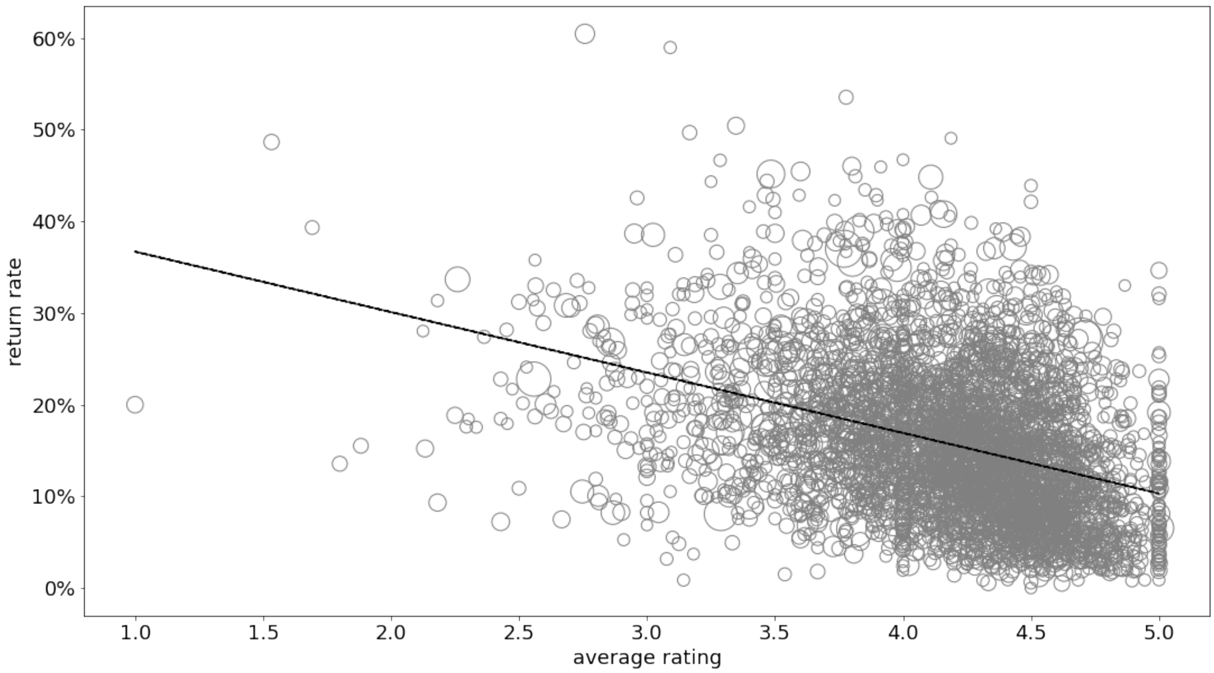

This scatter plot shows the ratings (in $x-$axis) and return rates in percentages (in $y-axis$) with the observation numbers represented by the radius of circles. The linear relationship (estimated using weighted least squares) is shown by the dashed line.

B Simulation of consideration sets

Suppose we have the following observed consideration sets for product 2:

$$\{2, 1, 3, 4\}$$

$$\{2, 3\}$$

$$\{2, 5, 6\}$$

$$\{2\}$$

$$\{2, 5\}$$

$$\{2\}$$

$$\{2, 1\}$$

Based on these sets, we have the following $binom_{2k}$'s:

$$binom_{21} = 2/7, \quad binom_{22} = 7/7, \quad binom_{23} = 2/7,$$

$$binom_{24} = 1/7, \quad binom_{25} = 2/7, \quad binom_{26} = 1/7$$

The probabilities are generally lower in the above example, so it is very unlikely to have a consideration set larger than actual consideration sets. Since $binom_{22}$ is 1, whenever we take a draw from a Bernoulli distribution, it will be 1 and therefore product 2 will be in the consideration set as expected. For other products depending on the result of the draw based on the respective probability, we will include them in the consideration set. Repeating this procedure will yield simulated consideration sets for product 2.

C Optimal pricing

If we take the derivative of the profit function in (2) with respect to the price of item j, we get the following first-order condition for p_j to be optimal.

$$\frac{\partial \pi_f}{\partial p_j} = \sum_{l \in f} \left(\frac{\partial s_l(p)}{\partial p_j} \cdot (1 - r_l(p)) + \frac{\partial (1 - r_l(p))}{\partial p_j} \cdot s_l(p) \right) \cdot (p_l - mc_l) + s_j(p) \cdot (1 - r_j(p)) - \sum_{l \in f} \left(\frac{\partial s_l(p)}{\partial p_j} \cdot r_l(p) + \frac{\partial r_l(p)}{\partial p_j} \cdot s_l(p) \right) \cdot rc_l = 0$$

We can replace the summation notation over the firm's products and cross derivatives of scalars with vector notation and express all first-order conditions in a compact form:

$$\left(\Omega \circ \left(\frac{\partial s(p)}{\partial p} - \frac{\partial (r(p) \circ s(p))}{\partial p} \right) \right) \cdot (p - mc) + s(p) \circ (1 - r(p)) - \Omega \circ \frac{\partial (s(p) \circ r(p))}{\partial p} \cdot rc = 0$$

Rearranging the terms to isolate p yields:

$$p = mc + \left(\Omega \circ \left(\frac{\partial s(p)}{\partial p} - \frac{\partial (r(p) \circ s(p))}{\partial p} \right) \right)^{-1} \cdot \left(-s(p) \circ (1 - r(p)) + \Omega \circ \frac{\partial (s(p) \circ r(p))}{\partial p} \cdot rc \right)$$

Finally, by multiplying the terms within two big parentheses by -1, we get the pricing rule in a comparable way to the original pricing rule as in (4).

C.1 Derivation of Illustrative Example

We have

$$\frac{\partial \dfrac{(s'(p) \cdot (\bar{r} + \alpha_r \cdot p) + s(p) \cdot \alpha_r)}{s'(p) - (s'(p) \cdot (\bar{r} + \alpha_r \cdot p) + s(p) \cdot \alpha_r)}}{\partial \alpha_r} = \frac{(p \cdot s'(p) + s(p)) \cdot s'(p)}{(s'(p) - (s'(p) \cdot (\bar{r} + \alpha_r \cdot p) + s(p) \cdot \alpha_r))^2}.$$

The denominator is always positive, so we have to check the sign of the numerator. We can notice that the first term is the price elasticity of demand plus 1 when it is divided by $s(p)$, hence that part is negative for normal goods. $s'(p)$ is also negative and therefore we see that marginal cost is increasing in α_r.

For the markup we have

$$\frac{\partial \dfrac{-s(p) \cdot (1 - \bar{r} - \alpha_r \cdot p)}{s'(p) - (s'(p) \cdot (\bar{r} + \alpha_r \cdot p) + s(p) \cdot \alpha_r)}}{\partial \alpha_r} = \frac{s(p) \cdot (-\alpha_r \cdot p \cdot s(p) - (1 - \bar{r} - \alpha_r \cdot p) \cdot s(p))}{(s'(p) - (s'(p) \cdot (\bar{r} + \alpha_r \cdot p) + s(p) \cdot \alpha_r))^2}.$$

Again, we have to focus on the numerator and it is clear that the term inside the parenthesis is negative. This means that the markup term shrinks as α_r, the return rate's reaction to price, increases.

D Additional results

Table D.3: Estimation Table with Brand Fixed Effects

	both data sets
price	1.914
	(0.015)
average rating	0.692
	(0.018)
brand 1	-3.52
	(0.1)
brand 2	0.38
	(0.46)
brand 3	-1.82
	(0.14)
brand 4	-4.96
	(0.1)
brand 5	-4.07
	(0.11)
brand 6	-2.00
	(0.13)
brand 7	-4.75
	(0.12)
brand 8	-2.95
	(0.14)
brand 9	-2.27
	(0.15)
purchase week fixed effects	yes
log likelihood	- 87606
nr. observations	118706

Figure D.5: Price elasticity before and after accounting for returns

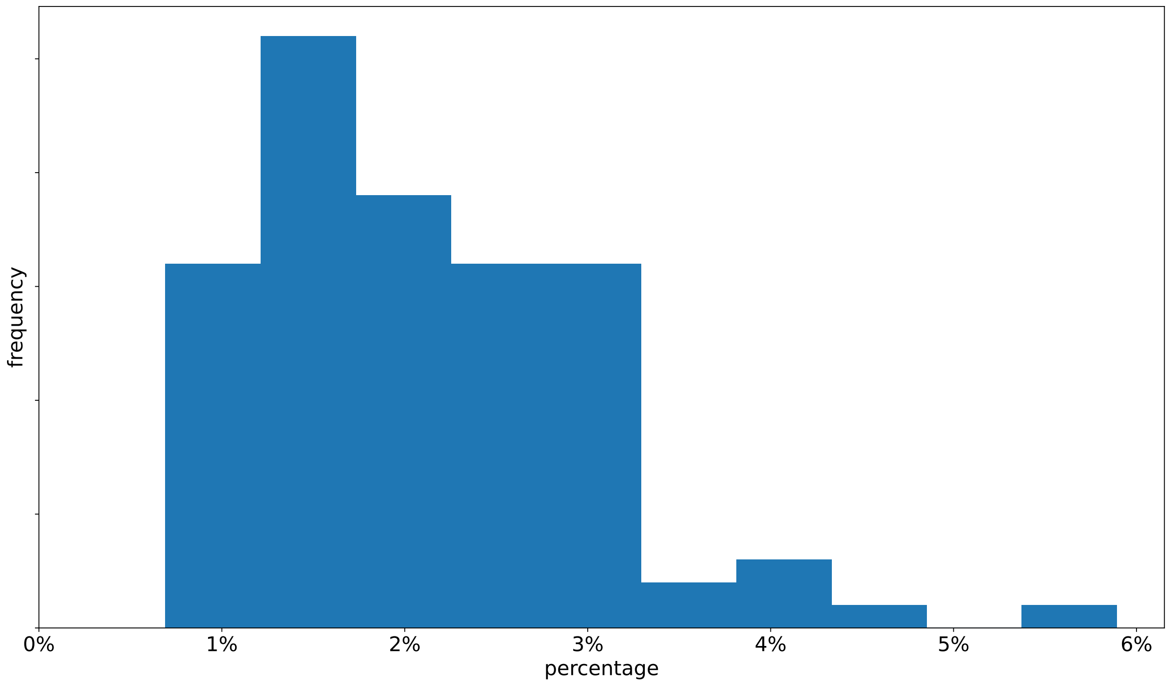

This histogram shows the distribution of the percentage deviation in the net demand elasticity of price with respect to the gross demand elasticity of price.

Figure D.6: Dependence of profit increase on return rate

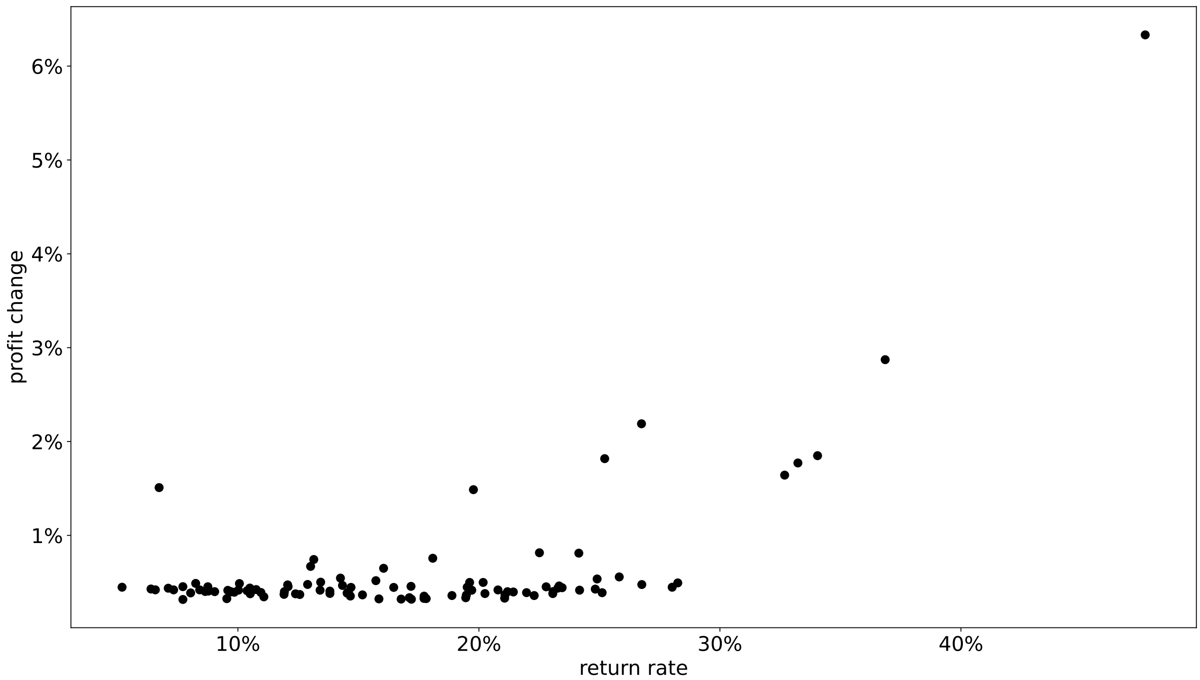

This scatter plot shows the original return rates of the products (in x-axis) and the profit change (in percentage) in the counter-factual equilibrium pricing scenario (in y-axis).

Figure D.7: Dependence of return cost change on return rate

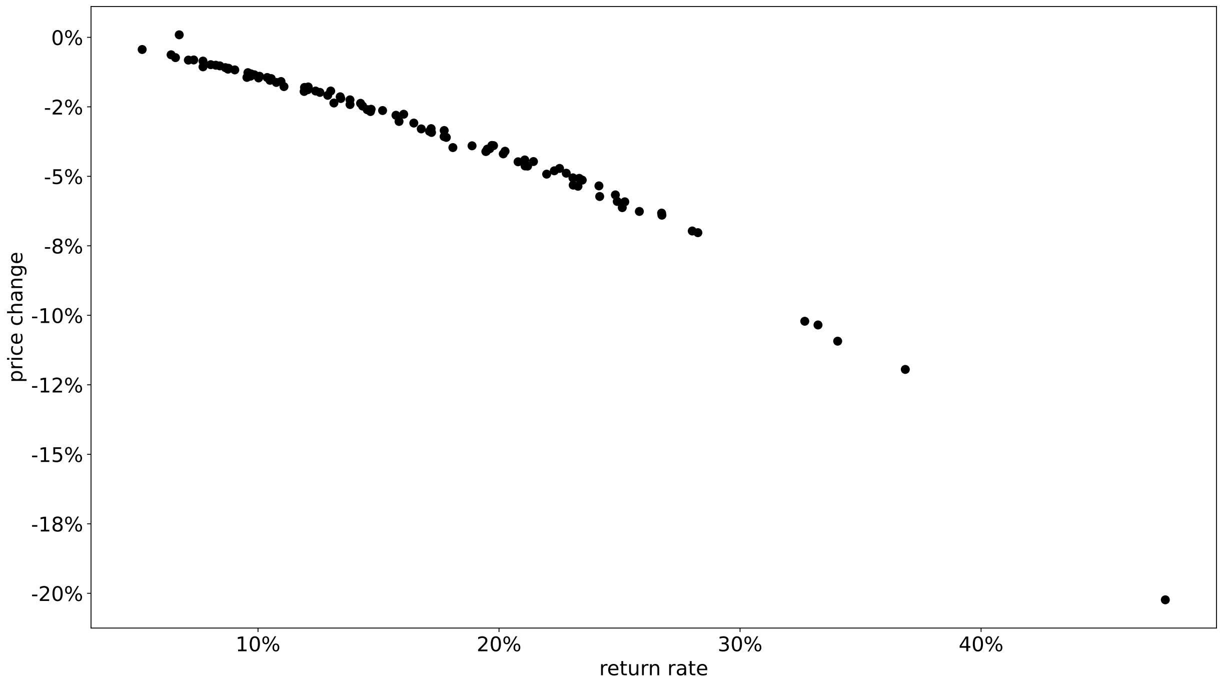

This scatter plot shows the original return rates of the products (in x-axis) and the return cost changes (in percentage) in the counter-factual equilibrium pricing scenario (in y-axis).