Product Returns Management in Online Retail

By

Marat Ibragimov

B.S., Moscow Institute of Physics and Technology (2015)
M.S., Moscow Institute of Physics and Technology (2017)
M.A., New Economics School (2018)
S.M., Massachusetts Institute of Technology (2022)

SUBMITTED TO THE DEPARTMENT OF MANAGEMENT IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN MANAGEMENT

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Authored by: Marat Ibragimov
MIT Sloan School of Management
3 May 2023

Certified by: John R. Hauser
Professor of MIT Sloan School of Management, Thesis Supervisor

Acepted by: Eric So
Sloan Distinguished Professor of Financial Economics
Professor, Accounting and Finance
Faculty Chair, MIT Sloan PhD Program

Product Returns Management in Online Retail

by

Marat Ibragimov

Submitted to the Department of Management
on May 3, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Management

Abstract

In Chapter 1, I and coauthors study the problem of predicting the product return rate using the products' visual information. In online channels, products are returned at high rates. Shipping, processing, and refurbishing are so costly that a retailer's profit is extremely sensitive to return rates. Using a large dataset from a European apparel retailer, we observe that return rates for fashion items bought online range from 13% to 96%, with an average of 53% – many items are not profitable. Because fashion seasons are over before sufficient data on return rates are observed, retailers need to anticipate each item's return rate prior to launch. We use product images and traditional measures available prelaunch to predict individual item return rates and decide whether to include the item in the retailer's assortment. We complement machine-based prediction with automatically extracted image-based interpretable features. Insights suggest how to select and design fashion items that are less likely to be returned. Our illustrative machine-learning models predict well and provide face-valid interpretations – the focal retailer can improve profit by 8.3% and identify items with features less likely to be returned. We demonstrate that other machine-learning models do almost as well, reinforcing the value of using prelaunch images to manage returns.

In Chapter 2, I consider customer search and product returns on the individual level. Previous research has focused on linking customers' purchase and return decisions. However, online retailers have access to the information which precedes the purchase decision – customer search. I demonstrate that customer search information provides important insights about product returns. Using data from a large European apparel retailer, I propose and estimate a joint model of customer search, purchase, and return decisions. I then provide theory and data indicating that using search filters, viewing multiple colors of a product, spending more time, and purchasing the last item searched are negatively associated with the probability of a return. Finally, I use the proposed model to optimize the product display order on the retailer's website.

Chapter 3 extends and reinforces the results obtained from previous Chapters. In the paper, I study the assortment planning problem in presence of frequent product returns. I develop a deep-learning model of customer search, purchase, and return. The model is based on a transformer framework and allows the recovery of important relations in the data. I use the estimated model to demonstrate that retailers could identify successful and unsuccessful products and modify the assortment. The modified assortment would increase the retailer's sales and at the same time decrease returns. Lastly, I provide qualitative insights on which products are most likely to be unsuccessful in online retail.

Thesis Supervisor: John R. Hauser
Title: Kirin Professor of Marketing, Professor of Marketing

Acknowledgments

During my five years at MIT, I was fortunate to meet brilliant scholars and friends.

I would like to express my deepest appreciation to my academic advisors John Hauser and Duncan Simester for their tremendous effort in raising me as a scholar. John and Duncan gave me a fantastic opportunity for academic freedom and were always willing to provide thoughtful comments and critiques on my research ideas. Their counseling strengthened my interest in applied marketing research. I want to especially thank John for guiding me through the review process of my first accepted paper and for teaching me how to shape and reinforce the idea in writing (and for supervising my job market paper!). I want to especially thank Duncan for teaching me the importance of being strategic in research and for giving very practical advice on how to identify a promising research area. John and Duncan are inspirational colleagues and friends, whose example encouraged me to work hard and truly enjoy my time at MIT.

I would also like to extend my deepest gratitude to my friends and co-authors Daria Dzyabura, Siham El Kihal, and Artem Timoshenko for their hard work, deep feedback on my research, thoughtful conversations, and for introducing me to the fun side of academic life! I also thank my friend Anton Popov for our great time spent at the gym and the cyber world.

My research greatly benefited from the research seminars and discussions at the MIT Marketing Group. I would like to thank the marketing faculty members Sinan Aral, Rahul Bhui, Sharmila Chatterjee, Dean Eckles, Renee Gosline, John Little, Drazen Prelec, David Rand, David Schmittlein, Catherine Tucker, Glen Urban, Birger Wernerfelt, Juanjuan Zhang, and my fellow Ph.D. students Jenny Allen, Cathy Cao, Matthew Cashman, Cathy Chen, Jason Du, Lei Huang, Graelyn Humiston, Madhav Kumar, James Li, Keyan Li, Chengfeng Mao, Cameron Martel, Reed Orchinik, Yifei Wang, Jeremy Yang, Shuyi Yu, Jerry Zhang, and Yuting Zhu. Many thanks to Nikolas Hill, Hillary Ross, and Davin Schnappauf for their great administrative support.

I am extremely grateful to my parents: my father Maruf, my mother Nailya, and my brother Renat for their constant support and encouragement during the lightest and toughest times of my Ph.D. study. Finally, I would like to extend my sincere gratitude to my partner Jenn Gu for her patience and tremendous support during my job market, and for the fun time spent together afterwards!

Contents

List of Figures

List of Tables

Chapter 1

Leveraging the Power of Images in Managing Product Return Rates

1.1 Introduction

[1]Online retailers are challenged by the high cost of product returns. Processing and refurbishing the returned item is so costly that large retailers such as Amazon and Walmart allow customers to keep the item, because it often costs more to ship and process the returned product than the product is worth (The Wall Street Journal, 2022). Nick Robertson, founder of the UK's largest fashion retailer, ASOS, stated that a 1% drop in ASOS' return rate could increase the firm's bottom line by an impressive 30% (Emma Thomasson, 2013).

In the $500 billion fashion industry, return rates are high, and vary greatly by item. The products upon which we focus are fashion items. For our focal retailer, a large European apparel retailer, we observe item return rates averaging 53% ranging from 13% to as high

[1]Joint work with Daria Dzyabura, Siham El Kihal and John Hauser

as 96% for some items. This is in contrast with the 3% return rate in the same retailer's offline channel, with the same set of items. Even with high margins, the items on the higher end of this return-rate spectrum generate a net loss for the firm's online store. In fashion, as in many industries, the product return rate is key input into any product management strategy. The problem in the fashion industry is that fashion seasons are short and return deadlines are generous. By the time an item's return rate is observed, the fashion season is well underway or almost over. To effectively manage item assortment in light of returns, it is critical that the retailer is able to predict item return rates using only data available prelaunch.

In this paper, we address this problem by leveraging image processing methods. We demonstrate that item images improve predictions of return rates, that policies based on predictions can improve profit, and that data-based insights are face valid, internally consistent, and suggest which items are returned at high and low rates. To do so, we develop a modeling framework to predict and interpret how product images relate to their return rates. Machine learning models produce accurate predictions of an item's return rate based on features of the product image and other characteristics available prelaunch. For example, including deep-learning image features in gradient-boosted regression trees (GBRT) predicts 13.5% better than a model based on traditional features alone. Using this model and the derived policy to decide on which items to display results in a profit improvement net of returns by 8.3% relative to displaying all items in the online channel. SHAP values (that relate automatically-interpretable image-based features to return rates) suggest how the firm might design (or otherwise source) items less likely to be returned.

We tested a variety of alternative machine-learning models and features to suggest which do well and which do not on our data. Among those tested are deep-learning features, human-coded features, hand-crafted automated pattern and color features, and automatically-generated image-based interpretable features. We find that many machine-learning models do well on our data providing evidence for the value of item images for managing item assortments.

Our contribution is to show that incorporating item images into models helps a firm decide, prior to launch, which products to include in its online store based on profitability net of returns. The approach is fully automated, scalable, and implementable prior to product launch, and an improvement on current practice that does not incorporate product images. The approach has the advantage that it can be easily implemented by a retailer for each fashion collection.

The remainder of the paper is organized as follows. We begin by reviewing relevant literature on managing product returns and leveraging image data (Section 1.2). Section 1.3 describes the data and provides empirical (model-free) evidence that image data predict returns. Section 1.4 demonstrates that image-based features improve predictions, explores alternative models, and develops a model-based policy for selecting which items to display/not-display in the online store. Section 1.5 complements the predictive model with automatically-generated image-based interpretable features which provide insights on how to source and design items. We conclude with a summary, limitations, and suggested future research (Section 1.6).

1.2 Related Literature

We build on and contribute to two streams of literature: managing product returns and leveraging image data.

1.2.1 Managing Product Returns

A rich literature in marketing and operations investigates firm strategies for managing product returns. One such strategy is to manage returns by optimizing the leniency of the return policy (such as fees, prices, or deadlines). Anderson et al. (2009) develop an individual-level model of purchase and return and use it to optimize the return costs for customers. Moorthy and Srinivasan (1995) suggest that a return policy is a signal of item quality. Shulman et al. (2011) show that the optimal policy (strict vs. lenient) balances sales and returns. See also Davis et al. (1998); Wood (2001); Bower and Maxham III (2012); Janakiraman et al. (2016).

Another approach for managing returns focuses on managing customers and understanding their return behavior. For example, Petersen and Kumar (2009, 2015) describe customer return behavior and how it affects future spending and how this can be accounted for in lifetime value calculations to target more profitable customers. Sahoo et al. (2018) study how product reviews decrease return rates by reducing consumers' uncertainty. Other studies link product returns to factors such as prices and price discounts, marketing instruments, free shipping promotions, the use of an app, or even the weather (e.g., Conlin et al. (2007); El Kihal et al. (2021); Narang and Shankar (2019); Petersen and Kumar (2010); Shehu et al. (2020); El Kihal and Shehu (2022). Other than at the broad category level (Hong and Pavlou, 2014), the literature has not explored the characteristics of products related to high return rates.

Rather than focusing on return policies, managing (and "firing") customers, or prices and marketing strategies, our research focuses on the products (items) themselves. Not only is this a gap in the literature, but it is clearly complementary to return policies, managing customers, prices, and marketing strategies. We focus on which items to display/not-display and item features that lead to high or low return rates. Not displaying an item is mathematically equivalent to charging an infinite price. With new data, future research might explore softer strategies such as setting a very high price. Pricing research is not feasible with our data and beyond the scope of this paper. We observe prices and use prices as traditional product features, but we do not observe the demand curve and cannot optimize prices.

1.2.2 Leveraging Image Data

Images have always been an important part of firms' marketing efforts. Over the past two decades, technical advances and the rise of digital platforms have created an abundance of visual data. Together with the development of image processing tools and advanced modeling techniques, these data have created unique research opportunities in marketing. For example, researchers have analyzed images in consumer reviews (Zhang and Luo, 2023), user-generated digital content (Hartmann et al., 2021; Liu and Toubia, 2018; Klostermann et al., 2018; Dzyabura et al., 2021), firm logos (Dew et al., 2022), and seller images on digital platforms (Zhang et al., 2021). For a detailed review of published and ongoing research, see Dzyabura et al. (2021).

He and McAuley (2016); Lynch et al. (2016); McAuley et al. (2015) demonstrated by example that images are valuable for making recommendations regarding clothing styles, substitutes, and personalized rankings. Shi et al. (2021) use machine learning to identify garments and classify fashion-item features from street snapshots, runway photos, and online stores. They use these tools to interpret fashion dynamics and conclude that machine learning can identify fashion features not discussed in fashion magazines. They do not use the fashion features to predict item sales or item returns.

The literature supports that images contain valuable information in many product categories, and specifically in fashion. Although images have not been used to forecast return rates for specific items, they have been proven valuable for other tasks. The literature also suggests that machine learning can identify independent variables ("features") that are used to forecast dependent variables (in our case return rates for each item).

1.3 Data Description and Empirical Evidence that Image-based Data Predict Return Rates

To study product returns, we chose an industry that is particularly challenged by high return rates – women's apparel. We obtain transaction data from a major European mono-brand retailer-manufacturer. We augment the transaction data with a study in which human judges label images from the retailer's two largest categories (dresses and shirts).

1.3.1 Retailer Transaction Data

The women's apparel retailer has a network of 39 retail stores in Germany complemented by a large online operation that accounts for 30.5% of its sales. All items appear in both channels and are always sold for the same price in the two channels. The retailer has a lenient return policy mandated by law: customers can return any purchased item for any reason within 14 days, without providing the reason. By the retailer's policy, items must be returned in the same channel in which they were purchased.

We use data on 1,231,055 transactions, including sales and returns, that occurred in online channels during the observation period from September 1st, 2014 until August 31st, 2016 (two full consecutive years). We exclude non-apparel items such as perfume, gift cards, or accessories. We observe returns for all orders made within the observation period. The data also include offline sales but returns are rare in the offline channel (with a 3% offline return rate for the focal retailer).

For each transaction, we observe the date, the channel (online/offline), item identifiers, and which items were returned. For each item, we observe the category (e.g., dresses), price, and four-to-six images. The images were taken by the same studio, using standardized procedures, resulting in consistent image quality. In our primary analysis, we include only the front image of each item, which is the most informative and always the first image displayed to the customer on the retailer's website. (A model using all images performs only marginally better – Appendix A) The images display the item by itself (not on a model or a manikin) against a white background. We include only items for which we have images (97% of items) and which were sold at least 20 times (Varying this threshold did not change our findings – Appendix A). On average, an Item fills 55% of the image (standard deviation 13%). The fill-rate mostly depends on the item category (for example, dresses are likely to be larger than shirts). Among all the item images, 99% have the same size (2200 x 1530 pixels). Figure 1-1 contains example images of four items.

The resulting data contain 4,585 distinct items from fifteen different apparel categories, as categorized by the retailer. Return rates for items sold via the online channel lie within

Figure 1-1. Examples of Images of Four items

(a) (b) (c) (d)

13–96% (56% average for this subsample, slightly above the overall average of 53%). These rates are well above return rates in the offline channel (3%) where consumers can touch, feel, and try on items. The processing and refurbishing costs of returns from the online channel, our focus, is well above returns from the offline channel (return-cost data are proprietary). For the interested reader, data on consumer characteristics complement our focus on displaying/not-displaying and designing items. For example, evening shoppers return more than morning and daytime shoppers; middle of the week shoppers return less than beginning and end-of-the-week shoppers; and price discounts are positively related to return rates. These data are only observed postlaunch and cannot be used to manage item returns prelaunch, hence postlaunch data are not used by our models. For completeness we provide these data in Appendix A.

1.3.2 Data Augmentation with Human-Coded Features (HCF)

The number of items per fashion season is large and fashion seasons change rapidly. We seek an automated way for the firm to manage assortments. It would be prohibitively expensive, and the firm may not have the time between fashion seasons, to ask humans to code the item images. Nonetheless, to explore whether or not image data are sufficient without human-coded data, we asked human judges to code illustrative fashion features in the two largest categories. These data provide evidence that image features are predictive of return rates and provide a benchmark with which to evaluate the predictive ability of the automatically-generated image-based interpretable features studied in Section 1.5. The automatically-generated fashion features are curated to apply across categories and, hopefully, subsume many specific fashion features such as the human-coded features for dresses and shirts.

We conducted a study in which human judges labeled 2,392 images from the largest two categories of items (dresses and shirts). Four independent judges, blind to the purposes of the study, labeled each clothing item with respect to symmetry (symmetric vs. asymmetric), pattern (solid, floral, striped, geometric/abstract), and additional details (text, metallic/sequin, graphic, lace). Three of the judges coded sleeve length (short, medium, long, sleeveless), and the presence of belts and/or zippers. The human-based label for an image is equal to 1 if the majority of the judges indicate attribute presence, and equal to 0 otherwise. Ties were uncommon and broken by across-item percentages. On average, judges agreed with the majority vote for 91.7% of the judgments.

1.3.3 Model-Free Motivation: Observable Variables and Image Data Relate to Return Rates

Non-image variables. Return rates are related to observable variables. For example, return rates vary by category as illustrated by Figure 1-2a. Dresses, the largest category in our data, are returned on average 72% of the time while cardigans are returned 37% of the time. Seasonality and price are the other variables used by the retailer that are related to returns (evidence in Appendix A).

Color. A minimal use of images is the color of the item. For example, consumers can more easily imagine themselves in common conservative colors such as blacks, blues, and greys, but often want to try fashion colors such as pinks, purples, and pastel colors. To examine whether return rates vary by image data, we begin with the color labels (twelve color bins) that the retailer uses to categorize each item. The bins are not perfect, for example, "pink" includes many shades of pink and a single color does not fully summarize multi-colored patterns or highlights. Nonetheless, Figure 1-2b suggests that color labels are related to return rates.

We will show that color labels augment traditional models based on category, price, and seasonality. We will also demonstrate that we can do even better with more comprehensive image features and models that account for interactions and non-linearities (and that are regularized).

Human-coded features (HCFs). HCFs are not designed to be scalable to all items in all categories for every fashion season. But they are valuable as indicators that image features are related to return rates. Figure 1-3 displays the correlations of the HCFs with return rates. Asymmetrical items are associated with a higher return rate, compared with symmetrical items. Items with patterns (floral, striped, or geometric/abstract) have a lower return rate compared with items without a pattern (solid items). Among the additional details, lace details, metallic/sequin details and belts seem to be associated with higher return rates, while the presence of a zipper and text or graphic details is associated with

Figure 1-2. Online Return Rates by Category and Color – Retailer's Traditional Classifications

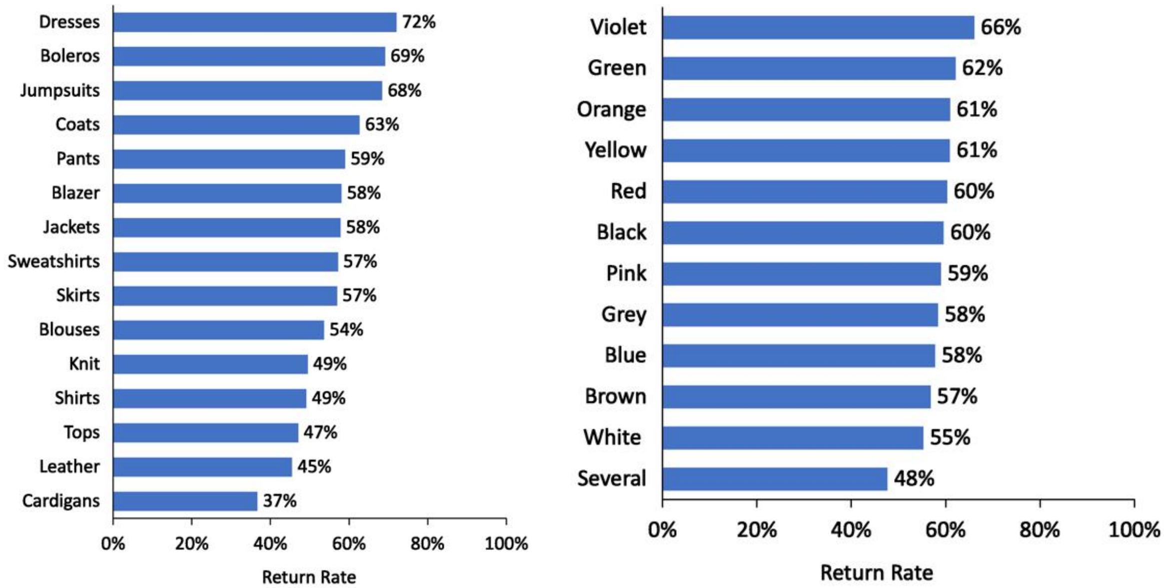

(a) Return Rates by Category Label **(b)** Return Rates by Color

lower return rates. Finally, the length of sleeves is negatively correlated with the return rate. When we regress the item return rate on the HCFs features, we get similar insights. Details are in Appendix A.

Figure 1-3 motivates the hypothesis that image-based features relate to return rates. Likely the relationship is more complex than simple correlations – the HCFs likely interact with each other and with the traditional measures such as fashion category. In the next section, we explore models to handle complex interactions. Because the HCFs do not scale, Section 1.5 develops more general and more comprehensive automatically-generated image-based interpretable features. Model-free evidence, Section 1.5.2, confirms that the interpretable features are correlated with return rates. Section 1.5 also explores model-based methods, called SHAP values, that relate the interpretable features to return rates.

Figure 1-3. Relationship between the Human-coded Features (HCFs) and Return Rates

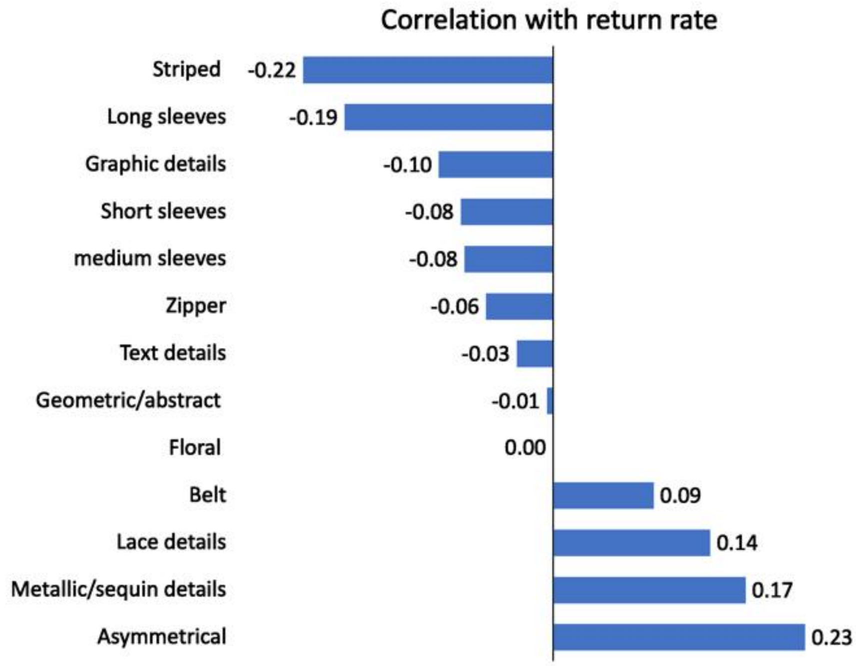

1.4 Prediction: Using Image Data to Predict Return Rates

The previous section suggests that images (image features) augment traditional measures when predicting return rates. We seek a good predictive model to support the retailer's decisions on which items to sell in its online store (display/not-display). For the retailer's policy, we focus on the profitability of individual items rather than the number of units per se. If an item is not profitable, then the retailer's decision whether or not to sell it does not depend upon the forecasted number of units sold. Consistently with the managerial goal, we summarize return behavior by a return rate for each item. The return rate per item varies between 0% and 100%.

Let r_i be item i's return rate, defined as the ratio of the number of returned units

($N_i^{returned}$) to the number of purchased units of the item ($N_i^{purchased}$):

$$r_i = \frac{N_i^{returned}}{N_i^{purchased}} \tag{1.1}$$

To manage product returns using predictions of r_i we make three modeling decisions. The first decision is which features to extract from the images. The second decision is which model to use to predict r_i as a function of the image features and the traditional variables. The third decision is the display/not-display policy used by the retailer – we show the policy is a function of the return rate and the model's predictive ability. We begin by defining the criteria we use to evaluate predictive models. In Section 1.4.6 we argue that our primary criterion is appropriate for the data-and-model based policy that the retailer can use to decide which items to display and which items not to display in its online store.

1.4.1 Criteria to Evaluate Predictive ability

Our primary criterion is the out-of-sample R^2, calculated on all items (i) in our sample (K_{all}). For ease of presentation, we multiply R^2 by 100.

$$R^2_{model} = 1 - \frac{\sum_{i \in K_{all}} (r_i - \widehat{r}_i^{model})^2}{\sum_{i \in K_{all}} (r_i - \widehat{r}_i^{average})^2} \tag{1.2}$$

where $\widehat{r}_i^{model}$ is the out-of-sample item return rate predicted by our model. We use twenty-fold cross-validation to generate out-of-sample predictions for each point in a sample. We randomly divided our sample into twenty non-overlapping folds, where we used 75% of folds to train the model, 20% to validate the model (optimized over a set of hyperparameters,

Appendix A), and the remaining 5% to compute out-of-sample predictions. By assigning different folds to training, validating, and testing the model, the cross-validation procedure allows us to construct out-of-sample predictions for all points in the sample.

To ensure reliable estimates of item return rates, we exclude items that were sold fewer than twenty times. We obtain the same results if we screen the data to require a minimum threshold of either 10 or 30 unit sales. The retailer's decisions are item-by-item because in the online store items are displayed singly and interactions are uncommon. This is one way in which managerial decisions differ from the offline store where items are displayed together and interactions are common. However, the number of purchases for each item does affect the precision of each r_i. The models do not improve when we use $N_i^{purchaesed}$ to weight the data for each item. Details are in Appendix A.

Because other policies might depend on other criteria, we examine the robustness of our methods to different performance measures: we supplement R^2_{model} with mean absolute deviation (MAD) and U^2_{model}. U^2_{model} is a common measure used in marketing that is based on information theory (and probabilities) and measures the amount of (Shannon's) information explained by the model relative to that explainable by perfect predictions (Hauser, 1978). Although derived for classification (0 vs. 1), U^2_{model} applies to more continuous measures such as r_i. It differs from R^2_{model} and MAD because it uses logarithms rather than squared or absolute error. Although derived from information theory, U^2_{model} is sometimes called a pseudo-R^2 . Other classification metrics, such as area under the curve (AUC), are derived for a 0 vs. 1 outcome. The extension of AUC to continuous measures is proportional to MSE and would be redundant with R^2_{model} (Hernández-Orallo (2013), Theorem 7 & Corollary 8).

1.4.2 Baseline Predictions (Item's Category, Seasonality, and Price)

Before we explore the use of images to manage returns, we explore non-image baseline predictions that use information routinely collected by the retailer. For each item in its inventory, the retailer observes the seasonality (month), the item category (e.g., dresses), and price. For price, we use the average price at which the item was sold. Other measures, such as price relative to average category price, do not improve predictions. For the purposes of this analysis, we treat price as exogenous to the decision on whether or not to display an item in the online store. Our data do not contain sufficient information on the demand curve to optimize price. We demonstrate that profitability is improved when price is exogenous. Future research with improved data could include price optimization in policies to improve profitability further.

We can choose a variety of prediction models with which to predict return rates as a function of image and non-image features. These methods vary from simple regression to highly nonlinear functions obtained with machine learning. In our data, we obtain the best predictive ability using gradient boosted regression trees (GBRTs). Bagging methods (random forest) and LASSO do not predict as well as a GBRT, although image-feature-based models using these methods provide incremental predictive ability relative to models based on non-image features alone. Details are in Appendix A.

Table 1.1 reports the predictive ability of the baseline model. To address the variance in the estimated R^2_{model} due to randomness of the division into folds, we generated twenty-five different sets of cross-validation folds (each set including twenty folds); we report the average

Table 1.1. Baseline and Color-Label-Model Predictions.

Model	Non-Image features	Image features	U^2	MAD	R^2	R^2 % change vs. baseline
Non-Image Baseline	Category, seasonality, price	None	52.75 (0.27)	8.59 (0.01)	41.31 (0.18)	–
Color-labels added to baseline	Category, seasonality, price, and color labels	None	53.56 (0.28)	8.48 (0.01)	42.50 (0.20)	+2.88%

Note: Models use LightGBM and differ only with respect to the set of features included. Standard deviations are reported in parentheses. Perfromance measures multiplied by 100 and computed out of sample.

and standard deviations of the estimated R^2_{model}

1.4.3 Improving Predictions with Images (Baseline Plus Color Labels)

Empirical model-free evidence in Section 1.3.3 suggests that color labels are related to return rates. Color labels are minimal image-based features and can be used without image-processing. Table 1.1 shows that color labels improve predictions slightly relative to the baseline. While the improvement is small, the color-label model is further evidence that there is information in images. We show next that deep-learning image features improve predictions substantially beyond predictions obtainable with the retailer's color labels.

1.4.4 Predictions Using Deep-Learning Image Features

Images are more than just color. Consider the three items in Table 1.2 . The first item, the white top, is easily categorized and a common color; the color-label model does well.

Table 1.2. Return Rates and Color-Label Predictions for Three Apparel Items

			
Actual Return Rate minus Color-label Prediction	+1.0%	-12.0%	+15.2%

Note: Actual return rates are not included for confidentiality reasons.

The second item, the top with stripes, is multicolored and hard to categorize by color; the color-label model does less well. The third item, the dress, is readily categorized as pink, but the color-label model does not do well, likely because the pink is not a prototypical pink and because the dress's shape does not work well for everyone.

To improve upon the color-label-based benchmark, we examine image-processing features identified with a convolutional neural network (CNN). In our data, CNN-based features predict best of all tested image-based features. We examine the predictive ability of other image-based features in Section 1.4.5 and more interpretable automatically-generated image-based features in Section 1.5.

Apparel images are often more complex as illustrated by the shape of the pink dress in Table 1.2. Other dresses might feature floral patterns or complex geometric shapes. Deep-learning algorithms have the advantage that they learn feature representations automatically and can be modified for particular applications. To explore the potential of deep learning for image-based predictions of apparel return rates, we use an established CNN. Through a series of nonlinear filters and transformations, the CNN learns highly complex nonlinear transformations to map an image to a set of deep-learning features. The tradeoff is that, while good for prediction, the CNN features are difficult to interpret. The CNN features likely capture the information provided by more-specific features (including the HCFs), but without interpretability, we do not gain insight into which features are associated with high return rates. For greater detail on each transformation and for an application of a CNN to unstructured marketing data, see Zhang et al. (2021).

Our 4,585 images are not sufficient to train a deep CNN from scratch, thus we use the second-to-last pre-output layer of the Residual Neural Network (ResNet; He et al. (2016)). ResNet won the 2015 ImageNet Large Scale Visual Recognition Challenge and was trained on the ImageNet data set (1.3 million images in roughly 1,000 categories). The ResNet network has 152 layers, making it one of the deepest networks yet presented on ImageNet. The second-to-last layer of the network contains 2,048 features. (The last layer is the output layer.) The 2,048 deep-learning features were used directly in the GBRT.

In Table 1.3, we see substantial improvement when using deep-learning features relative to the baseline and color-label models. This improvement in predictive ability leads to a substantial improvement in profit when using the derived display/not-display strategy (Section 1.4.6). Because the 2,048 deep-learned features are likely to encode well the image information, we expect little or no improvement when we add other machine-learned features to the deep-learning features (see next section). Returning to the images in Table 1.2, the GBRT based on deep-learning features predicts return rates better for the hard-to-predict items. The GBRT/CNN model predicts a return rate for the striped top within 4.7% of the true rate (color-label predictions are within 12%), and a return rate within 5.6% for the pink dress (color-label predictions are within 15.2%).

Table 1.3. Predictions Using Deep Learning Image-Processing Features.

Model	**Non-Image features**	**Image features**	$\mathbf{U^2}$	**MAD**	$\mathbf{R^2}$	$\mathbf{R^2}$ **% change vs. baseline**
Non-Image Baseline	Category, seasonality, price	None	52.75 (0.27)	8.59 (0.01)	41.31 (0.18)	–
CNN Features	Category, seasonality, price, and color labels	Deep -learning	56.70 (0.26)	8.15 (0.02)	46.88 (0.19)	+13.48%

Note: Models use LightGBM and differ only with respect to the set of features included. Standard deviations are reported in parentheses. Perfromance measures multiplied by 100 and computed out of sample.

The ResNet CNN is not the only image-processing model that does well on our data, but it is the best of those tested. For example, the VGG19 CNN does almost as well with an $R^2_{model} = 46.84\%$. This and the robustness analyses (Section 1.4.5) suggest that the $R^2_{model} = 46.84\%$ from ResNet CNN is hard to beat. If sufficient data were available, future research might try a custom deep-learning model to squeeze out slightly more predictive ability.

The results in Table 1.3 represent the performance of the model estimated on the entire data set including all product categories. We also explored per-category models for categories with sufficiently many sales and found that predictions in all five largest categories benefit from images, e.g., predictions in "Shirts" improves from $R^2_{model} = 14.78\%$ to 24.08%; predictions in "Dresses" improves from $R^2_{model} = 28.11\%$ to 31.13%. All per-category models are in Appendix A. Machine-learning models are notoriously data hungry – the best predictions are obtained with a model that merges data from all categories. Overall, per-category models reinforce the value of images, but the most managerially-relevant results come from a model that uses the entire data set to train the machine-learning models.

1.4.5 Robustness of Incremental Predictive Ability Due to Image Features

Robustness of the basic hypothesis. Our basic hypothesis is that automated image-based features improve predictions and enable retailers to make more profitable display/not-display decisions. Table 1.3 is based on a particular set of image-based features (CNN) and a particular predictive model (GBRT). We have already summarized that for our data (1) the GBRT predicts best, but other machine-learning models are feasible, (2) alternative deep-learned image-based features are feasible, (3) the results are robust to evaluative criteria, (4) robust to precision weighting by the number of units purchased per item and (5) alternative data screening (minimum threshold of 10 or 30 rather than 20 items). The basic insights also hold for (6) dimensionality reduction of the 2,048 CNN features with various forms of principal component analysis (PCA) and (7) measures of uniqueness and distance from prior fashion seasons – details in Appendix A.

(Black box) automated pattern & color features. We test one more level of robustness. Researchers in machine learning often use automated pattern & color features as an alternative to deep-learning image-based features. Such color & pattern features might improve return-rate predictions. RGB color histograms provide one popular automated color feature. Figure 1-4 Illustrates an RGB coding of the color of an example fashion item as heavily based on red, but with mid-level peaks in green and blue. The number of bins in Figure 1-4, $256 \times 256 \times 256 \simeq 16$ million, is too large for a GBRT. For feasibility we

Figure 1-4. Example RGB Color Histogram Encoding of an Apparel Item

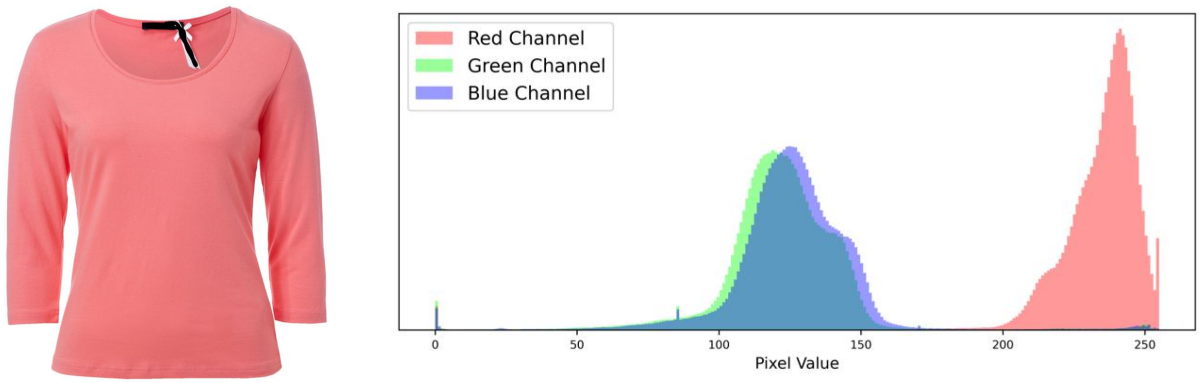

use $5 \times 5 \times 5 = 125$ bins. To capture pattern, we use Gabor filters. Gabor filters use frequency-domain transforms to isolate the periodicity and the direction of that periodicity with sinusoidal waves (Manjunath and Ma, 1996). Although Gabor filters are difficult to interpret, they might improve prediction. See Liu et al. (2020) for an application.

These automated pattern & color features do not predict as well as CNN-based features $R^2_{model} = 45.28\%$). Adding these features to a model based on CNN features and color labels is redundant (does not improve predictions, $R^2_{model} = 46.84\%$). Alternative automated color features (HSV features, ORB features) do not change the basic message – details in Appendix A. Although automated pattern & color features provide an alternative to CNN features in the predictive model, they do not greatly enhance interpretability. We examine more interpretable features in Section 1.5.

Results for human-coded features (HCF). The HCFs improve predictions relative to the non-image baseline (6.85% for dresses and shirts), but do not predict as well as the models with automated CNN-based features (9.92% dresses and shirts)[2]. Given the added time and cost of HCFs, the CNN features appear to be a better choice for the predictive model. The

[2]We re-estimated all models for the two categories for which HCFs were coded (dresses and shirts). The absolute predictive ability, but not the relation among models, varies when we limit the data to the two largest categories.

interpretable features in Section 1.5 are curated to generalize the HCFs. Section 1.5 suggests they predict better than the HCFs.

Summary of robustness tests. The GBRT/CNN model appears to be robust to alternative predictive models, alternative deep-learning image-processing features, alternative performance metrics, alternative data cleaning, dimensionality reduction, and the use of automated pattern & color features. The GBRT/CNN model appears to be a reasonable proof-of-concept. Its predictive ability does not seem to be due to chance. It is of course possible that some retailers will adopt alternative models or image-based features for reasons outside our analysis. The performance of many alternative models reinforces the basic hypothesis that image-based features help manage returns.

1.4.6 The Relationship between a Model's Predictive Accuracy and Profitability

Because all items are already inventoried for the bricks-and-mortar stores (Section 1.3.1), the marginal fixed costs for displaying the items online are minimal. As long as we do not compromise overall variety, we can consider removing (not displaying) items that have negative expected profits based on the predicted return rate, $\widehat{r}_i^{model}$, and a measure of the uncertainty in $\widehat{r}_i^{model}$. For the remainder of this section, we simplify notation and use $\widehat{r}_i$ as shorter notation for $\widehat{r}_i^{model}$. r_i continues to denote the true return rate.

We make online display/not-display decisions item by item. Because our estimate,$\widehat{r}_i^{model}$, is independent of the number of items sold, $N_i^{purchased}$, and because fixed costs are negligible for displaying an item, we focus on profit per item sold.

This section evaluates whether we should not display items with negative expected profit per sale or whether we should take the precision of the model into account. For example, perhaps we should be more aggressive with a model that predicts better. On the other hand, if predictions were no better than random noise, perhaps we should be more cautious about not displaying any items.

There are inventory costs for carrying an item, but those are well-studied, present no new insight, and can easily be added to the profit-maximizing model. The firm's allowable-return policies are set by law and considered fixed for this analysis. To a first order, we ignore interactions among items and assume that the small percentage of items removed does not affect the demand for the remaining items. Fortunately, our derived policy removes a small fraction of items, but it is a boundary condition of the policy. To the extent that removing some items increases demand for other items, the increased demand improves profits further. To the extent that removing some items decreases overall demand for the online store, profits might decrease – an issue we do not have the data to address.

Naïve policy that ignores uncertainty in predictions. The return costs for returned items consist of two components: a flat processing cost for shipping and handling the return (c_{fix}) and a cost that is proportional to price of the item (c_{var}) because returned items must be discounted or discarded if they are damaged or out of season. Let p_i be the price of item i and c_i be its cost. If r_i were predicted perfectly, the profitability of item i, π_i, would be given by:

$$\pi_i = (1 - r_i)(p_i - c_i) - r_i(c_{fix} + p_i c_{var}) \tag{1.3}$$

The retailer's exact costs are proprietary. For illustration, we used a fixed return cost

Figure 1-5. Distribution of Items' Profitability in Our Data

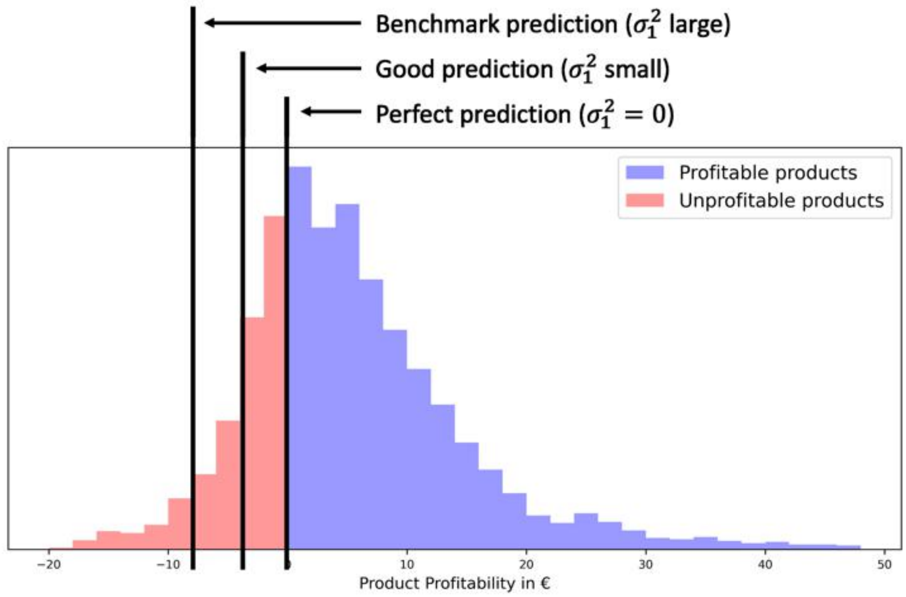

c_{fix} of 5.31€ (iBusiness 2016) and a variable return cost c_{var} of 13.1% of an item's price (Asdecker, 2015). The naïve policy using these illustrative costs would imply that 27.2% of the items in our data are unprofitable. The naïve policy would remove these items. However, the predicted profit from the naïve policy is not achievable because our predictions are uncertain. The naïve policy would overestimate true profits as a 25% improvement. The naïve policy would also violate the assumption that a small percentage of items is not displayed. Figure 1-5 illustrates this naïve policy. "Profitable" products are retained (shown as blue in Figure 1-5) and the "unprofitable" products not displayed (shown as pink in Figure 1-5).

Policy taking uncertainty into account. Our empirical model produces imperfect predictions of the return rates, $\widehat{r}_i$. We assume that $\widehat{r}_i$ is unbiased. This implies that our estimate of profitability $\widehat{\pi}_i$ has a mean equal to the true profits with a variance based on the uncertainty in the predicted return rate. In symbols, $\widehat{\pi}_i|\pi_i \sim \mathcal{N}(\pi, \sigma_1^2)$. We examine whether the retailer's decision depends on the ability of the model to accurately predict return rates, that is, we examine whether the policy depends on σ_1^2.

Because decisions are made for each item, i, we temporarily drop the item subscript, i. We assume the retailer's prior beliefs about profits are normally distributed across items: $\pi_i \sim \mathcal{N}(\mu_0, \sigma_0^2)$. Let $\mathcal{P}$ be a policy such that the retailer displays the item if $\mathcal{P} = 1$ and does not display the item if $\mathcal{P} = 0$. Let $\phi = (\widehat{\pi}, \mu_0, \sigma_0^2, \sigma_1^2)$, then the uncertainty-dependent policy is based on solving the following mathematical problem:

$$\max_{\mathcal{P}(\phi)\in[0,1]} \mathbb{E}[\mathcal{P} \cdot \pi + (1 - \mathcal{P}) \cdot 0] \tag{1.4}$$

The policy that maximizes the mathematical expression in Equation (1.4) is a threshold policy given by Equation (1.5). Equation (1.5) yields intuitive policies as σ_0^2 and σ_1^2 approach zero (perfect information) or infinity (no information) and that expected profits using the policy decrease in σ_1^2[3].

$$\mathcal{P}(\phi) = \begin{cases} 1 & \text{if } \widehat{\pi} \geq \mu_0 \frac{\sigma_1^2}{\sigma_0^2} \\ 0 & \text{if } \widehat{\pi} < \mu_0 \frac{\sigma_1^2}{\sigma_0^2} \end{cases} \tag{1.5}$$

Assuming that the retailer has positive priors, we added the thresholds for uncertainty-based policies to Figure 1-5. (1) For perfect predictions ($\sigma_1^2 = 0$), launch all items for which $\widehat{\pi} > 0$. (2) For good predictions (σ_1^2 small), launch only items for which $\widehat{\pi} > 0$ exceeds the threshold. And (3), when predictions are extremely noisy (σ_1^2 large), launch almost all items even those with expected negative profits. As predictive uncertainty σ_1^2 increases, the uncertainty-based policy screens out fewer items and achievable profits decline. The

[3]Derivation of the threshold policy, a proof that expected profits decrease in σ_1^2, and limiting cases as σ_0^2 and σ_1^2 approach zero or infinity are provided in Appendix A

Table 1.4. Expected Profit Improvement Using Different Predictive Models

Model	Features	Percent Items not Launched	Profit % change vs. Launch All Items
Non-Image Baseline	Category, seasonality, and price	5.98% (0.11)	6.81% (0.18)
Color labels added to baseline	Category, seasonality, price, and color labels	6.26% (0.13)	7.16% (0.19)
CNN Features	Category, seasonality, price, CNN from image	7.13% (0.12)	8.29% (0.23)

dependency on σ_1^2 motivates MSE and R^2_{model} as appropriate criteria with which to judge the predictive model. The better the R^2_{model}, the better is the achievable profit.

Using our data, we simulate the model-based policies (see Table 1.4). When the GBRT/CNN model is used to determine the data-based display/not-display policy, the retailer chooses not to launch 7.13% of the items. The expected profits increase by 8.29% relative to launching all the items. Even compared with the non-image baseline, the improvement in profits is important to fashion retailers with many items in many categories over many fashion seasons. This is especially true for fashion items that are high-priced and high-volume. The potential for profit improvement is even greater if retailers were able to source and/or improve items at the design stage. To that end, we next examine interpretable features, both image and non-image, that are associated with high and low return rates.

1.5 Generating Interpretable Insights

The retailer might improve its profits further if it were to use information available in images to make decisions when sourcing or designing new fashion items. To help the retailer's buyers source items and to help the retailer's designers design new items, we complement the predictive model with an interpretable model that identifies item features that are linked to high and low return rates. To deal with large assortments and rapid fashion seasons, we seek automatically-extracted image-based interpretable features that do not require consumer tests, surveys, or experiments. When the retailer can invest in HCFs, the HCFs enhance interpretability to the extent they help buyers and designers visualize the image-based features.

1.5.1 Automatically-extracted Image-based Interpretable Features

Each of the proposed image-based features is based on insights, experience, and expectations from the fashion industry. We seek features that are interpretable by the retailer's buyers and designers but can be generated at scale automatically. Automatic generation allows the retailer to use the features for large assortments in every fashion season. We extract features related to color (color clusters, color dominance, brightness, horizontal and vertical color asymmetry), pattern (pattern direction, pattern complexity), shape (shape asymmetry, shape ratio, shape triangularity), and item uniqueness (uniqueness). These features are chosen to be as general as feasible and, hopefully, subsume more specific image-based features such as the HCFs. We describe each of these features in detail below and provide examples in Figure 1-6. These features illustrate the type of automatically-generated features that are feasible. At the end of this section, we examine whether this set of features captures sufficient variation in return rates.

Color clusters. To visualize the basic color composition of an item, we use weighted K-means clusters in RGB-pixel space. For each item's image, we calculate the proportion of pixels closest to the mean of the color cluster (thirty clusters in our data). Unlike the retailer's color labels, the color clusters are more-nuanced and data-driven.

Color dominance. Some items have many colors but none dominate; other items have a dominant color with patterns, say flowers, of different colors. Color dominance is the maximum value of a color share for the item.

Brightness. The perceived brightness of an apparel item affects sales and return rates. Brightness might be partially redundant with color clusters, but that is an empirical question. **Brightness** is defined as the average intensity of the image after converting it to greyscale. Brightness varies over a garment. For example, if an item has a uniform color, the brightness variation is close to zero; if the item has a complex pattern of light and dark stripes, the **brightness variation** is larger. Computationally, we use the standard deviation. Both brightness and brightness variation are allowed to enter the model. Figure 1-6a illustrates fashion items with low and high brightness and brightness variation.

Pattern direction. Gabor pattern features had moderate success in predicting returns (Section 1.4.5), but they are very difficult to interpret. Pattern direction and pattern complexity are more interpretable. Pattern direction is summarized by applying a Sobel filter to each direction (X for horizontal and Y for vertical) to the greyscale images Gonzales and Woods (2018). This is equivalent to a partial derivative with respect to movement orthogonally along either the horizontal or vertical axis. For example, horizontal stripes have a high derivative in the vertical direction and vertical stripes a high derivative in the horizontal direction (see Figure 1-6b).

Figure 1-6. Examples of Automatically-extracted Image-based Interpretable Features

(a) Illustration of Items with High and Low Brightness and Brightness Variation

Item with high brightness | Item with low brightness | Item with high brightness variation | Item with low brightness variation

(b) Illustration of Pattern Direction (Sobel X- and Y- Directions)

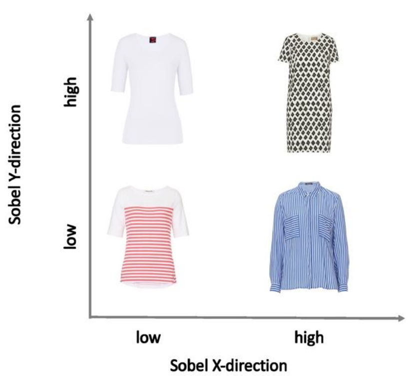

(c) Illustration of Pattern Complexity

Item with high pattern complexity | Item with low pattern complexity

(d) Illustration of Shape Ratio

Dress with high Shape ratio | Dress with low Shape ratio | Shirt with high Shape ratio | Shirt with low Shape ratio

Pattern complexity. Some apparel items have checkered patterns (high derivative in both the horizontal and vertical directions), while others have more complex patterns. To represent pattern complexity, we extract edges from the image using the Canny edge detector and we extract straight lines using Hough transformations (Duda and Hart, 1972). Each line is represented by the orthogonal distance from the top left corner of the image to the line and by the angle of the line relative to the X-axis. Two features are extracted: pattern complexity is the standard deviation of the angles of the extracted lines; the other feature is the number of extracted lines. Pattern complexity is extracted if there are more than twenty lines, otherwise it is set to zero. All such meta-parameters are tuned. Figure 1-6c illustrates (a) an item with high pattern complexity (lines of varying angles) and (b) an item with low pattern complexity (horizontal stripes with a zero angle).

Asymmetry. The HCF analysis suggests that asymmetric items have higher return rates (review Figure 1-3). Shape asymmetry, horizontal color asymmetry, and vertical color asymmetry are likely to affect return rates. Dresses, shirts, and other apparel items are naturally asymmetric vertically. To extract shape asymmetry, we compare the left half of the image to the mirror image of the right half of the image. The percentage of non-overlapping pixels indicates shape asymmetry. For example, if the item is perfectly symmetric horizontally, then there will be no non-overlapping pixels; if the fashion item is highly asymmetric, there will be many non-overlapping pixels. To extract horizontal color asymmetry, we use KL-divergence to compare the RGB histograms for the right and left halves of the image. Vertical color asymmetry compares the top and bottom halves.

Geometric shape. In fashion, the shape of an item is likely to be important in predicting return rates. For example, long dresses are often bought for more formal wear where fashion fit might be extremely important, while shorter dresses are bought for more casual wear where the consumer is less discerning. Shape ratio, the ratio of median width to the median height, captures both sleeveless and item-length phenomena. Because the GBRT allows interactions between the shape ratio and category, the impact of this variable can vary by category such as dresses (length matters more) versus shirts (sleeves matter more). Figure 1-6d provides examples of high and low shape ratios for dresses and shirts. Shape triangularity, the ratio of the median width of the bottom 25% of the item to the median width of the top 25% of the item, differentiates many fashion items. For example, an A-line dress has high shape triangularity while a pencil dress has a shape triangularity close to 1. Because triangularity is easy to visualize, for brevity we do not provide examples in Figure 1-6.

Uniqueness. Uniqueness might contain information not otherwise captured by the automatically-extracted image-based interpretable features. For consistency with the GBRT /CNN model, we define uniqueness as the Euclidean distance between the CNN-learned features of the item and the category mean of the CNN-learned features. Although uniqueness did not improve the GBRT/CNN predictive ability, the lack of improvement may have been because the CNN features already contain a (black-box) measure that captures uniqueness.

1.5.2 Model-Free Evidence Motivates the Use of Interpretable Features

Before we examine formal models, we examine whether or not the proposed automatically-extracted image-based interpretable features are related to return rates. Figure 1-7a reports the correlations between return rates and the interpretable features (other than color clus-

Figure 1-7. Model-Free Evidence: Correlations of Return Rates and Interpretable Image-based Features

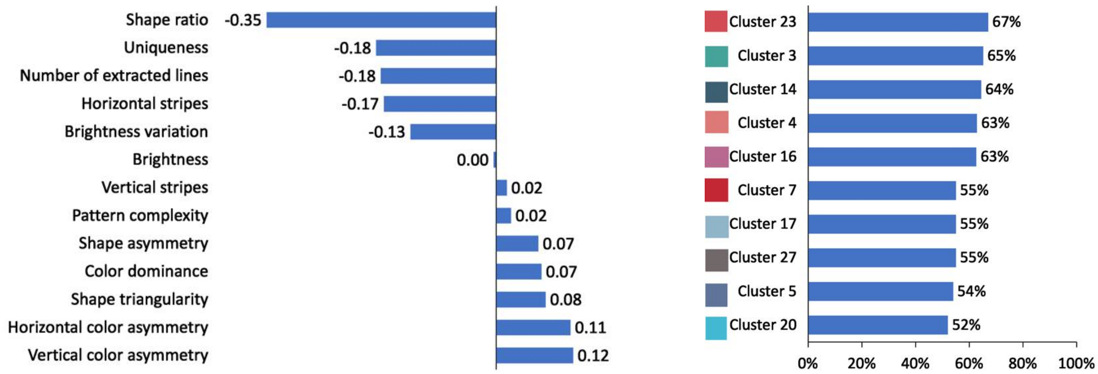

(a) Correlations of Return Rate and features **(b)** Return Rate by Color Clusters

ters). Figure 1-7b reports the return rates for the top five and bottom five color clusters. These model-free analyses motivate a more-complex machine-learning model[4].

1.5.3 Summarizing the Marginal Effect of the Interpretable Features

Our analyses in Section 1.4 suggest that features (image and non-image) interact and that a GBRT (or another machine-learning model) is a good model with which to predict return rates. For comparison to the CNN-based predictive model, we estimate a GBRT model with automatically-extracted image-based interpretable features added to non-image features. Interpreting the impact of features in a tree-based model with hundreds of trees is challenging. The machine-learning literature uses the SHAP (SHapley Additive exPlanations) framework to interpret the marginal impact of each feature on the predicted target variable (Lundberg and Lee, 2017). The SHAP value is based on Shapley values from game theory and enables us to interpret feature impacts in an arbitrary black-box model. The

[4]Note: For illustration, we use in (b) the largest color cluster present in the product image.

SHAP value for feature j for item i (denoted as ϕ_{ij}) indicates the marginal change in the predicted return rate $\widehat{r}_i$ due to a change in the value of interpretable feature j while taking into account all other features in the model. Mathematically, the predicted value $\widehat{r}_i$ of the model for item i could be decomposed as $\widehat{r}_i = \overline{r} + \sum_j \phi_{ij}$ (where $\overline{r}$ is the average return rate for all items). By computing SHAP values for all items, we obtain a sample of SHAP values that can interpreted as the marginal impact on predicted $\widehat{r}_i$ given a random set of all feature values.

To determine the relative importance of each interpretable image feature, we use the mean absolute SHAP value, $F_j = I^{-1} \sum_i |\phi_{ij}|$, where I is the number of items and we sum ϕ_{ij} over all items i for feature j. Intuitively, F_j measures how far on average the given feature j pushes the predicted value of $\widehat{r}_i$ from the sample mean $\overline{r}$. For ease of interpretation, we use Pareto charts that rank the features by F_j and display the most impactful features first.

To illustrate the use of SHAP values, Figure 1-8 ranks the color cluster centers by their impact on the predicted return rate, F_j, from highest to lowest, measured by the average absolute SHAP value within the cluster. Figure 1-8 provides more nuanced interpretations on an item's color composition than do retailer pre-defined color labels. For example, prototypical red has a high impact, but other shades of red have a low impact. Some shades of blues have a high impact, but the prototypical blue has a low impact.

1.5.4 Automatically-Extracted Image Features and Return Rates

Figure 1-9 provides the Pareto Chart for the automatically-extracted image-based interpretable features. We address non-image features in Section 1.5.6. Figure 1-9 reports

Figure 1-8. The Effect of Color Clusters on Return Rates (Clusters Ranked by F_j)

mean-absolute SHAP values for all features. We aggregate the impact of color clusters by the sum of their SHAP values. Appendix A provides more detail on color clusters.

The F_j do not indicate the direction of the impact of a feature, nor do the F_j illustrate variation across items. For example, a feature may have the same F_j value if it is high on a few items and low on many, or just moderate for all items. Figure 1-9 complements F_j with the correlation between SHAP values and standardized feature values to indicate the direction of impact on predicted return rate and to suggest whether the feature affects many items (high correlation) or just a few items (low correlation). See Appendix A for the directionality and variation of impact of color clusters.

Consistent with experience in fashion apparel, color clusters have the greatest impact on return rates. Shape ratios are the next most important and the direction is as expected. More formal items (lower shape ratio) have higher return rates than more casual items (higher shape ratio). As expected, shape ratio has different interpretations for different

Figure 1-9. Pareto Chart of SHAP Values for the Automatically-Extracted Image Features

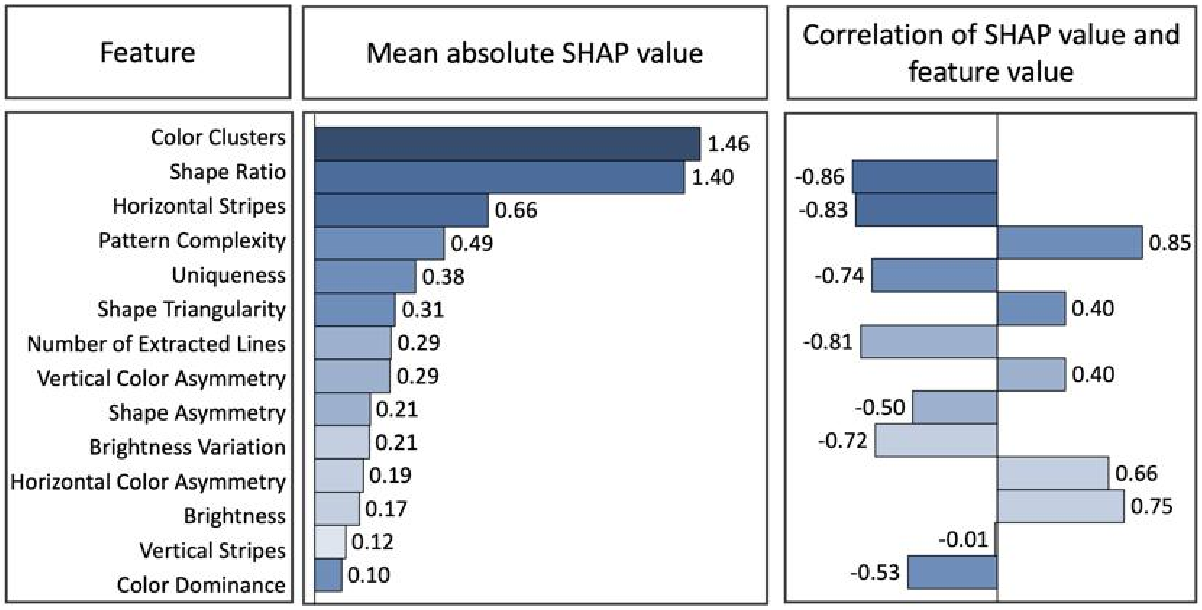

categories (captured in Figure 1-10 below). Sleeveless dresses (lower shape ratio) are returned more often, consistent with the implications of the HCFs (see Section 1.3.2). Interestingly, horizontal stripes are important and are associated with low return rates, while vertical stripes are much less important. Pattern complexity is important and positively correlated with return rates, while uniqueness is negatively correlated with return rates. Uniqueness was redundant with the CNN features in the predictive model (Section 1.4.5), but provides incremental predictive ability in a model with interpretable features (Figure 1-9).

The brightness features are less important, likely because some brightness information is extracted by the color clusters. However, bright products have higher return rates while products with a high variation of brightness (many contrasting colors) have lower return rates. Interestingly, brightness has low impact on predicted return rate, but high correlation. When we examine variation among items (Figure 1-10), this is explained because the SHAP values tend to be small in magnitude, but consistent in their impact on return rates.

To provide further insight to buyers and designers about the variation in SHAP values within items, we use a method known as "bee-swarm charts" to visualize the SHAP values, ϕ_{ij} for all items i for all features j. A bee-swarm chart details, for each item, the impact on return rate predictions of high vs. low values of the feature. Features are ranked by the mean absolute SHAP values, F_j.

Figure 1-10 provides the variation in impact (bee-swarm chart) for the automatically-extracted interpretable image-based features. For example, on average items with higher values of the shape ratio are less likely to be returned, but this relation is not homogeneous. Buyers and designers can examine the detailed points, each of which corresponds to an item, to determine the impact of that item's shape ratio for that item. This can be done for any of the automatically-extracted interpretable image-based features, including color clusters. (Appendix A provides bee-swarm charts for color clusters.)

1.5.5 Using the Interpretable Features to Source or Design Fashion Items

By combining the insights from Figures 1-6 to 1-10, we can predict items that are likely or not likely to be returned. For example, shirts with higher shape ratios, horizontal stripes, and darker colors (red dots to the left in the bee-swarm chart) are less likely to be returned. Shirts with lower shape ratios, solid colors (no horizontal stripes or patterns), and a pinkish color (cluster 24) are more likely to be returned (blue dots to the right in the bee-swarm chart). Examples of such shirts and dresses are shown in Figure 1-11. For confidentiality, we do not provide the predicted or actual return rates, but they are consistent with the

Figure 1-10. The Impact of Automatically-Extracted Interpretable Image-based Features (Bee-swarm chart)

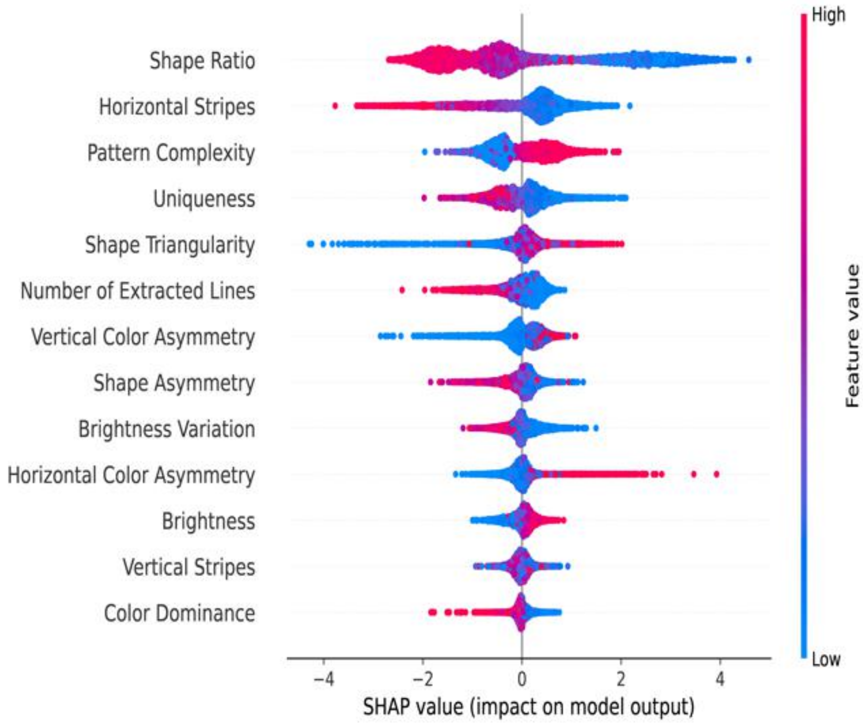

expectations from the interpretable model.

1.5.6 Non-image Features and Return Rates

Item category, price, and seasonality are all important features. With mean absolute SHAP values of 2.89, 2.78, and 1.37, respectively, they are, on average, more impactful than the automatically-generated image-based interpretable features. Non-image features are best included in the GBRT model from which SHAP values are computed, both as controls and because of their interactions with the automatically-extracted image-based interpretable features. The non-image features also provide valuable diagnostic information. Appendix A provides greater detail on the non-image features, e.g., sales and return rates by category.

Figure 1-11. Illustration of Combining Interpretable Image Features on Expected Return Rates

(a) Shirt with low predicted return rate **(b)** Shirt with low predicted return rate **(c)** Dress with low predicted return rate **(d)** Dress with high predicted return rate

1.5.7 Comparison of Predictive Models: Deep-learned vs. Interpretable Features

Recognizing the tradeoff between predictive ability and interpretability, we expect an interpretable-feature GBRT model to predict better than either the non-image baseline or the color-label model, but we do not expect the model to predict as well as the GBRT/CNN model. This is indeed the case: a GBRT model based on the automatically-extracted image-based interpretable features has an $R^2_{model} = 45.81$ which is less than that the $R^2_{model} = 46.88$ for the GBRT/CNN model. Both predictive abilities are well above the non-image baseline and the rudimentary image (color-label) model. Predictive ability is slightly better than the more-difficult-to-interpret automated-pattern-&-color-features model (Table 1.1) and better than the model based on HCFs (Section 1.4.5).

Our automatically-generated image-based interpretable features were curated carefully to provide insight while predicting return rates, but such choices are not unique. Retailers and researchers may wish to explore other interpretable features or combinations of HCFs and interpretable features.

1.6 Discussion and Further Research

Product returns generate considerable costs for online retailers – a large and growing retail channel. We propose that images, available prior to a fashion season, enable retailers to select which fashion-items to display online. We demonstrate, by example, that image-based features in a machine-learning model provide substantial incremental predictive ability relative to models based on traditional measures available to the retailer prior to launch. The predictive ability appears to be robust to a large number of variations. The display policy depends on the accuracy of the predictions and demonstrates that increased profits are feasible.

We augment predictions with automatically-extracted image-based interpretable features that can be used quickly and repeatedly for every fashion season and that scale to large assortments and many categories of items. The interpretable model sacrifices a small amount of predictive ability to provide diagnostic information valuable to the retailer's buyers and designers. Both the predictive and interpretable models, once developed and trained, run quickly and scale well.

Our application focuses on fashion-item returns in the apparel industry. This industry is important by itself, but we expect the approach to apply more broadly. Incorporating product images has the potential to improve predictive accuracy prior to product launch and generate important insights for design in industries such as hospitality, furniture, real estate, and even groceries.

Our analyses are illustrative and ceteris paribus. Researchers might explore (1) policies in which items are displayed online but not offline, (2) the implications on overall demand (online and offline) of not displaying items online, (3) interactions among items, (4) policies in which online items can be returned offline and thus increase offline traffic, (5) analyses that combine prelaunch features with postlaunch features, (6) how item features and prices jointly affect return rates, and (7) models that predict and provide insight jointly about sales and returns.

Chapter 2

Customer Search and Product Returns

2.1 Introduction

Online retailers are challenged by frequent product returns. Product returns often significantly decrease the profits of the firm by reducing the revenue (the firm must refund the returned product) and increasing the cost (backward logistics, dry cleaning, etc.)[1]. Return costs are so high that major online retailers such as Amazon and Walmart have begun to allow customers to keep the item because it sometimes costs less to refund the purchase price than bear the return costs (Wall Street Journal 2022). Zara announced it would be charging online shoppers for returns unless the items are returned to the physical store (BBC 2022). Managing product returns is critical to retailers' profitability. For example, L.L. Bean spent $50 million per year on returns costs, amounting to about 30% of the retailer‘s annual profits (Abbey et al., 2018).

Product returns are typically studied in a "purchase/return" framework where researchers

[1] According to the Wall Street Journal (2022), online returns can cost $10 to $20 per returned item, excluding freight.

assume the product purchase event as the starting point of the customer journey. In the purchase/return framework, research has established that product characteristics jointly affect the probability of purchase and return because the option to return a product has value to the customer and impacts purchase decisions. From a managerial perspective, research suggests that changes in the return policy (for example, towards a more lenient policy) would impact customers' purchase behavior.

I extend (1) purchase/return frameworks and (2) search/purchase frameworks to a unified search/purchase/return framework. Online retailing allows me to track the customer journey from the moment the customer starts looking for a product (pre-purchase), to the decision to purchase or not, and then to the decision on whether to keep or return the product (post-purchase). I seek to understand better the more-complete customer journey and gain insight on whether observing customer's search (pre-purchase) explains data-based stylized facts and illuminates mechanisms by which search and returns are related. Improved understanding provides insights about the customer journey and informs product return management.

Data on the entire search/purchase/return journey are rare. Using browsing sessions linked to data on purchasing and returning items at a major European apparel online retailer, demonstrate data-based stylized facts connecting search, purchase, and returns. The stylized facts motivate a rational model of the customer journey. However, the extended rational model presents challenges for parameter estimation. I discuss alternative strategies, some of which are not feasible, and I demonstrate a practical estimation strategy. The data and analysis suggest strategies by which a retailer can maximize its profits.

The focus of the paper is on understanding the search/purchase/return customer journey. By grounding a formal model empirically, I gain a deeper understanding of the mechanisms throughout the journey and identify opportunities such as targeted strategies that depend upon observed search. Although improved return-probability prediction is not the primary goal, improved insight improves prediction and enhances a firm's ability to manage backward logistics.

To motivate the value of including customer search when managing returns, consider Nelly and Wendy, who both purchased the same pair of jeans. Nelly kept the jeans while Wendy returned them for a full refund. From the purchase data, these customers are indistinguishable, however, search data might reveal that Nelly used search refinement tools (e.g., filtering products by color), searched many colors of the chosen option, and spent considerable time reviewing the product webpage – I will show that all of these observations reveal that Nelly will have a lower return probability. Wendy, on the other hand, purchased jeans from the front page without using filters, nor spending much time searching. Both Nelly and Wendy made rational decisions based on their needs, but such decisions provide insights about their subsequent behavior. The use of filters is not causal – we should not force Wendy to use filters, rather the use of filters tells us about Nelly's likelihood of returning an item.

The use of filters, the number of options searched, and search time reveal aspects of customers' underlying costs and benefits. Knowing these costs and benefits, the firm might develop policies to alter costs and benefits so that Wendy searches more and more efficiently and is then more likely to find a good apparel "match" and less likely to return the apparel item. Because the model is empirically grounded, policy simulations based on the model and the estimated parameters provide an initial test of strategies such as refinement tools, incentives to evaluate more colors, to search longer, and better website layouts that make it easier to search many items. In future research, such strategies can be subjected to A/B testing. As an example of how the empirically-grounded model might be used, I demonstrate an example policy simulation to evaluate the implications of targeted reductions in search costs. Policy development with a search/purchase/return model, combined with A/B testing, are particularly important in countries such as those in Western Europe that dictate return policies but do not dictate low-cost changes to website design. Similarly, I use the model to explore the changes on the retailer's website: the product display order on the retailer's website.

2.2 Related Literature

This paper synthesizes two fields of research: product returns and customer search. Research on product returns has been both theoretical and empirical. Theoretically, researchers demonstrate that the option to return products serves as a risk-reducing mechanism that encourages the customer to experience the product (Che (1996); also studied empirically in Petersen and Kumar (2015)) or as a signal of item quality (Moorthy and Srinivasan, 1995). Empirical research has focused on the optimization of return policies by firms. Researchers recognize the tradeoff between higher demand and higher return rates when firms use lenient policies and suggest that the optimal return policy must be balanced (Davis et al., 1998; Bower and Maxham III, 2012; Abbey et al., 2018) because overly strict return policies lead to a decrease in purchases (Bechwati and Siegal, 2005). Janakiraman et al. (2016) provide an extensive review of the effect of return policy leniency on purchases and returns.

Anderson et al. (2009) propose a structural model where the option to return is embedded in a customer purchase decision – the customer learns private information only after purchasing the product. Other empirical studies demonstrate that a variety of factors affect the probability of product returns including price, discounts, marketing instruments (e.g., free shipping), or the truthfulness of product reviews (Petersen and Kumar, 2009, 2010; Sahoo et al., 2018; Shehu et al., 2020; El Kihal and Shehu, 2022). Empirical studies have focused on concrete instruments, such as visualization systems and online product forums. These instruments decrease product uncertainty in the match of the product to the customer and, therefore, product returns (Hong and Pavlou, 2014). Recent research has used machine learning to accurately predict returns and identify product-related features to be considered when selecting and designing fashion items for the retailer's website (Cui et al., 2020; Dzyabura et al., 2019)

The second field of research, customer search, is an established and mature field of research both empirically and theoretically. The literature typically follows either sequential (Weitzman, 1979) or simultaneous (Stigler, 1961) approaches. In both approaches, the customer knows the distribution of the rewards and searches to fully resolve uncertainty. For example, Weitzman shows that the optimal (dynamic programming) search strategy for some relatively simple search problems is an index strategy. Most of the literature focuses on sequential search buttressed by Bronnenberg et al. (2016) who report strong evidence to support sequential search. Recent papers allow for flexible preference heterogeneity (Morozov et al., 2021), add learning (Ke et al., 2016; Branco et al., 2012; Dzyabura and Hauser, 2019), multiple attributes (Kim et al., 2010), intermediaries (Dukes and Liu, 2016), search duration (Ursu et al., 2020), and search fatigue (Ursu et al., 2023). Like many of these models I assume the customer searches optimally as if the customer solves a dynamic program that anticipates acting optimally for the remainder of the customer journey – in this case, subsequent search, purchase, and potential returns.

The availability of click-stream data has enabled researchers to study empirically customer search behavior (Bronnenberg et al., 2016; Chen and Yao, 2017; Ursu et al., 2020) and provide detailed insights on search-to-purchase customer behavior. For example, Bronnenberg et al. (2016) examine consumer search behavior for cameras and show that early search is highly predictive of consumer purchase and that the first-time discovery of the purchased alternative happens towards the end of the search. Chen and Yao (2017) show that refinement tools significantly impact consumer behavior and the market structure. Ursu et al. (2020) study search duration, quantify consumer preferences and search costs, and develop insights on how much information to provide on a platform. The data enable us to build on these insights to examine the full customer journey.

The product returns literature focuses on the purchase-to-return portion and the search literature on the search-to-purchase portion. The research expands the focus by considering both the pre-purchase and post-purchase customer journey in online retailing – search-to-purchase-to-return. In doing so, the paper links the two fields of research, contribute to both, and investigate new phenomena. First, I embed the "pre-purchase" events in a traditional "purchase/post-purchase" framework studied in returns literature. I seek to improve the knowledge on why product returns happen and how the retailer could manage them. Second, the model extends the existing search models by accounting for returns. The analyses in the fashion industry generalize directly to cancellations in the travel industry, such as for hotels, Airbnbs, airlines, cruise ships, and resorts. I demonstrate that the option to return impacts the way customers search for the product. By modeling the search-to-purchase-to-return customer journey, I gain insight on customer behavior in the full customer journey with implications of how to manage the entire process more profitably. Not only do search patterns help predict returns, but nudging search might influence net purchases after accounting for returns.

2.3 Data and Model-Free Evidence

2.3.1 Descriptive Statistics

I sought and obtained online-channel individual-level data from a large apparel retailer in Western Europe. I focus on the online channel because (1) most returns are through the online channel (the retailer has in total 53% of sold items being returned – typical for the European apparel industry) and (2) the online channel is an ideal situation in which to observe search, purchases, and returns for each customer. Ultimately, insights from online shopping should be relevant to offline shopping, but I leave that to future research. I preprocessed data by removing noise and outliers (for example, extremely short/long sessions). In this paper, I focus on orders which had at most one product purchased. This focus isolates the impact of search on product returns by excluding situations when the customer purchases several colors or variations of a product with an intention to keep only one. The model-free evidence suggests that qualitative behavior of customers does not change for situations in which multiple products are purchased. I leave the extension to multiple-item orders for future research. A detailed description of data pre-processing can be found in Appendix B.1.

The retailer sells medium-priced fashion products for women, men, and children. (The retailer sells mostly adult apparel, which compromises 95% of the purchases.) As is typical for Europe, the retailer has a generous return policy, where items can be returned for free, for a full refund, within 60 days after the purchase with or without providing a reason.

The data include both mobile and desktop usage and consist of three main components:

- ***Search*** records the sequence of actions made by the customer during the browsing session. I observe products listed on the website for the customer and the set of products considered (clicked to view the detailed product page), as well as the sequence of these clicks. I also observe all actions (e.g., clicking on a product, sorting by price) and the timing between the different actions, allowing us to observe how much time a customer spends on a specific product page.
- ***Purchases*** includes the products purchased (if any) by the customer during browsing sessions. These data include product characteristics such as price, category, fabric, size, brand, color, and product image. I also observe each product's base price and can infer whether it was sold with a discount.
- ***Returns*** contains information on whether the purchased product was kept or returned by the customer. The data also include the date and stated reason for the return.

All three components of the data are matched by a unique identifier. For each session, I observe the complete customer journey from opening the retailer's website to deciding whether to keep a fashion item. This allows us to build and estimate the model which combines customer search, purchase, and return decisions. The observation period is between 1 October 2019 and 15 May 2020. Over this period, in the final dataset I observe 482,962 search sessions, of which 54,585 (11.3%) result in a purchase. In 32.5% of these single-item cases, customers decide to return the product purchased. As anticipated, the return rate for the single-item subsample is lower (32.5%) than the return rate for the multiple-item subsample (53%), but the qualitative implications remain intact. Figure 2-1 provides summary statistics for the data.

Search Descriptives. Customers can access the retailer's website through desktop or a mobile device (61.1% access through a mobile device in the data). On the website, the customer observes a product list, which displays an image of the product, its price, category, and a small front view picture of the product. Customers can use search refinement tools to select a more specific product list. The website offers two types of refinement tools: filtering and sorting. Customers can filter by brand, color, products on sale, new products, or product size. They can sort the product list by price (ascending or descending), new products, or top sellers. When the customer clicks on a specific product, further information is revealed on the product page, such as more (and higher quality) product images, and detailed product description. Amongst the 54,585 search sessions, median length of customer search is 263 seconds before purchasing or leaving the retailer's website. On average, they browse 5.5 product pages, spending 46.7 seconds per page (Figure 2-1). 45% of customers use at least one refinement search tool (Figure 2-1).

Purchase Descriptives. Out of the 482,962 search sessions, 54,585 ended with a purchase. Customers choose among 16 product categories, as predefined by the retailer (e.g., jeans, blouses, dresses, coats, shoes). The most popular purchased product categories are "jackets and coats" (26.6%) and jeans (16.3%).

Return Descriptives. Dresses and jumpsuits have the highest probability of return (46.4% and 47.9% respectively) and t-shirts have the lowest return rate (11.1%).

Figure 2-1. Summary Statistics of Data

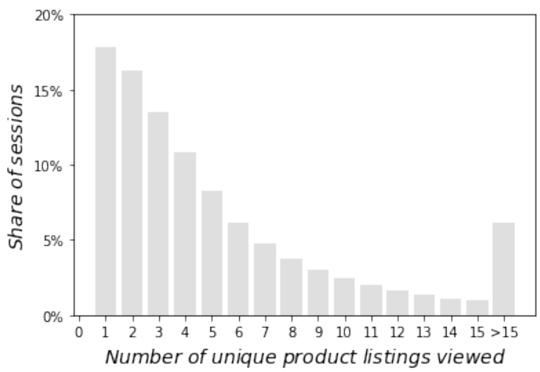

(a)

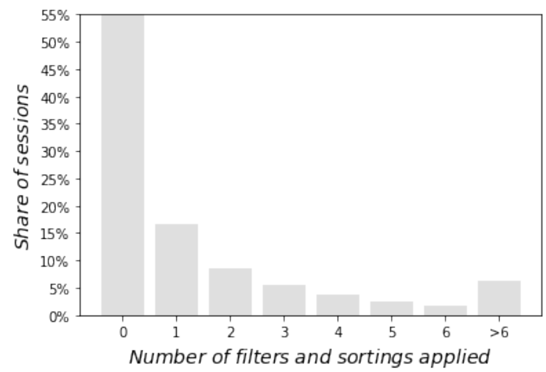

(b)

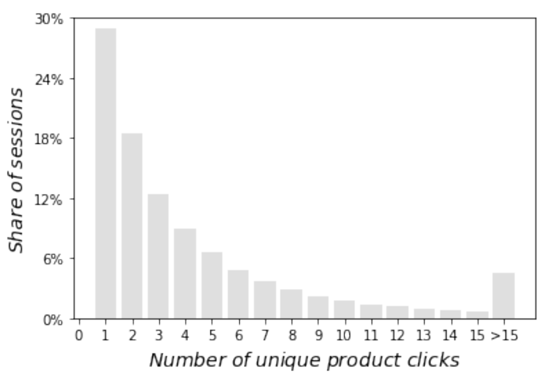

(c)

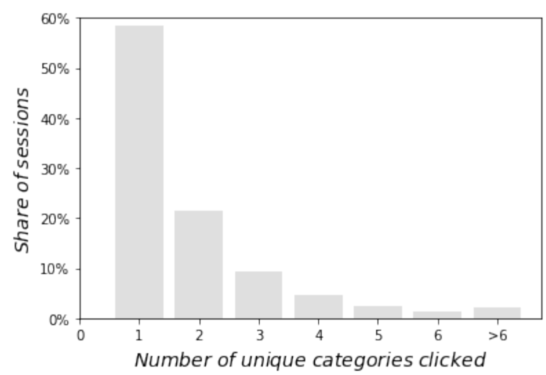

(d)

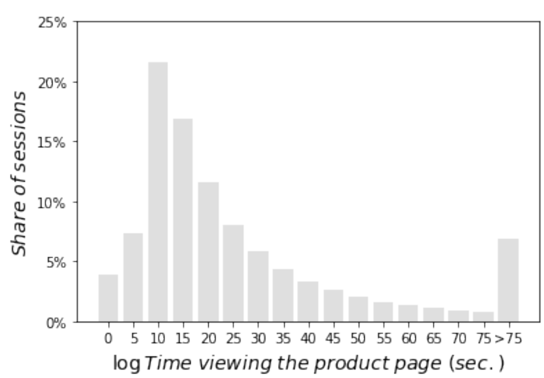

(e)

2.3.2 Model-Free Evidence

Before I turn to the model of rational behavior throughout the purchase journey, I provide illustrative model-free evidence. The probability of returning a product is linked to different aspects of customers search journey (Figure 2-2). The difficulty with search data comes from the facts is that it is very unstructured and in empirical setting could not be summarized into one variable "search". To address this issue I compute aggregate statistics of customer related to different aspects of customer search plot their relation with returns.

First, I look at the number of products on which the customer clicked to reveal additional information during the search session. In Figure 2-2a, I observe a strong positive correlation between the number of products clicks and the probability of return. Customers who browse more products return on average more frequently.

Actions which precede product views (or clicks) are also related to product returns. Specifically, customers who use tools to refine their search experience (for example, filtering product by price, size, color, etc.) are less likely to return the purchased product. Figure 2-2b relates the number of filters applied by the customer to return probability. Moreover, the way customers review the product page is also related to product returns. Figure 2-2c suggests that customers who spend more time reviewing the product page are less likely to return the product.

The results in Figures 2-2a to 2-2c demonstrate the relation of customer search and probability of returns. However, all these measure ignore an important aspect of customer search – which products the customer was clicking. For example, consider two customers who both clicked on five products during their browsing session. However, the first one was

Figure 2-2. Informativeness of Search on Product Returns

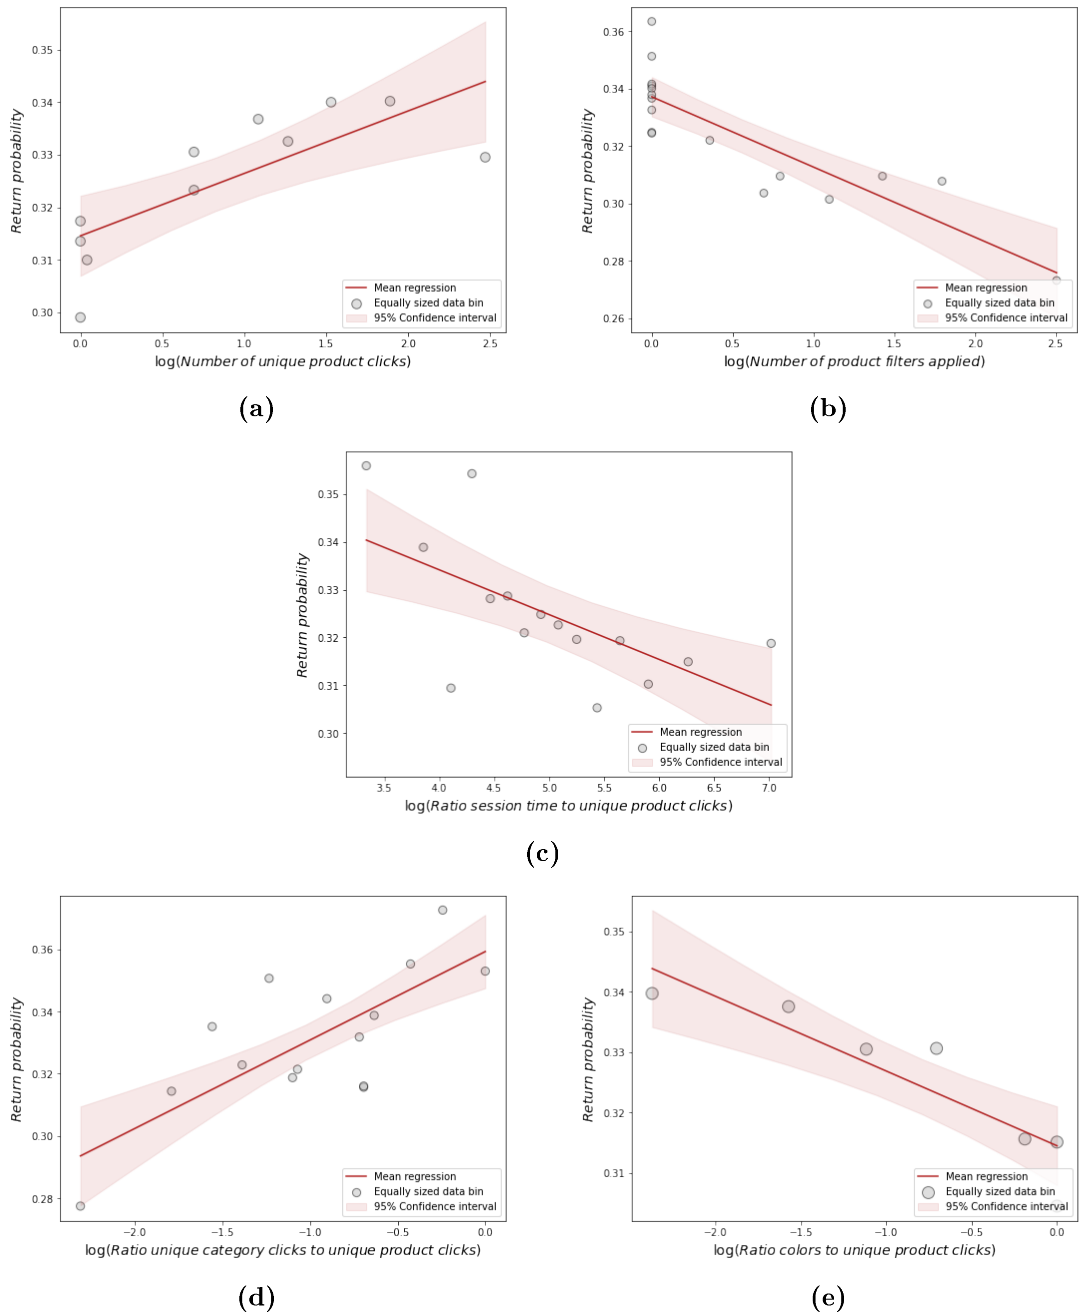

clicking only at similar T-shirts, while the second customer clicked on two T-shirts, two pairs of jeans and dress. Their search is very different and these customer would likely have a different return behavior (the first one is more focused and knows what he or she looking for). Notice that Figure 2-2a would treat them as the same customers.

Figure 2-2d and Figure 2-2e address this issue by shedding light on the type of product the customer searches. First, customers whose click set mostly consists of different color varieties of purchased product are less likely to return the product (Figure 2-2d). While at the same, customers who click at products from many different product categories are more likely to return the product (Figure 2-2e). These two results demonstrate the difference between the deep search (looking for something specific) and the broad search (casually browsing the website). The former are less likely to return the product after the purchase as they purchase is more informed and less impulsive. In Appendix B.2 I use deep learning product embedding to demonstrate that customer who look high-variety products are more inclined to return the product.

2.4 Model

Building on the literatures in search-to-purchase and on purchase-to-returns, I model search as sequential and rational – customers act as if they solve a dynamic program. Consider a customer i, who is searching for a product on the retailer's website. When visiting the website, the customer observes products on the website from which to choose (or click)[2]. The total number of products available to customer i is S_i. Each product view from this list

[2]Starting here, I use product click to describe the event when a customer views the product page by clicking on it from the list on the website.

contains information about the product (for example, product category, price, etc.). After evaluating these products, the customer may choose to click on a product to reveal additional information. For instance, a click on a product page could reveal a detailed image of the product or additional information such as the fabric. I capture this by assuming that the customer observes characteristics of the product x_{ij}, before customer i clicks on product j, and receives a signal ψ_{ij} after clicking on the product. This signal is a noisy estimate of how much the customer would like the product when it arrives home. I assume that the customer pays search costs to reveal the signal ψ_{ij} (mental costs, mouse navigation, or time costs). Unlike many, but not all, search models, e.g., Weitzman, search reveals a signal (noisy information) about the product rather than product quality or final utility.

After evaluating a clicked product, the customer has two options: continue the search by clicking on the next product or stop the search. In case of stopping the search, the customer purchases one of the previously clicked products (or chooses the outside option). In case the customer decided to purchase a product, the customer must wait until the product arrives at home before receiving additional information about the product. At home, the customer receives this additional information and decides whether to keep the product or return it to the retailer. I assume that the product inspection at home reveals a true customer's individual preference shock ϵ_{ij} for the product, which captures all the product information revealed to the customer upon physical product evaluation (for example, the customer understands the physical fit of the product, its expected use, or simply how close it is to the customer's fashion taste). Returns are costly. The customer pays return costs R_i. Despite the "free" returns policy return, return costs include printing the return label, bringing the package to the postal office, mental costs, and time spent returning the item.

Figure 2-3. Sequence of Customer Actions and Observed Information in the Search, Purchase, and Return/Keep Stages

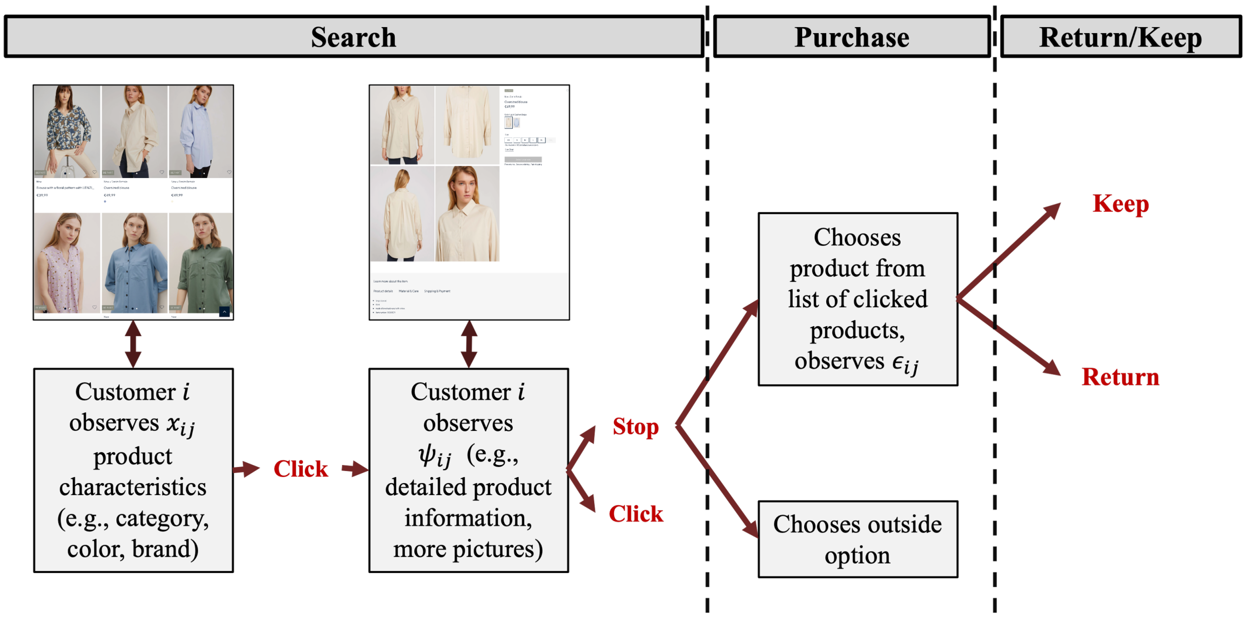

In Section 2.4.1 I specify the utility of the customer and the effect of the returns option. I elaborate search costs in Section 2.4.2. I examine the optimal search rules under the availability of returns in Section 2.4.3. Figure 2-3 illustrates the timing of the customer search-purchase-return journey as well as all the information in each stage.

2.4.1 Utility and Returns

For ease of notation without loss of generality, I number products such that j is the index of a sequence in which customer i searches ($j = 1$ implies the first clicked product). The customer utility of purchasing the product j from the website could take of three possible forms in Equation (2.1). Search costs are paid prior to the realization of this utility.

$$u_{ij} = \begin{cases} \mu^u(x_{ij}) + \epsilon_{ij} & \text{if customer purchases product } j \neq 0 \text{ and keeps it} \\ -R_i & \text{if customer purchases product } j \neq 0 \text{ and returns it} \\ 0 & \text{if customer chooses the outside option } j = 0 \end{cases} \tag{2.1}$$

where x_{ij} is a vector of product characteristics (e.g., category, color, brand, etc.); $\epsilon_{ij} \sim \mathcal{N}(0, \sigma_{\epsilon_{ij}})$ is the individual preference shock; and R_i return costs (e.g., shipping fee, travel time to the postal office, etc.). Because j is an index of a sequence, x_{ij} takes on different values for different customers. Without loss of generality, I set the utility of the outside option to zero and specify $\mu^u(x_{ij})$ to have a nonzero intercept.

As outlined in Figure 2-3, x_{ij} represents information readily available on the website before starting the search and ψ_{ij} represents information revealed after the customer clicks on the product page (additional information about the product like reviews, high-quality pictures, etc.); ϵ_{ij} is revealed after the customer makes a purchase decision (for example, at home the customer tries the received product and evaluates the product as an addition to the customer's wardrobe). To understand how the return option impacts the search, remember that the customer observes ϵ_{ij} only after receiving the product, therefore the customer must make a purchase decision without knowing ϵ_{ij}, but anticipating that ϵ_{ij} will be revealed. Clicking on the product on the website only reveals ψ_{ij} which is a noisy signal of the individual preference shock ϵ_{ij}, which could be expressed as:

$$\psi_{ij} = \eta_{ij} + \epsilon_{ij} \tag{2.2}$$

where $\eta_{ij} \sim \mathcal{N}(0, \sigma_{\eta_{ij}})$ is the noise of the signal. I assume that η_{ij} and ϵ_{ij} are independently distributed and that the customer knows the distribution of η_{ij} and ϵ_{ij} before search as well as all model parameters.

Because the customer knows the distribution of $\psi_{ij}|\epsilon_{ij}$, it is feasible to compute for each clicked product the expected utility, v_{ij}, given the signal ψ_{ij} (see Appendix B.3 for details):

$$
\begin{aligned}
& v_{ij} \\
=& \mathbb{E}_{\epsilon_{ij}} \left[(\mu^u_{ij} + \epsilon_{ij}) \cdot \mathcal{I}(\mu^u_{ij} + \epsilon_{ij} \geq -R_i) + (-R_i) \cdot \mathcal{I}(\mu^u_{ij} + \epsilon_{ij} < -R_i) | \psi_{ij} \right] \\
=& \sigma_{v_{ij}} \cdot T\left(\frac{\mu^u_{ij} + \epsilon_{ij}}{\sigma_{v_{ij}}} + \frac{\psi_{ij} \cdot \sigma_{v_{ij}}}{\sigma^2_{\eta_{ij}}} \right) - R_i \\
& \sigma_{v_{ij}} \sqrt{\frac{\sigma_{\epsilon_{ij}} \cdot \sigma_{\eta_{ij}}}{\sigma^2_{\epsilon_{ij}} + \sigma^2_{\eta_{ij}}}}
\end{aligned}
\tag{2.3}
$$

where $T(x) = \Phi(x) \cdot x + \varphi(x)$; $\Phi(x)$ and $\varphi(x)$ are standard normal cdf and pdf respectively and $\mathcal{I}(Condition) = 1$ if $Condition$ is $True$

Equation (2.3) demonstrates how the return option impacts customer search – because v_{ij} depends on the random variable ψ_{ij} that is unobservable before search (or click), the return option changes the distribution of the reward from search. One can show that $T(x) \geq x$ and thus $v_{ij} \geq \mu^u_{ij} + \frac{\sigma^2_{\epsilon_{ij}}}{\sigma^2_{\epsilon_{ij}} + \sigma^2_{\eta_{ij}}} \psi_{ij}$. This implies that for any product characteristics, the option to return improves the customer's expected search utility. Intuitively, by the principle of optimality, having the option, but not the obligation, to return a product is weakly better than a strategy of "always keep the product."

2.4.2 Search Costs

Each additional search (or click) requires the customer to pay a search cost. For example, search cost might include mental or physical effort of reviewing the additional information, or clicking/moving mouse. I denote the search costs, c_{ij} , and model its relationship to the search environment by:

$$\log c_{ij} = \mu^c(z_{ij}) + \xi_{ij} = z_{ij}'\beta^c + \xi_{ij} = \mu_{ij}^c + \xi_{ij} \tag{2.4}$$

where z_{ij} is a vector of search-object characteristics affecting the search costs (position on page, device, page number, age, etc.), β^c are levels of search costs coefficients, and $\xi_{ij} \sim \mathcal{N}(0, \sigma_{\xi_{ij}})$ is an individual product-level shock on search costs that follows a normal distribution. The search-object characteristics are important for understanding how website design interacts with the customer journey. Changing these characteristics changes costs and affects the optimal search.

I assume that the customer observes both z_{ij} and ξ_{ij} before the search (or before clicking on the product). In Section 2.5.2, I explain that the assumption of heterogeneous search costs is essential for model parameter estimation.

2.4.3 Optimal Search Strategy when Returns are Available

Because the return option changes the distribution of the reward, the decision rules common in the search literature must be updated. Conceptually, the selection, stopping, and purchase rules retain the property of maximum expected utility, but the rules anticipate more, are much more complicated, and they equations change. I summarize these decision rules below and provide derivations in Appendix B.4:

- **Selection rule.** If the customer is going to search (or click), the click would be the option with the highest reservation utility ω_{ij} derived from the system in Equation (2.5).

$$
\begin{aligned}
c_{ij} &= \sigma_{v_{ij}} \int_{\theta}^{+\infty} T\left(\frac{\mu_{ij}^{u}+R_i}{\sigma_{v_{ij}}} + \frac{\psi_{ij}\cdot\sigma_{v_{ij}}}{\sigma_{\eta_{ij}}^{2}}\right) - T\left(\frac{\mu_{ij}^{u}+R_i}{\sigma_{v_{ij}}} + \frac{\theta\cdot\sigma_{v_{ij}}}{\sigma_{\eta_{ij}}^{2}}\right) dF(\psi_{ij}) \\
\omega_{ij} &= \sigma_{v_{ij}} T\left(\frac{\mu_{ij}^{u}+R_i}{\sigma_{v_{ij}}} + \frac{\theta\cdot\sigma_{v_{ij}}}{\sigma_{\eta_{ij}}^{2}}\right) - R_i
\end{aligned}
\tag{2.5}
$$

  Notice that the second equation is a 1-to-1 mapping $\theta \to \omega_{ij}$ and thus, to find ω_{ij}, I need only solve the first equation for θ

- **Stopping rule.** The customer continues to search until the customer's maximal expected utility of searched options (Equation (2.3)) exceeds the maximal reservation utilities of unsearched options. This stopping rule is conceptually similar to most of the search literature; the manner in which it is computed is changed.

- **Purchase rule.** When the stopping rule is reached, the customer purchases either the option from the set of searched ones, which yields the highest expected utility, or the outside option.

Equation (2.5) illustrates the difference between the standard search model and the model with a return option. Specifically, by changing the reward distribution, the return option changes the reservation utilities for the customer. These changes are substantial: the return option could change the order in which the customer searches products and could change the stopping rule. The option to return also changes the purchase rule, as in this case the customer would compare $v_{ij}(\psi_{ij})$ rather than $\mu_{ij}^{u}(x_{ij}) + \epsilon_{ij}$. For face validity, it is helpful to compare Equation (2.5) with a classic case of "no returns allowed". When $R_i \to +\infty$ and $\sigma_{v_{ij} \to 0}$, Equation (2.5) converges to the standard equation for reservation utilities.

2.5 Empirical Specification and Estimation Strategy

In this section, I describe in detail the estimation strategy and empirical specification. Because the return option complicates Equation (2.5), estimation becomes more difficult.

2.5.1 Utility

I assume that the utility of customer i's purchase of product j is in Equation (2.3) a linear function of product characteristics. To focus on search and returns and not overparameterize the model, I assume that customers have homogeneous preference vectors for product characteristics:

$$\mu^u(x_{ij}) + \epsilon_{ij} = x_{ij}'\beta^u + \epsilon_{ij} \tag{2.6}$$

where x_{ij} is a vector of product characteristics; β^u is a vector of the customer's sensitivity to product attributes; and $\epsilon_{ij} \sim \mathcal{N}(0, \sigma_{\epsilon_{ij}})$ follows i.i.d. standard normal distribution. Without loss of generality, because utility is unique only to a positive linear transformation, I set $\sigma_{\epsilon_{ij}} = \sigma_\epsilon = 1$.

Recall that the customer receives a noisy signal ψ_{ij} before purchase. In Equation (2.7), I assume that the noisy signal is normally distributed $\eta_{ij} \sim \mathcal{N}(0, \sigma_{\eta_{ij}})$ and that this signal depends upon the product and session characteristics. Specifically,

$$\log \sigma_{\eta_{ij}} = y_{ij}'\beta^\eta \tag{2.7}$$

where y_{ij} is a vector of product/session characteristics which could affect the quality of the signal; β^{η} is a vector of weights for the strength of signal quality. The vector, y_{ij} enables us to model website design (and the product itself) affect the quality of the signal that the customer receives.

From Equations (2.5) to (2.7), it follows that the expected purchase utility under the return option is a function of model parameters β^u, β^{η}, R (return cost); observable variables x_{ij}, y_{ij} and the unobservable shock ψ_{ij}. For ease of notation for the remainder of the paper, I make this dependency explicit:

$$v_{ij} = v_{ij}(\beta^u, \beta^{\eta}, R, x_{ij}, y_{ij}, \psi_{ij}) \tag{2.8}$$

2.5.2 Search Costs

I capture heterogeneity in search costs with the vector of search-object characteristics y_{ij} that depends upon both the search object j and the customer i. For identification, I set $\sigma_{\xi_{ij}} = \sigma_{\xi} = 1$. The random cost assumption is important because without it, the likelihood function could be equal to zero for some customers. To see why the model would collapse, consider the case when the costs are not random. Constant costs imply that the reservation utilities in Equation (2.5) would be deterministic. For a given list of products, all products would have the same fixed reservation utilities for every consumer. Therefore, given the optimal search rules, all customers viewing this list would click on products in the same order which contradicts of data. Mathematically, the observed data would make infeasible the set of constraints regarding the search order (Equation (2.11) left part). Without randomness in costs, there is no combination of values of model parameters such that there is a non-zero probability of observing data where the search-purchase varies among consumers.

2.5.3 Likelihood

Let S_i denote the number of products presented to the customer (for example, the number of products the customer sees on the main page of the website). From this set of products, the customer searches (clicks) on C_i products according to optimal search rules in Section 2.4.3. Recall that the index j represents the order in which the customer searches for the product (e.g., $j = 2$ denotes the second searched/clicked product and $j = C_i$ denotes the last searched/clicked product). Notice that this notation implies that the customer did not click on products with $j > C_i$ and thus I can enumerate non-clicked products randomly.

The customer acts as if the customer's actions were described by knowing all parameters of the model, as well as observing the shocks $\epsilon_{ij}, \psi_{ij}, \xi_{ij}$, which are not observable by the researcher. Because data enable us to observe all other information that is available to the customer, I use the optimal search rules from Section 2.4.3 to write constraints on model parameters.

Consider a customer who searched for C_i products; purchased a product with index b and decided to keep it. In (Equations (2.9) to (2.12), the following constraints on the customer parameters are implied where $\mathcal{I}(Condition) = 1$ if the condition is satisfied.

Return. The customer keeps a purchased product b if the product utility is greater than the negative return costs R:

$$\mathcal{I}(x_{ib}'\beta^u + \epsilon_{ib} \geq R) \tag{2.9}$$

Purchase. The customer purchases the product if its expected utility from Equation (2.3) is greater than the expected utility of all other products in the consideration set ($j = 0$ is the outside option).

$$\mathcal{I}(v_{ib} \geq \max_{s=0..C_i} v_{is}) \tag{2.10}$$

Search continuation. After searching option j, the customer continues the search if the best option on hand, $\max_{s=0..j} v_{is}$, is worse than the value of searching the unsearched options. The customer would choose the option with maximal reservation utility ω_{ij}:

$$\forall j < C_i \quad \mathcal{I}(\omega_{ij+1} \geq \max_{s=j+2..S_i} \omega_{is}) \cdot \mathcal{I}(\max_{s=0..j} v_{is} \leq \max_{s=j+1..S_i} \omega_{is}) \tag{2.11}$$

Search stopping. The customer stops searching when the maximal expected utility of searched options is higher than the value of searching the remaining options:

$$\mathcal{I}(\omega_{iC_i} \geq \max_{s=C_i+1..S_i} \omega_{is}) \cdot \mathcal{I}(\max_{s=0..C_i} v_{is} \geq \max_{s=C_i+1..S_i} \omega_{is}) \tag{2.12}$$

Equations (2.9) to (2.12) define the set of constraints that must be satisfied to observe the given search sequence. Multiplication of these equations produces a binary variable W_i which takes 1 if and only if all constraints are satisfied. The case when the customer decides to return a product (or chooses the outside option) closely follows derivations in Equations (2.9) to (2.12). In Appendix B.6, I show that the variable W_i can be rewritten in a more compact way[3]:

$$
\begin{aligned}
W_i &= W_i(\beta^u, \beta^c, \beta^\eta, R, x_i, y_i, z_i, \epsilon_i, \psi_i, \xi_i) = \\
&= \left[\prod_{j=1}^{C_i-1} \mathcal{I}(\omega_{ij} \geq \omega_{ij+1})\right] \cdot \mathcal{I}(\omega_{iC_i} \geq \max_{j=C_i+1..S_i} \omega_{ij}) \times \\
&\times \left[\prod_{j=0}^{C_i-1} \mathcal{I}(v_{ij} \leq \min\{\omega_{iC_i}, v_{ib}\})\right] \cdot \mathcal{I}(v_{iC_i} \leq v_{ib}) \cdot \mathcal{I}(v_{ib} \geq \max_{j=C_i+1..S_i} \omega_{ij}) \times \\
&\times \mathcal{I}(x'_{ib}\beta^u + \epsilon_{ib} \geq -R)
\end{aligned}
\tag{2.13}
$$

Because the researcher does not observe $\epsilon_i, \psi_i, \xi_i$, I obtain the probability of observing the given search sequence of customer i by integrating out these variables:

$$
P_i = P_i(\beta^u, \beta^c, \beta^\eta, R, x_i, y_i, z_i) = \int \cdots \int W_i \cdot dF(\epsilon_i, \psi_i, \xi_i) \tag{2.14}
$$

Given this probability for each observation in a sample, I compute the log-likelihood function:

$$
LL(\beta^u, \beta^c, \beta^\eta, R) = \sum_{i=1}^{N} \log P_i(\beta^u, \beta^c, \beta^\eta, R, x_i, y_i, z_i) \tag{2.15}
$$

2.5.4 Estimation

Variables x_i, y_i, z_i are observable $\forall i$, thus, if computations were feasible, I would find the estimate of the parameter vector $\beta = (\beta^u, \beta^c, \beta^\eta, R)$ by maximizing the log-likelihood function in Equation (2.15). Unfortunately, this maximization is not feasible for two reasons.

First, the reservation utilities ω_{ij} from Equation (2.5) are not computed directly because

[3] I drop the product index for compactness. For example, x_i should be read as a set of variables for all products in the search set $\{x_{ij}; j \in 0..S_i\}$

they are defined through implicit functions. Without further simplification and approximation, solving for the maximum-likelihood estimates would not be feasible with today's computers. (Working from simplified problems, I estimate it would take more than 1,000 years.) To make computation feasible, I approximate the function $\omega_{ij}(\mu_{ij} + R, \sigma_{\eta_{ij}}, c_{ij})$ with interpolation techniques. Details are given in Appendix B.4.

Second, when integrating over unobserved shocks $\epsilon_i, \psi_i, \xi_i$, there is no known closed-form solution for this integral. I must approximate the integral. I considered the following approximations for integration.

Accept-reject simulator (Chen and Yao, 2017). An accept-reject simulator replaces the true probability P_i with a simulated probability $\widehat{P}_i$. In this approach, for given parameter estimates, I simulate B random draws of shocks from corresponding distributions and calculate the share of draws in which $W_i = 1$ (all constraints in Equation (2.13) are satisfied). The parameter vector corresponding to the largest share of draws is the maximum-likelihood estimator. The challenge with this approach is that the nature of search data makes P_i close to zero and thus requires large values of B with a correspondingly substantial increase in computation time. Compounding the computational limit is the fact that this approach produces a non-smooth objective function which requires the use of non-gradient optimization methods (e.g., Nelder-Mead method). Such methods are substantially slower. In Appendix B.9 I demonstrate that this approach is not feasible for data because this approach produces highly imprecise estimates.

Accept-reject simulator with smoothing (Honka and Chintagunta, 2017; Ursu, 2018). I considered replacing the sharp constraints in accept-reject simulation, such as $\mathcal{I}(a < b)$, with a continuous function of differences $b - a$. This approach punishes large violations of the constraint but allows small differences. While this approach is often feasible, it is not feasible when the full customer journey is modeled. First, most of constraints of the form $\mathcal{I}(a < b)$ have arguments a and b bounded from below (for example, v_{ij} is bounded below by $-R$ because $T(x) \to$ if $x \to \infty$). Because of these bounds, the difference $b - a$ does not translate well into a probability that the constraint was violated. Second, returns are represented by a single constraint and I observe returns only for the searches that ended with a purchase (approximately, 10% of the sample). This implies that the "returns constraints" constitute a small proportion of all the constraints in the model. Thus, violation of the "return constraints" would have a much lower impact on the final objective function. Artificially, the optimal solution to Equation (2.13) would rely almost exclusively on the "non-return constraints" and accept solutions that violate the "return constraints." Because of this imbalance, smoothing of all constrains poses threats to the estimation quality, especially on the critical parameters related to returns.

Partly closed form integration. The third approach to feasibility recognizes that some, but not all variables, in the constraints can be integrated out, at least theoretically. For example, the return constraint $\mathcal{I}(x_{ib}'\beta^u + \epsilon_{ib} \geq -R)$ could be replaced with probability $1 - \Phi\left(-(x_{ib}'\beta^u + R) \cdot \sqrt{\frac{\sigma^2_{\epsilon_{ib}} + \sigma^2_{\eta_{ib}}}{\sigma^2_{\epsilon_{ib}} \cdot \sigma^2_{\eta_{ib}}}} - \frac{\sigma_{\epsilon_{ib}}/\sigma_{\eta_{ib}}}{\sqrt{\sigma^2_{\epsilon_{ib}} + \sigma^2_{\eta_{ib}}}} \cdot \psi_{ib}\right)$ by integrating out the shock $\epsilon_{ib}|\psi_{ib}$. In Appendix B.7, I prove that I can rewrite the binary variable W_i as:

$$\widetilde{W}_i = \mathcal{I}(Condition1) \cdot \mathbb{P}[Condition2] \tag{2.16}$$

where $Condition1$ are constraints which cannot be integrated out and thus represented by binary variable; $Condition2$ are constraints which could be integrated out and represented by continuous variable from 0 to 1.

In Appendix B.7, I demonstrate that $Condition 1$ has only one sharp constraint that cannot be integrated out. Hence, the requirement on the number of draws for the sharp constrain is significantly smaller than it would be if all constraints were simulated.

To mitigate the concern that the optimal solution might ignore the returns data due to imbalance of the number of constraints, I implement a two-stage approach:

- In the first stage, I estimate only the purchase-return model.

$$\begin{aligned} W_i^{FS} &= W_i^{FS}(\beta^u, \beta^\eta, R, x_i, y_i, \epsilon_i, \psi_i) \\ &= \left[\prod_{j=0}^{C_i} \mathcal{I}(v_{ij} \leq v_{ib})\right] \cdot \mathcal{I}(x'_{ib}\beta^u + \epsilon_{ib} \geq -R) \end{aligned} \tag{2.17}$$

- In the second stage, I fix the parameters related to product returns (R and β^η) and estimate all other parameters by simulating the variable from Equation (2.17).

The two-stage approach is an approximation that works well with partly-closed-form integration, but it is an approximation. To examine the implications of the approximation, I create synthetic data with known parameters and apply the three alternative methods. In Section 2.6 I demonstrate that the two-stage approach with partly-closed-form integration achieves the highest accuracy and recovers the values of parameters quite well for the synthetic data. See Appendix B.9.

2.6 Parameter Identification and Monte Carlo Simulations

2.6.1 Parameter Identification

In Section 2.5 I introduced the empirical specification of the model and demonstrated that it depends on six groups of parameters: $\beta^u, \sigma_{\epsilon_{ij}}, \beta^\eta, \sigma_{\xi_{ij}}, R_i$. To assure that critical parameters can be uniquely recovered from empirical data, I follow a standard approach in search and purchase models; I normalize the variance of individual unobserved product fit $\sigma_{\epsilon_{ij}}$ to 1 and the variance of individual search costs $\sigma_{\xi_{ij}}$ to 1 for each customer i and product j. For the analyses, I set the return costs R_i to be constant between customers. Extending the model to include heterogeneity in return costs when a retailer multiple browsing sessions for each customer, is straightforward. Considering these constraints, I summarize in Table 2.1 the parameters that need to be estimated.

From Table 2.1 it follows that in addition to the classical search model parameters (β^u and β^c), the model also has the parameters β^η and R, which are identifiable from the customer return transactions data. Prior research has argued that parameters β^u and β^c can be identified from search data. I need only argue that product returns data allows us to identify additional parameters β^η and R without compromising identification of β^u and β^c.

In classical models, β^u is identified from the purchase data. Similarly, for the model β^u, β^η, and R identified from purchase and return data. However, I need to impose additional constraints on β^η similar to the model considered by Anderson et al. (2009). Specifically, the

Table 2.1. Parameters to be Estimated

Par.	**Intuition**	**Example**	**§**
β^u	Customer's sensitivity of mean utility to product attributes	Customers may prefer on average products made of natural fabric	Section 2.5.1
β^η	Vector of weights for the strength of signal quality, revealed to the customer upon product clicking	A simple white T-shirt might be easier to evaluate online than a nuanced night dress	Section 2.5.1
β^c	Levels of search costs coefficients	Clicking on the product at the top of the website requires less movement of the mouse, scrolling, or reading	Section 2.4.2
R	Return costs paid by the customer	Customer need to spend time packaging the item and posting it	Section 2.4.1

set of features on which depends, must exclude the intercept. Intuitively, I need to assume that the variance of the signal y_i the baseline product category has signal variance equal to 1; the signal variances of the other categories represent the relative quality of the signal in comparison to the baseline category.

To illustrate how β^u, β^η, and R are identified from purchase and returns data, consider the most simplified example. Assume, that the retailer's assortment contains only one product and thus the customer has only three options: choose the outside option; choose a product and keep it; choose a product and return it. In this case, the model would be described by two parameters: the constant mean utility of choosing a product β_0^u and the return costs R (notice that since I have only one product category its variance of signal quality is normalized to 1). In Appendix B.8 I provide a formal proof that β_0^u and R are identified, but intuitively, I use the two observed probabilities (choosing outside option/choosing product and returning it) to estimate two model parameters.

Next, consider adding to this assortment one additional product and assume that the customer randomly gets one product to purchase or choose outside option. In this case, customer utility would depend on two parameters: $\beta_0^u + \beta_1^u \cdot x_{ij}$ ($x_{ij} \in 0, 1$ indicator of new product). The signal variance would depend on one parameter: $e^{\beta_1^\eta \cdot x_{ij}}$ (the old product is used as baseline and hence its signal variance is normalized to 1). Returns costs depend on one parameter R. In this case I identify β_0^u and R by using the part of the data which involves only the baseline product as discussed in the previous paragraph. Similarly, I use the data containing the new product to identify β_1^u and β_1^η. Adding more products in the assortment repeats the reasoning above, while taking into account the choice set leads to more restrictions and adds more information at the same time. Thus, β^u, β^η, and R are identified from purchase and returns data.

In Section 2.6.2, I demonstrate that the two-stage approach discussed in Section 2.5.4 allows to recover the true parameter values from synthetic data supporting the identification of the model. In Appendix B.9, I use the same synthetic data to demonstrate that maximizing the likelihood function via the one-stage approach allows us to recover the true parameters. I also compare the approach to other popular approaches used in existing literature to justify the choice of the estimation procedure.

2.6.2 Synthetic Data (Monte Carlo Simulations)

I use synthetic data (Monte Carlo simulations) to examine the feasibility of model estimation, its properties, and identification. Synthetic data analyses demonstrate that the likelihood in Equation (2.17) enables us to recover the true parameters of the model. Syn-

Table 2.2. Simulation Data Summary

Variable	Simulation Specification	§
Product mean utility $\mu^u(x_{ij})$	$\mu^u(x_{ij}) = \beta_0^u + \beta_1^u \cdot x_{ij}$	Section 2.5.1
Variance of fit signal quality $\sigma_{\eta_{ij}}$	$\log \sigma_{\eta_{ij}} = \beta_1^{\eta} \cdot x_{ij}$	Section 2.5.1
Product mean search cost $\mu^c(z_{ij})$	$\mu^c(z_{ij}) = \beta_0^c + \beta_1^c \cdot z_{ij}$	Section 2.4.2
Return cost paid by customer R_i	$R_i = R$	Section 2.4.1

thetic data also enable us to compare the estimated parameters of the purchase journey model to classical models of customer search (without a return option). These comparisons illustrate how the option to return products modifies parameter estimates, i.e., some parameter estimates are biased if the model does not allow returns.

I simulated 10,000 synthetic customers according to the search model described in Section 2.4. So that the synthetic data are relevant, I chose parameters that closely mimic the real data. Specifically, the synthetic data have the same marginal probabilities of purchase and return as the real data. I assume that the retailer has two product categories ($x_{ij} \in 0, 1$ is dummy coded product category). The retailer provides the customer with the list of 20 products on the website ($z_{ij} \in 1, 2, \ldots, 20$ is the position of the product on the website where lower number represents higher position on the listing on the website). I summarize the specification of the simulation in Table 2.2.

The simulation has six parameters to be estimated from the synthetic data. In Appendix B.9, I describe in detail the parameters of the synthetic data. For models with returns and without returns, I use the zero vector as a starting point. The results of the estimation with the two-stage approach are shown in Table 2.3.

Table 2.3 demonstrates that the search-purchase-return model, and the model without returns, can be estimated using the two-stage approach and that the true parameters can

Table 2.3. Results of Model Estimation on Simulated Data

Par.	**True parameter value**	**Estimates of the model with returns**	**Estimates of the model without returns**
β_0^u	-1.4	-1.347	-1.568
β_1^u	-0.3	-0.280	-0.099
β_1^η	-0.5	-0.443	–
β_0^c	-4.0	-3.944	-3.360
β_1^c	0.3	0.312	0.312
R	-1.4	-1.427	–

be recovered well for the search-purchase-return model. I expect precision to improve with a larger sample of synthetic customers – precision improves from 5,000 to 10,000 synthetic customers. I tried alternative approaches discussed in Section 2.5.4. Partially closed form integration provided the most precise estimates and required fewer computational resources and time. Alternative estimation strategies performed substantially worse. Details are provided in Appendix B.9.

Table 2.3 compares models with and without modeling returns. For the chosen synthetic data, not modeling returns (a) overestimates the utility of a product, (b) underestimates the absolute value of a category dummy, and (c) slightly underestimates search costs. Naturally, when returns are not modeled, there is no estimate of return costs. These differences are intuitive. Returns have an option value. If they are not modeled, this option value is absorbed in the base utility. Also, when returns are not modeled, it appears to the researcher that search is less costly because the return option makes it more attractive to search. Underestimation of the absolute value of the category dummy is due to the particular parameters of the simulation. Specifically, category decreases the expected utility directly (-0.3 is less than zero) but also increases the utility by providing better signal quality (-0.5 is less than zero). As a result, the customer would purchase this product more frequently because the customer would expect to keep it. As a result, the no-return model would overestimate the absolute value of the utility of this product.

2.7 Estimation Results

Having demonstrated that I can recover known parameters from synthetic data, I proceed to the actual data. Table B.2 reports the estimation results for the model. The first column reports the sensitivity to attributes of the mean utility β^u from Equation (2.6), the second reports sensitivity of the variance of the signal β^η from Equation (2.7), and the estimates of search costs β^c from Equation (2.4). The return costs are at the end of the table.

Table B.2 suggests variation in mean product utilities and quality of the signals across categories. For example, the category "Dresses" has a substantially higher variance of the signal than "Polo Shirts." This difference is not surprising (based on fashion experts) because the choice of a dress is more nuanced than choosing a polo shirt. Customers need more information to make a correct purchase decision. Consistently with Dzyabura et al. (2021), colorful items are harder to evaluate than black and blue items, resulting in a higher return rate online. Interestingly, apparel products made from natural fabrics provide a better signal of the customer preference shock. Likely, synthetic fabrics look attractive based on the online images, but, for some customers, look less attractive when the customer inspects the product at home. Natural fabrics, on the other hand, look more consistently attractive online and in the home.

Table B.2 suggests that the product position on the website affects the search costs substantially. Specifically, a product in the middle of the list has 8.8% higher search costs compared with a product at the beginning of the list. The result is qualitatively consistent with the results in Ursu (2018), where the position of the product on the website affects the search costs in a randomized setting. This implies, for a given set of product characteristics, customers are more likely to click on the first displayed item. Interestingly, mobile device users have lower search costs implying that either it is easier for customers to navigate the website through a mobile device or that those customers are more experienced with search. This is an interesting finding that retailers could use to implement different policies for different versions of the website (desktop vs. mobile).

2.8 Model Implications

2.8.1 Insights Obtained from the Estimated Search-Purchase-Return Model and Comparisons with Data

In Section 2.3.2 I introduced model-free evidence that customer search and product returns are related. These model-free relationships are valuable, but I must not over-interpret them. Some relationships are valuable in the sense that if I observe a type of search behavior, the product is less likely to be returned (selection of customers), others could have causal interpretation. The analytical model and the parameter estimates clarify the qualitative insights about the interrelation between search and returns because the relationships are tied to the underlying parameters about product quality, the quality of the signals, and search-object characteristics. The underlying parameters shed light on the mechanism behind the model-free relationships. By focusing on the underlying parameters, the retailer could use the model to develop potential policies that can be evaluated with policy simulations. As in every model, assumptions were made in developing the model, I recognize that the firm needs A/B tests to confirm the superiority of any proposed policy.

For example, consider the following insight: if the last-clicked product is purchased (the product whose product page the customer viewed prior to purchase), then it is less likely to be returned. This is a description of an empirical relationship in the data. The last click does not cause a lower return, but rather, if I observe that the last-clicked product is purchased, then I observe a lower probability of return. Because of the nature of the purchase journey model, the net utility of the last clicked product is above all non-clicked products; the customer likely found a hidden-gem that matched the customer's tastes. In the model-free evidence, I observe that if the customer searches a variety of colors, the customer is less likely to return the product. A partial explanation of this data-based observation is likely due to idiosyncratic customer characteristics, but part of the explanation is due to empirical search costs, which are lower for additional colors than for additional products.

Table 2.4 summarizes the insights of the model and model-free correlational evidence from Section 2.3.2. For each correlational insights from Section 2.3.2 I provide intuition from the theoretical search-purchase-return model and discuss it in detail in corresponding paragraph in Section 2.8.1. Furthermore, I use counterfactual simulations in Section 2.8.2 to test potential policies that the retailer could use to better manage product returns.

Customers who purchase the last clicked product are less likely to return it. This implies that the customer decided to purchase a product right after viewing it and did not need to look at additional items afterwards. Intuitively, this means that the customer

Table 2.4. Summary of Model Insights on Customer Search and Product Returns

Search Variable	**Corr. sign**	**Model explanation**
Last Viewed Item Purchased	$-$	Customer finds a perfect match product and stop search to make a purchase
# of Product Clicks (Figure 2-2a)	$+$	Customer struggles to find a good product and need to look at many alternatives
# of Search Refinement Tools (Figure 2-2b)	$-$	Customer likes a particular attribute and can search products with that attribute quickly
Time Spent on Product Page (Figure 2-2c)	$-$	Customer extracts a lot of information about the product by consuming the information
# of Product Colors Clicks(Figure 2-2d)	$-$	Customer likes the overall product style and searches for best color fit
# of unique Product Categories (Figure 2-2e)	$+$	Customer browsing the website without particular goal

found a product that matched well the customer's preferences (for example, a "dream dress"). The customer does not want to continue the search as it would only drain valuable time. Having a perfect-match purchase results in a lower return probability. I now provide formal motivation.

Recall that C_i is the index of the last product clicked (viewed) and imagine customer i purchased this product. Thus, from Equation (2.13), I write down all constraints which involve the utility of a purchased product $v_{iC_i}(\psi_{iC_i})$ which is an increasing function of the signal ψ_{iC_i} received after click (Equation (2.3) for reference):

$$\left[\prod_{j=0}^{C_i-1} \mathcal{I}(v_{ij} \leq \min\{\omega_{iC_i}, v_{iC_i}\})\right] \cdot \mathcal{I}(v_{iC_i} \geq \max_{j=C_i+1..S_i} \omega_{ij}) \cdot \mathcal{I}(x'_{ib}\beta^u + \epsilon_{ib} \geq -R) \tag{2.18}$$

All constraints in Equation (2.18) bound the value of $v_i C_i$ from below but do not impose any upper bounds. Similarly, from the same Equation (2.13) it follows that the customer who did not purchase the last item viewed would have the expected purchase utility v_{ib} bounded from above by $\omega_{iC_i}(b \neq C_i)$.

From Equation (2.3) it follows that $v_{ib}(\psi_{ib})$ is an increasing function of ψ_{ib}, thus ψ_{ib} is bounded from above only if $b \neq C_i$. Therefore, customers who purchase the last item viewed ($b = C_i$) received on average higher value of the signal ψ_{ib}. Because the signal ψ_{ib} is positively correlated with product fit ϵ_{ib}, same customers would have on average a higher value of product fit ϵ_{ib} and would be less likely to return the product. Practically, this implies that the information about whether the last clicked product was purchased allows the firm to identify customers who received a very good signal and hence are consequently less likely to return the product.

Customers who make more clicks prior to purchase are more likely to return the product. Intuitively, the customer who clicks on many products is hesitating between different options and does not have a strong preference for any of them. For example, imagine a customer who cannot find an item he or she likes, but ends up purchasing the item with a utility just above the outside option.

More formally, recall the simplified equivalent constraints explaining the search behavior of the customer in Equation (2.13). After dropping the less relevant constrains, I get:

$$\left[\prod_{j=1}^{C_i-1} \mathcal{I}(\omega_{ij} \geq \omega_{ij+1})\right] \left[\prod_{j=0}^{C_i-1} \mathcal{I}(v_{ij} \leq \min\{\omega_{iC_i}, v_{ib}\})\right] \tag{2.19}$$

The right set of constraints implies that all the clicked options' expected purchase utilities (except the last one) are bounded from above by at least ω_{iC_i}. At the same time the left set of constraints implies that the ω_{ij} is a decreasing function of option order click j (according to optimal search rules in Section 2.4.2 customer clicks options in a decreasing order of their reservation utilities). Therefore, on average ω_{iC_i} is a decreasing function of number of searched options C_i (Recall that C_i is the index of the last product clicked). Thus, using the same logic as in discussion of previous insight, a customer, who searched longer (or made more clicks), would likely have a lower upper bound ω_{iC_i} and hence a lower purchase utility with consequent higher return probability.

The retailer observes this information online and may implement policies which would reduce the need to search additional options, say by a recommendation system based on observed search costs. The retailer could also reduce the return rate by showing additional random products for the customer at zero search costs. I explore this policy in Section 2.8.2.

Customers who apply search refining tools have a lower probability of return. Intuitively, consider two customers (A and B) looking for a dress. Customer A does not have particular preferences while B wants a black dress made of a natural fabric. Customer B applies search filter to narrow the search while Customer A is content to search from the default website. Suppose both customers buy the same dress. Customer B is more likely to find the correct match to Customer B's preference and is thus less likely to return the product. Customer A is less likely to find the best match.

More formally, by using search filters, the customer changes the distribution from which to sample the products (for example, browse only products made of natural fabrics). Typically, the application of search filters requires paying an additional search cost beyond those estimated in Table B.2 (for example, navigating through the menu, reading, clicking). This implies that the customer faces a tradeoff: sample from a better distribution by paying additional search costs or sample from the default distribution for free. The model suggests that reducing these additional search costs of applying filters might result in a decrease in returns and an increase in purchases. This change could be tested by the retailer through A/B tests, for example, by encouraging customers to use filter tools with a pop-up window on the website.

Customers who spend substantial time reviewing the product page are less likely to return the product. Intuitively, consider the scenario when a customer opened the product page and started to explore information about products. The customer reads the descriptions, looks at pictures, inspects at the patterns, etc. If the customer searches for a long time, then it is likely that the customer obtained a substantial amount of information about the product – enough to be confident that the product will match customer's needs. Thus, I expect customers who spend on average more time reviewing the products are less likely to return the product.

Mathematically, by consuming each unit of information, the customer gets a signal ψ_Δ about the variance of noise quality σ_{η_Δ} for the cost of c_Δ. The customer could choose how many of these signals to obtain before he or she makes a purchase decision. For example, each signal ψ_Δ might be an additional picture or an additional line of item description. If the customer receives T signals, then the final signal would be $\psi = \sum_t \psi_{\Delta t}/T = \sum_t \eta_{\Delta t}/T = \eta + \epsilon$ Because $\eta \sim \mathcal{N}(0, \sigma_{\eta_\Delta}/T)$ the variance of the noise of the final signal is a decreasing function of time spent. Thus, more time spent leads to higher overall quality of the signal and thus lower return probability.

Customers who browse many colors of the purchased product are less likely to return the product. Intuitively, the customer already likes the style of the product, say a dress, and is now searching for the best color match.

Although the model assumes homogeneous customers, I still can model this scenario by assuming homogeneity in preferences for all characteristics except for color, where I assume heterogeneous preferences. Formally, I assume other characteristics are identical (e.g., shape, style, etc.) and I can write the product utility as the sum of non-color and color sub-utilities $u^{final} = u^{nc} + u^{c}$. The fact that the customer clicked on different colors of the product implies that the expected utility of these different colors was higher than other products. Because the non-color-related characteristics are identical, it is likely that u^{nc} is high and the customer tried to maximize u^{c}.

Although the data and model suggest a relationship that may or may not be causal, the theoretical model suggests that the retailer may enhance purchase and reduce returns by reducing the search costs of alternative colors of a chosen product (increase u^{c} and u^{final}). For example, subject to experimental A/B tests, the retailer might suggest alternative colors to the customer at the time of search or at the checkout.

Customers who browse many colors of the purchased product are less likely to return the product. Intuitively, the customer already likes the style of the product, say a dress, and is now searching for the best color match.

Although the model assumes homogeneous customers, I still can model this scenario by assuming homogeneity in preferences for all characteristics except for color, where I assume heterogeneous preferences. Formally, I assume other characteristics are identical (e.g., shape, style, etc.) and I can write the product utility as the sum of non-color and color sub-utilities $u^{final} = u^{nc} + u^{c}$. The fact that the customer clicked on different colors of the product implies that the expected utility of these different colors was higher than other products. Because the non-color-related characteristics are identical, it is likely that u^{nc} is high and

the customer tried to maximize u^c.

Although the data and model suggest a relationship that may or may not be causal, the theoretical model suggests that the retailer may enhance purchase and reduce returns by reducing the search costs of alternative colors of a chosen product (increase u^c and u^{final}). For example, subject to experimental A/B tests, the retailer might suggest alternative colors to the customer at the time of search or at the checkout.

Customers who click on products from many different categories are more likely to return the product. Intuitively, the customer does not have a particular product in mind that he or she wants to purchase, therefore the customer would browse the website hoping to find something interesting.

As with the color the explanation requires heterogeneous customers, however, with the small update to a homogeneous model I can provide an intuitive mechanism. Customers with a particular product in mind would have a high preference for a particular type of product (for example, short sleeve T-shirts). At the same time, the "uncertain" customer would have more uniform preferences.

Given the same set of products in the search set the former customer would be more likely to click on the products of the preferable category, making the number of unique clicked categories smaller. Moreover, because this customer on average has a higher preference towards this category than the average customer his or her return probability would be lower. Intuitively, customer search helps to infer the preferences of the customer which in turn affect the probability of return. Appendix B.2 extends the empirical result beyond the product category.

2.8.2 Policy Simulations: Implications for Retailer Actions

Explanations of the stylized model-free evidence are based mostly on the structure of the theoretical model, although in some cases the relative size of the estimated parameters. However, other policies represent tradeoffs. Whether or not these policies are profitable depends upon the specific estimate parameters.

Caveats: The data are observations of customers searching in the empirical retail environment. Although the retailer has indicated that it has not designed its website to influence returns, the data are not policy experiments and could be subject to the unobserved retailer's actions. By fitting the search-purchase-return model to the data and estimating underlying parameters of search-object costs, signal variances, and product utilities, I hope that the parameters are close to those which would be obtained if I could have taken the retailer's website design decisions into account. As a result, all of the policy simulations indicative and subject to future A/B empirical tests.

With these caveats, I use the counterfactual simulations to evaluate what would happen if the search environment was changed by the retailer ('what if" simulations) in a manner that affects the estimated parameters of the customer journey model. I assume that the new customer visiting the website would behave according to the estimated model discussed in Section 2.4 and Section 2.5. I sample randomly sets of unobserved (for the researcher) shocks $\epsilon_{ij}, \psi_{ij}, \xi_{ij}$ from corresponding distributions and let the "modeled customer" search for the product.

First, I simulate customers with default parameters of the model, compute average values for the variables of interest (for example, overall return/purchase rate or profit), and use these

Table 2.5. Summary of Policy Simulation Results

Policy	**Main Findings**
Improving the website design to decrease customer search costs	Increase in purchase rate and decrease in return rate
Modifying the effort and time to return a product in the online channel	Increase in purchase rate and increase in return rate
Changing the product ranking on the webpage	Order of displaying product impacts the return rate

as the baseline. Next, I tweak the model parameters (as indicators of changes in the website) and evaluate how this change in parameters affects the average values for the variables of interest compared with the baseline.

Table 2.5 summarizes the results of three illustrative policy simulations. In brief, in each policy simulation, I change a specific aspect of the retailer's website (as reflected in model parameters) and observe the resulting (simulated) customer behavior. I present these recommendations as hypotheses consistent with the data and analytic model, recognizing that I have not fully modeled, nor do I have the data to explore, endogenous retailer decisions.

Some of these implications are intuitive and make sense independently of the model, e.g., website design enhances sales. Others are relatively new – website design can decrease return rates as can reordering the display of products. The latter is important in Europe. Many return policies are dictated by regulation, but not website design. Other implications are intuitive, but the estimated parameters provide a means to evaluate competing interests. For example, lowering the effort to return a product will clearly increase returns at the margin and increase sales at the margin (option value of a return). The model quantifies the effects sufficiently so that managers and judge whether or not the net effect will be profitable.

Improving the website design to decrease customer search costs. The retailer may invest in modifying the website to decrease customer search costs. I consider a hypothetical scenario where the retailer can modify the website to reduce fixed search costs by 5%. Policy simulations predict that lower search costs would benefit both for the customer and the retailer. Specifically, a 5% decrease in search costs results in 9.1% increase in purchase probability and an increase of 1.9% in customer surplus.

Moreover, the model provides additional insights on the positive effect of decreasing search costs. It shows that decreasing the search costs would lead to a decrease in the return rate via allowing the customer to choose a better option. To test this, I assume that the retailer could make the search costs for 1 random product equal to 0 (for example, through a recommendation system). This implies that the customer can always make an additional cost-free click on the product shown. Adding this one additional searched product leads to a 3.8% decrease in the return rate.

Modifying the time and effort to return a product in the online channel. In many countries, product returns are mandated by law and customers have the right to return any product within a specified deadline. However, even in strictly regulated environments, the retailer can make the returns process harder or easier for customers. For example, the retailer may make customers print the label themselves, complete a complicated return form, or require customers to return to a bricks-and-mortar store (as recently implemented by major retailer Zara, BBC 2022).

In the returns literature, it is well documented that the return option (and how lenient or strict it is) affects both the purchase and the return decision of the customer. The model supports and extends these insights. Specifically, the model allows us to distinguish between two reasons for a change in purchase probability. One reason is the improved expected utility of the product as well documented in Anderson et al. (2009). A second reason is that decreased return costs improve the reservation utilities of products and encourage the customer to click on more products. Specifically, a 10% decrease in return costs leads to an 7.8% increase in the number of clicks which translates into a 6.3% increased purchase probability. In the data, the model shows that 37.4% of the total purchase probability increase is explained by the second reason – greater search.

Changing the product ranking on the webpage. From Table B.2 it follows that the customer is more likely to click on products displayed at the top of the website because the search costs are an increasing function of product position. This implies that changing the order in which the retailer displays products on the website could substantially impact customer's search behavior and, by implication, returns. One way to encourage customers to click on a higher-utility products is a greedy algorithm: rank the products based on their mean utility. In this case, the products which are liked by customers are displayed at the top of the website, allowing customers to browse these more easily and thus improve the customer surplus (Ursu, 2018). However, such a greedy product ranking might have a negative effect on product returns.

I use the estimates of the model from Table B.2 to rank products according to their estimated mean product utility without considering returns. The policy simulation suggests that, although the purchase probability increases substantially (by 11.7%), the customer surplus increases only by 0.4%. However, ranking on mean product utility increases the return rate by impressive 9.6% and may negate the revenue improvement.

I can likely do better by taking returns into account. Unfortunately, reservation utilities (considering returns) are nonlinear functions of parameters negating a simple approach.

Figure 2-4. Policy Simulation – Changing the Product Ranking

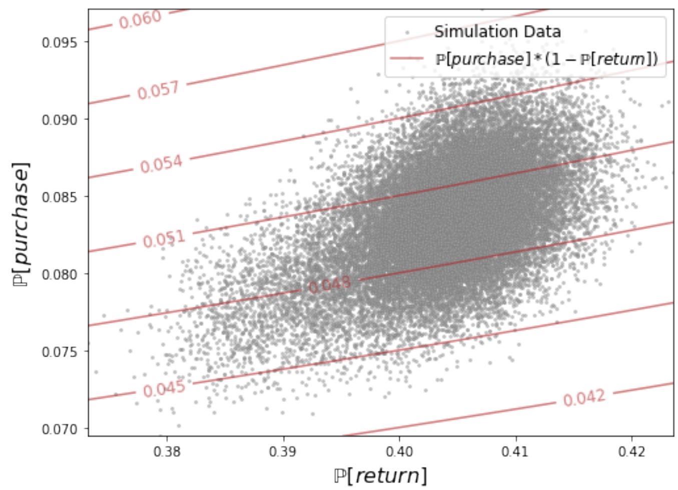

Brute force is also not feasible because the number of rankings is factorial in the number of products. The retailer has 48 products per page on its website, resulting in $48! \simeq 10^61$ possible combinations. Exhaustive exhaustive enumeration is infeasible with current computational power.

To illustrate the potential improvement that is possible with alternative rankings, I randomly selected 48 products and plotted in Figure 2-4 the purchase and return probabilities for 10,000 different rankings of these products. Figure 2-4 illustrates the potential improvement. For this set of 48 products, I observe that the $\mathbb{P}[return]$ ranges within [37%, 42%] and $\mathbb{P}[purchase]$ within [7%, 9.5%]. Without knowing the retailers profit margins and return costs, I cannot choose the optimal ranking from this set. As a further illustration I consider a simple metric, $\mathbb{P}[return] \cdot (1 - \mathbb{P}[return])$, which maximizes sales net of returns. I plot $\mathbb{P}[return] \cdot (1 - \mathbb{P}[return]) = const$ in Figure 2-4 for different value of the *const*. In Appendix B.10, a greedy algorithm based on this metric and applied to 100 random product sets increases this metric by 6.4% relative to a random ranking. Optimization, which is beyond the scope of this paper, would increase the metric even more. Based on these results, I expect that, for any metric the retailer uses to balance purchases and returns, I can improve profits, as evaluated with policy simulations, with a greedy algorithm based on that metric. This improvement is easily achievable and could be optimized. Furthermore, the revised ranking is relatively easy to evaluate with A/B testing vs. ranking by utility alone.

2.9 Conclusions and Future Research

Managing product returns is highly relevant but also challenging. Online retailers, particularly fashion retailers, face high return rates and high return costs. Improving how a retailer can manage product returns has a direct and considerable impact on the firm's bottom line. To the best of my knowledge, retailers and researchers have not investigated the complete search-to-purchase-to-return customer journey to generate insights and suggest strategies by which a retailer can maximize profits. Modeling the customer journey reveals mechanisms by which search, purchase, and returns are related and provides insights that help retailers develop profit improving strategies. Such strategies may not be obvious, or would be evaluated incorrectly, if the entire purchase journey were not modeled.

Using an empirical-theoretical framework, I developed a rational model of customer search in the presence of a return option. I obtained data from a major European apparel online retailer to estimate the model and demonstrate the importance of modeling search and returns jointly. The data and analysis suggest insights on the interrelationship of search, purchase, and returns and suggest strategies by which a retailer can maximize its profits. The model provides explanations of model-free evidence on how search, purchase, and returns are related. More specifically, the model shows that purchasing the last clicked product, browsing multiple product colors, searching more, and using refinement tools are linked to a lower return probability. The retailer's data provide supportive evidence for these insights.

In policy simulations, I illustrate strategies to help the retailer better manage product returns. Some of these strategies such as improving some aspects of the website design are relatively low-cost in comparison to changes in the return policy. In countries with strong customer protection legislation, it is impossible to change the return policy but, generally, firms are not limited in the way they design their websites. I show that reducing the search costs through a more efficient website design decreases the return rate, while increasing the purchase rate. Policy simulations also show that rank-ordering products on the website based on their mean utility increases purchase probabilities, but also return probabilities. The empirical data and estimated parameters enable us to simulate the sizes of effects individually and assess whether, based on prices and profit margins, specific rank-ordering enhances.

Future research. Data on the complete purchase journey are rare. This paper illustrates what can be done with a formal model combined with such data. Hopefully, other researchers will obtain purchase journey data in other product categories and explore implications further. Although the retailer claims it did not design its website to minimize returns, the insights of the model suggest that such designs are feasible. Future research might model the retailers' decisions as endogenous and enable estimation in such data regimes.

Chapter 3

Online Assortment Planning in Presence of Frequent Product Returns

3.1 Introduction

Assortment planning is a major problem for a fashion retailer. The successful assortment must balance various customer needs: excite customers by providing a high variety of products; follow the fashion and seasonal trends; prevent the stock-out in a period of high demand; provide a wide range of available sizes, etc. All of these should be addressed under the physical constraints of the store size and the massive assortments of the retailer.

Given that the total number of active products could reach up to 20,000 and growing (Quelch and Kenny, 1994), choosing which products to put on the website is a very challenging problem. The offline retailer typically cannot precisely know which products the customer was considering during the store visit – thus it is hard to identify whether the product with low sales does not match customers' needs or customers just can't find it on the shelf. A/B testing with assortment is barely feasible. For example, if the retailer wants to remove one product from the assortment it not only has to remove all units from the shelves but also fill the empty shelves with alternative products substantially increasing the price of the experiment. This situation becomes even more complicated given the huge assortments and interrelation between products in the assortment (removing the product from the shelf could negatively impact complementary products).

In online channels the situation is different. The retailer is not constrained by the physical characteristics of the store and can put any number of products on its website. Indeed, many retailers tend to put all available products on online websites. However, due to a large number of products, customers typically review only a small subset of the assortment (for example, the user searches on average 164 products, while making on average 2.834 clicks in my data)[1]. This implies that online retailer is constrained not by the physical characteristics of the store but by the customer's willingness to spend time on the website. The latter is endogenous (customer's decision) and could depend on the assortment itself.

Moreover, the online retailer observes an important part of the customer journey – search for a product. Typically, retailers store all the browsing data and know exactly, which product the customer viewed, which areas of the webstie the customer clicked, and the order. This means that the retailer could distinguish between" low sales because a product does not serve customer needs" and "low sales because the product is hard to find." This information could be crucial in assortment planning.

The online channel is more flexible and randomized experiments with assortment are

[1]By searched product I imply any product which appeared on the customers web screen (see Figure 3-1a); by clicked product I imply any product which the customer clicked on to reveal additional information (see Figure 3-1b)

substantially more feasible. Specifically, to test the new assortment the retailer may just remove the product link from the online website without any impact on the physical product. Moreover, the firm can conduct an experiment on the customer level by manipulating assortment individually. Offline only store level policies are possible. This implies that the online retailer has lower costs and risks associated with experimenting with the assortment as all changes are completely reversible.

Lastly, knowing the exact customer history would allow the retailer to design a new approach to assortment planning infeasible in offline channels, for example, personalized assortments.

Finally, online the customers' journey typically does not end with a purchase. Indeed, in online after a purchase, customers return products with a very high probability (for example, 55% online vs. 3% offline). These frequent product returns could substantially negatively impact the retailer's profit due to the costs associated with product returns (dry cleaning, return label, damage to product, shipping, restocking). Therefore, the retailer must take into account the possibility of returns in assortment planning.

In this paper, I combine the advantages and challenges of online channels and propose a deep-learning-based solution to the online assortment planning problem in presence of product returns. Specifically, I develop a predictive model of the customer search-to-return journey which models how customers would make their click/purchase/return decisions given the structure of the website. Further, I use this model to quantify the changes in the retailer's assortment and give recommendations on whether to keep a product in the online assortment or not. Lastly, I give qualitative insights on which products are the least successful on the website.

3.2 Related Literature

In Section 1.1 and Chapter 2 I discuss customer search and product returns literature in detail. This section provides a brief overview of product assortment planning literature. An extensive review of the field could be found in Kök et al. (2009) and Mantrala et al. (2009).

Assortment planning is a mature multidisciplinary field. Scholars and practitioners recognize the impact of an assortment on the firm's well-being. Methodologically, product assortment planning is studied both from a static and dynamic perspective. Static models consider the problem as a single period, and based on a different variation of customer choice models (Mahajan and van Ryzin, 1999; Kök and Xu, 2011) The main challenge of these models is the discrete nature of the problem with a huge set of possible combinations (2^M where M is the total number of products). Thus, the authors propose various approximations and heuristics to make the solution feasible. The dynamic model extends the finding from static models by allowing the firm to change the assortment between periods and learning customer demand (Chen et al., 2021; Sauré and Zeevi, 2013; Caro and Gallien, 2007). Typically, these models focus on customer purchase behavior and ignore customer search and product returns. However, Cachon et al. (2005) demonstrates that ignoring the customer search in assortment planning could lead to reduced profit and unnecessary narrow assortments. Wang and Sahin (2018) considers the assortment model with customer search costs. Alptekinoğlu and Grasas (2014) shows that firms' product return policy would impact the optimal assortment. Recent empirical research utilizes advances in machine learning and provides a data-driven solution to assortment optimization problem (Chen et al., 2023). In

Rooderkerk and Kök (2019) authors discuss the challenges faced by an omni-channel retailer with regard to assortment planning.

This paper contributes to assortment planning literature by providing a deep learning framework that takes into account three stages of the customer journey – search, purchase, and returns, and allows us to evaluate how changes in the assortment would affect the retailer's well-being.

3.3 Data and Model-Free Evidence

3.3.1 Data Used in the Paper

To train and test the model I obtained individual-level data from a large European apparel retailer that distributed products through online and offline channels. In accordance with the goal of this paper, I focus exclusively on the online channel. The retailer specialized in mid-price fashion products for men and women[2]. It has a generous return policy typical for Western Europe – each product could be returned by the customer without a fee within 60 days after the purchase. Data capture detailed information about customer search-to-return journey and consists of three main parts:

- ***Click stream data*** captures a sequence of actions made by the customer during the browsing session. It includes the list of products that customers view during the browsing session as well as products' positions on the webpage Figure 3-1a); the order of customer product clicks to reveal additional information (see Figure 3-1b); time

[2]The firm also sells apparel for kids and bedroom accessories. However, the share of these products is approximately 1% of total sales.

Figure 3-1. Example of Information Observed by a Customer Visiting the Website

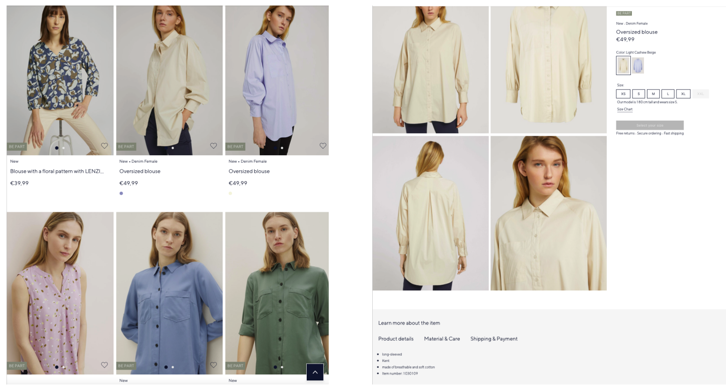

(a) Listing of products **(b)** Detailed product page

spent reviewing each product page; tools used to reinforce the search (filter/sort by price, category).

- ***Transaction data*** captures products purchased by the customer and identificator if the customer returned the product to the retailer. It is matched to the ***Click stream data*** via the unique identifier.
- ***Product reference data*** captures product level information and includes category, price, size, color, and product image as it is displayed on the website.

The observation period for this dataset is from 1 October 2019 to 29 February 2020[3]. The original dataset was prepossessed to remove outliers (for example, sessions without product views and longer than 3 hours were excluded from the observation). The final dataset includes 2,897,667 unique browsing sessions, where customers were presented with 11,523 unique products to choose from. The customer purchased at least one product in 142,302 (5.63%) cases and among these transactions, the customer returned at least one product in 78,563 (55.21%) of cases. This high number of product returns is consistent with the existing

[3]The retailer provided us with the data up until 15 May, however, due to the Covid-19 pandemic I excluded all months after February

Figure 3-2. Summary Statistics of Search-Purchase-Return

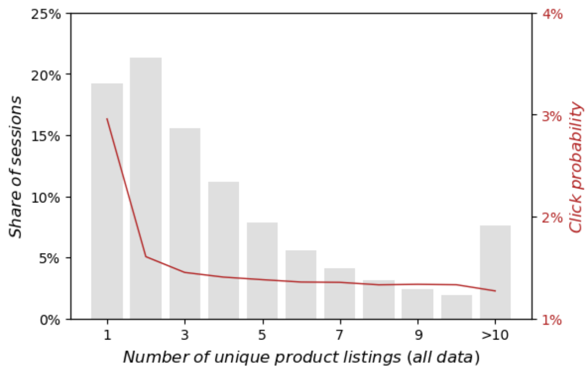

(a) Listings and clicks

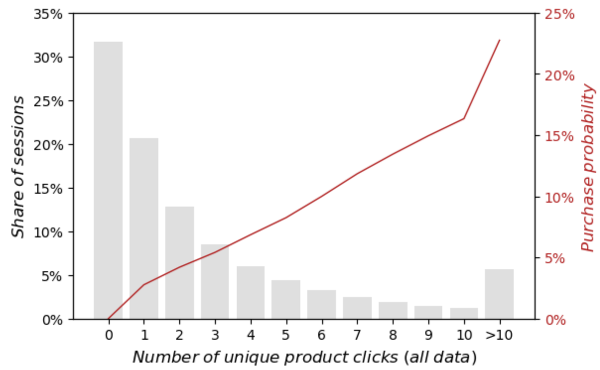

(b) Clicks and purchases

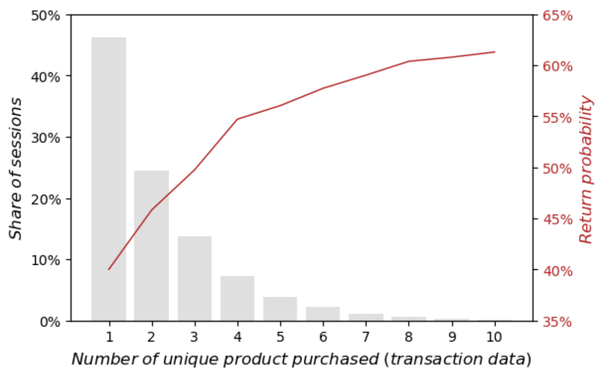

(c) Purchases and returns

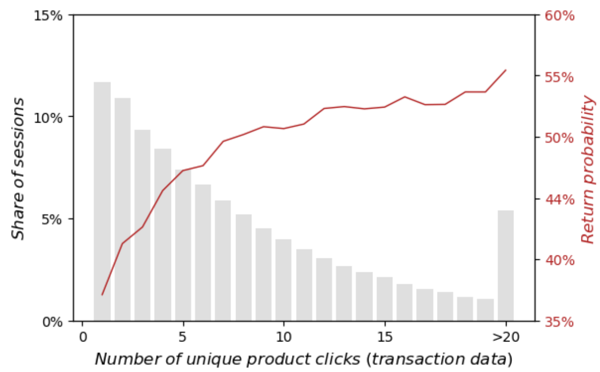

(d) Clicks and returns

literature.

Figure 3-2 provides additional information about the data used in the paper. As we can see half of the customers view less than three product listings (approximately 144 products)[4]. Figures 3-2b to 3-2d highlight the importance of modeling the entire search-to-return journey. Specifically, the figures demonstrate that information observed on one stage (for example, search) could be relevant to the subsequent stages. For example, customers who purchase many products are substantially more likely to return the product (Figure 3-2c. Therefore, the successful model of the customer journey must take into account the set of products the customer purchased but not look at purchased products as independent events.

[4] The retailer's website layout places 48 unique products in one product listing using the 3x12 grid.

3.3.2 Model-Free Evidence

Figure 3-2 provides the summary statistics on the session (or individual) level and suggests that search, purchase, and return stages are interdependent. However, assortment planning is typically a product-level policy. Therefore, in order to plan the assortment there must be enough heterogeneity between products. For example, consider an extreme scenario when all products in the assortment are identical. In this case, there is no possible improvement to the assortment due to all products being identical. However, it is still possible that customers who purchase several (identical) products would have a higher probability to return any of these products.

To see that there is enough variation between products consider three product-level summary statistics, each corresponding to one of the stages in the customer journey:

- ***Click rate*** – the share of the customers who click on the product to review additional information about the product after they see it in a product listing
- ***Purchase rate*** – the share of customers who purchased the product after they clicked on the product page [5]
- ***Return rate*** – the share of customers who returned the product after the purchase for a full refund

These three measures correspond to the quality of the product at different stages of the customer journey. Click rate demonstrates how "catchy" the product is on the website, the purchase rate illustrates how good the product is after the revealing of additional information, and the return rate summarizes whether the product serves the customers' needs after the

[5]For this paper I assume that the purchase could happen only after the click. The extension to fast-checkout case is straightforward and only constrained by the computational capabilities of the machine.

Figure 3-3. Summary Statistics of Product-Level Characteristics

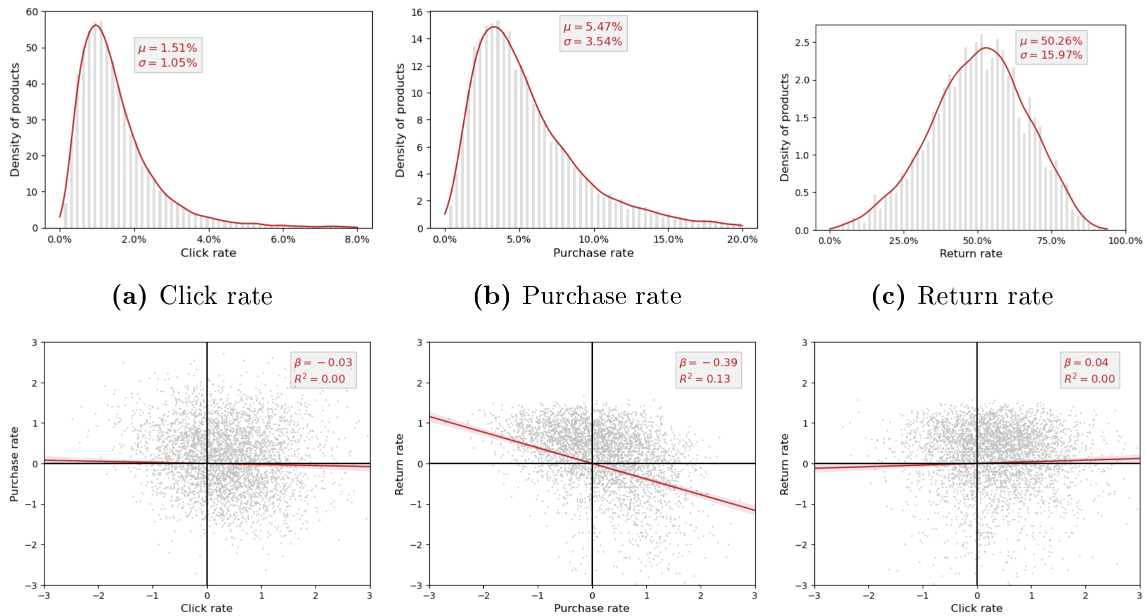

(a) Click rate **(b)** Purchase rate **(c)** Return rate

(d) Click and purchase rate **(e)** Purchase and return rate **(f)** Click and return rate

customer experiences the product. For display, all measures are all normalized to the [0,1] interval and could be interpreted as efficiencies (for example, efficiency of converging the customer from view to click or click to purchase)

Figure 3-3 presents the distribution of each outcome. The variation in these summary statistics is quite high, indicating the potential improvements of removing some products from the assortment. For example, the return rate ranges from 10% to 90% which is consistent with the previous research.

Finally, it could be tempting to concentrate all optimization efforts on one outcome or consider only a one-stage search model. For example, one may think that the retailer should just remove the products which are not popular among customers (have a low *Click rate*). Such a focused optimization effort assumes that all three variables should be highly correlated – a good product might attract many clicks, be purchased frequently, and be rarely returned.

But these measures and not highly correlated. The empirical results in Figures 3-3d to 3-3f demonstrate that looking only at the search component is not enough. Removing low-click products would cause the retailer to lose many high-sale (and low-return) products.

3.4 Deep Learning Model of Customer Search, Purchase, and Return

3.4.1 Model Overview

This paper considers a three-stage process of the customer journey: search, purchase, and return. The search stage starts when the customer visits the website. On the website, the customer observes the sequence of product listings each containing 48 products Figure 3-1a). The customer may click on any of the products in the listing to obtain additional information about the product (for example, a high-quality image of the product, or a detailed description, seesee Figure 3-1b. After the customer finishes the clicks on the products, the customer decides which products to purchase from the list of clicked products. The customer may purchase any number of products from the set of clicked products. Finally, when the customer receives all products at home he or she makes a decision about which products to return back to the retailer.

he described journey suggests that the predictive model should have three components that would mimic each stage of the customer journey. Figure 3-2; Figure 3-3 and established research suggest all of these blocks are interrelated (Ibragimov, 2023; Petersen and Kumar, 2009). A good predictive model should share information from different stages. The high-

Figure 3-4. The Structure of the Predictive Model

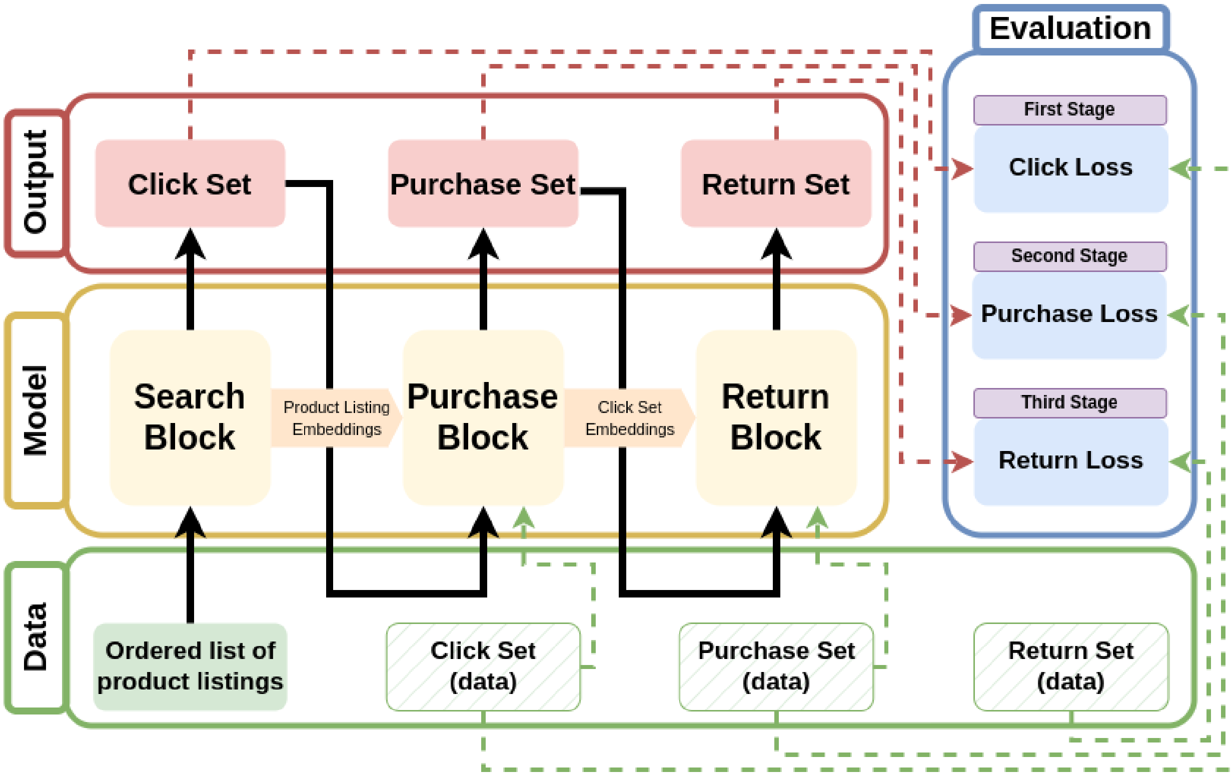

level structure of the model is shown in Figure 3-4.

The model consists of three blocks (yellow) and outputs three sets of variables (red) – click, purchase, and return sets of products. The input of the model (green) depends on whether the model is trained (dashed arrows) or used for modeling (solid arrows).

During the training process, for example, the purchase block has access to the true click set and thus was predicting the purchase set given the true click set. Notice that the true click set is unobservable to the model during the inference stage as it is one of the outputs and formally should not be used as input. Indeed, theoretically, we can use only the true input and generate three sets of variables, and simply compare those with the true sets (for example, exact match). However, such approach would be prone to error propagation as errors in generating the click set would result in bigger errors in generating the purchase set.

Therefore, it is important to use all the available information during training.

Search block. he search block takes as input the list of product listings observed by the customer (each product listing contains at most 48 products). The block has two outputs: (1) an ordered list of products on which the customer clicks during the browsing session; (2) embeddings of the product listings viewed by the customer.

The block structure is based on transformer models (Vaswani et al., 2017) which were successfully used in machine translation problems. In these tasks, the machine is required to transform the ordered sequence of words from the language $\mathcal{I}$ to the ordered sequence of words from language $\mathcal{O}$. Notice the similarity of the translation problem to the prediction of the click set. The customer views the product on the website in batches – product listings (Figure 3-1a). From each product listing the customer can clicks on any on the products. Intuitively, the machine must "translate" the ordered sequence of product listings (language $\mathcal{I}$) to the ordered list of customer-clicked products (language $\mathcal{O}$). As in classical translation problems clicks could have long dependencies both on the viewed product listings and previous clicks (for example, if a customer clicked several times on a T-shirt he may be looking for a T-shirt and thus probability of clicking in this category would be higher). The paper considers the listings-to-clicks translation to better reflect the customer search process on the provided data and correspond to a more general case. The model could be straightforwardly updated to the case when the customer observes all products at once (or "one-by-one" design). by making each product listing size one. Practically, many retailers prefer to split the product into product listings to help customer navigation.

The main idea behind the transformer model is the attention mechanism. This mechanism automatically learns the relevance of each input word in a very flexible manner.

The module itself operates on three important components: key, query, and value. Each of these components represents a mapping (typically linear) from the embedding space to key/query/value spaces respectively. These mappings are applied to each input word and thus the sentence with five words would be represented by five key/query/value combinations. The output of the attention module is represented as a weighted sum of values, where the weights are some increasing function of the similarity between key and query. The advantage of this approach is that it (1) allows propagating the relevant information both within the input by mapping input to key/query/value (self-attention); (2) extracts relevant information between two inputs by mapping input one to value/key and input two to query. Combining multiple attention blocks allows us to learn different relevancies and approximate the output.

Although the transformer approach performs extremely well on machine translation tasks it could not be directly implemented in the customer click set prediction problem due to the number of unique product listings the customer observes. Specifically, in the data customers visited 7,083,413 unique product pages, while typically the machine translation literature works with vocabulary sizes of 20,000[6]. This huge size of unique-product vocabulary results in a very sparse input sequence and an enormous number of parameters. To address this issue, I notice that the product listing is an ordered list of products or in translation terms – each word in the language $\mathcal{I}$ is an ordered list of words from language $\mathcal{O}$. This implies that there exists a structural dependence between different product listing embeddings. Instead of learning this structural relation in a model-free way, we can implement those in the structure

[6]Note that the product listing is an ordered list of products – different order of the same set of products would result in the different product page.

of the model. For example, if two product listings are different only in one product on the last position, it is very likely that their impact on the model would be very similar. A good model would learn these relations with infinite data, however, due to sparsity it could be hard on finite data. Therefore, we can help the model learn the dependencies by preserving the structure of the relations. This would substantially reduce sparsity and the number of parameters.

To preserve and learn the relation between product within one product listings I use an additional attention mechanism to construct the product listing embedding. This approach allows one to learn the dependencies of products on the product listing (for example, the variety of products) and helps to preserve the information on product position on the web page. In contrast, consider averaging the product embedding within the product listing – in this case, all the products have equal weight and all information about the product position is lost.

Purchase block. The purchase block takes as input the ordered sequence of customer clicks and product listing embedding. It has two outputs: (1) probability of each product being purchased by the customer; (2) embedding of customer click set.

The structure of this block is substantially simpler and reminds the classical prediction problem. Specifically, for each clicked product the model predicts the probability with which the customer would purchase this product. As in the click prediction, the decision to purchase the product may depend on any of the clicked products. Therefore, I implement an additional attention mechanism on the click set embeddings to preserve and learn the distant information.

Return block. The return block takes as input the set of purchased products by the customer clicks and clicks set embeddings. It outputs the probability with which each purchased product would be returned.

The return block copies the structure of the purchase block except that it has fewer parameters due to the fewer data (only 3.82% sessions results in a purchase).

3.4.2 Model Training

The three-block structure of the model poses additional challenges to estimating the model. Firstly, the model has three different outputs each of which has its own loss function. Secondly, each block has a different efficient size of the data which could be used for estimation. For example, the customer makes at least one click only in 68.50% of cases, and purchases approximately in 3.82% of cases with at least one product click. This substantially decreases the efficient sample sizes for purchase and return blocks (for example, before the return decision the customer should first click and then purchase a product).

These two reasons imply that simply minimizing the sum of losses would allow the model to ignore parameters that lead to smaller changes in the total loss, or in our case, the parameters needed for smaller parts of the data. This could lead to the situation when the model, in order to marginally increase the performance on predicting customer clicks, would completely ignore the accuracy of predicting the return probability. Moreover, the convergence speed of different blocks could be different, and thus an optimal number of epochs could vary depending on the block.

To address these issues I estimate models sequentially. First, estimate the click model using the data on customer search. Fix the weights of the search block and estimate the purchase block on a subset of the data where the customer made at least one click. Finally, fix the weights of the search and purchase blocks and estimate the returns block on the subset of data where the customer made at least one click. In this case, each step involves only one loss function and thus the estimation is straightforward. During the training, I did only one pass of the whole model training (however, each block has multiple epochs of training). This was done to increase the generalizability (out-of-sample performance) of the model, however, the framework could be straightforwardly extended to multiple-pass training.

Notice that despite the model blocks being estimated separately it is still a joint model. The subsequent block still takes as input one output of the previous block (for example, click set embedding for returns block). The only difference is that the previous block information is fixed for each customer during the training but still used as input. Moreover, product embedding vectors are fixed at the click stage and used by all blocks.

Individual-level prediction is a hard problem. The situation is even worse as in our case the set outcome is very imbalanced: the most frequent click option is the outside option (stop search); the most frequent purchase option is no purchase (leave a website without a purchase). This imbalance substantially shrinks the predictions towards outside options, making the simulations very short (customer clicks on fewer products). Shrinkage makes the search sessions substantially shorter. Given the consecutive block structure of the model, the next section is impacted by the previous one. As a result, if the modeled customer would make fewer clicks, then the modeled customer would have a small choice set and thus the transaction probability would be much smaller. To address this issue I implement the label smoothing technique (Vaswani et al., 2017), where we replace the binary outcome of the correct product with the probability distribution over possible products. Intuitively, it

Figure 3-5. Train and Validation Loss

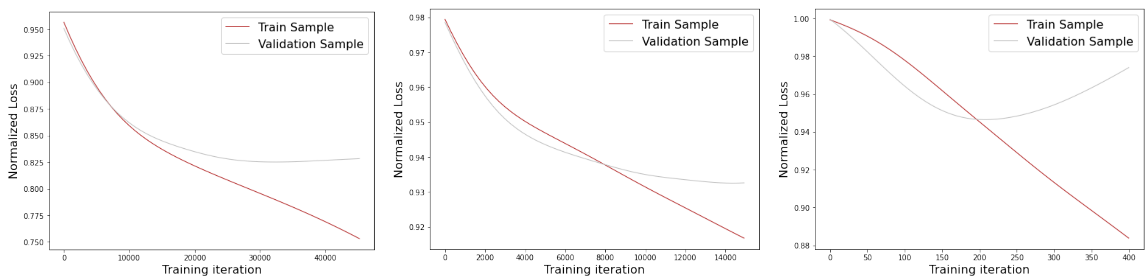

(a) Click block **(b)** Purchase block **(c)** Return block

prevents the model from becoming misleadingly confident about the predictions and helps the model learn the dependencies more uniformly. To evaluate model training, I report both smoothed and true loss functions and show that the model improves on both.

During the training, the data were split into 3 parts: training, validation, and testing. The training (estimation of the parameters) was done on the training data, the validation data was used to stop the training process. All simulations and model evaluations are made on the testing data.

Figure 3-5 plots the performance of the model on training and validation data as a function of training time. Because the performance of the model in all three cases has a form of a U-shape for validation loss, we stop the training when the performance on the validation sample started to decrease.

3.4.3 Model Performance

The estimated model was able to extract relevant information from data. Table 3.1 reports the results of the model estimation with the smoothed and true performance measures. In Table 3.1 the uniform benchmark represents the customer who chooses each product randomly with equal probability except for the outside option which is chosen with the prob-

Table 3.1. Performance of the Model

	Click block	Purchase block	Return Block
Uniform benchmark	3.331	0.100	0.704
Model (% of uniform benchmark)	81.97%	93.20%	94.39%
Model (smoothed outcome)	2.008	–	–
Model (smoothed outcome, %)	79.92%	–	–

Note: first and third rows represent the negative loglikelihood (the smaller the better), while the second and fourth represent the share of loglikelihood remained unexplained by the model (the smaller the better).

ability observed on the data (this benchmark performed better than a completely random choice model)

The model successfully improves in all blocks indicating that the model successfully recovered important relations in the data. Table 3.1 demonstrates that optimizing the smoothed likelihood function does not hurt the performance measured by a true likelihood.

From Table 3.1 it follows that the performance on the individual level is far from perfect. This observation is not surprising as click/purchase/return prediction on an individual level is particularly hard due to individual-level unobserved preference shocks. This makes the behavior of the customer not completely deterministic. Contrast this to natural language models, where the language rules make the problem deterministic with a very limited number of ground truth possibilities (given the input sentence there are typically a couple of correct translations). To see how the model performs on average, I put the aggregated summary statistics in Table 3.2 as well as their true values on the test data[7].

The results in Table 3.2 demonstrate that, on average, the model improves predictions in comparison to random mode. Interestingly, the model captures non-trivial relationships in the data as plotted in Figure 3-6, replicating the observations on the true data in Figure 3-2[8].

[7]The R^2 is measured as weighted sum according to the denominator. For example, for a return rate higher contribution to R^2 would give products with higher sales

[8]Note that the dependencies represent per item value. For example, in Figure 3-6c the y-axis demonstrate

Table 3.2. Aggregate Performance of the Model

	Validation Prediction	True Value
Average Number of Clicks	2.91	2.834
Purchase Rate	3.67%	3.74%
Return rate	48.71%	49.45%
Product click rate R^2	71.50%	–
Product purchase rate R^2	61.87%	–
Product return rate R^2	26.68%	–

Figure 3-6. Customer Level Relation between Clicks, Purchases and Returns

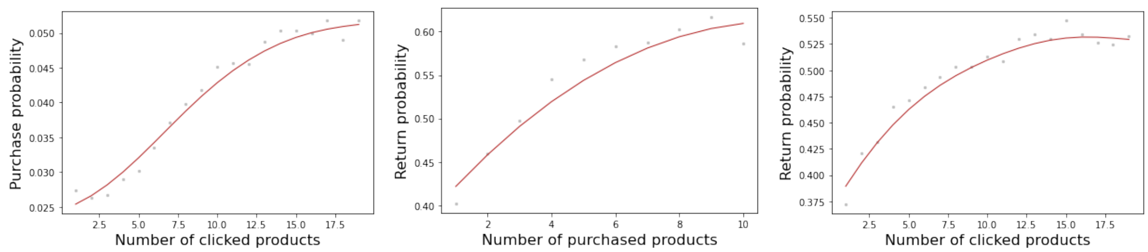

(a) Number of clicks and predicted purchase probability

(b) Number of clicks and predicted return probability

(c) Number of clicks and predicted return probability

3.5 Simulation Results

3.5.1 Modeling Changes in the Assortment

Our model allows us to describe customer behavior from the visit to the website to the return of the product. We now use the model to address the optimization of assortment.

I implement a counterfactual simulation approach. I simulate customer behavior according to the trained model to evaluate what would happen if the retailer would change the assortment of products on the website. The data provided by the retailer does not have randomized variation in the assortments which poses threats to the external validity of the results. To address this issue I consider only a very small deviation from the trained model, which is unlikely to cause major disruption in the customer journey. Under the assumption

that in sessions with multiple products purchased the probability of return for **each** product would be higher.

that the model describes the customer journey accurately within minor changes to the website, we can implement 'what if" simulations. Further research would address the outlined issue by, for example, randomly changing the assortment for customers and introducing an exogenous source of variation.

The procedure used in the paper is (see Figure 3-7 for visualization):

1. Using the simulation results on the validation sample, compute product-level click, purchase, and return rates.
2. Split all products in the assortment into 7×7 equally sized bins based on predicted purchase × return rates. Denote all products in bin b as $\mathcal{S}_b$.
3. Remove products in set $\mathcal{S}_b$ from the website by hiding them from products that could be viewed by the customer(Figure 3-7b). However, to keep the models comparable, I replace the removed product with a random product from the list of kept products (Figure 3-7c)
4. Simulate customers on the updated website. Compute average values for variables of interest (for example, overall return/purchase rate or profit).
5. Repeat (1)-(4) for each bin b.
6. Simulate customers visiting the default version of the website (for example, if the website has four products on the web page, then the customer sees Figure 3-7a). Compute baseline values for variables in (4)
7. Compare values in (4) with corresponding values in (6)

There are several important observations with regard to the procedure discussed above. First, notice that the simulation tries to evaluate the past decision of the retailer to add products to the website. That is, in (4) I obtain the value for the website with old assortment

Figure 3-7. Example of Changing the Assortment Online

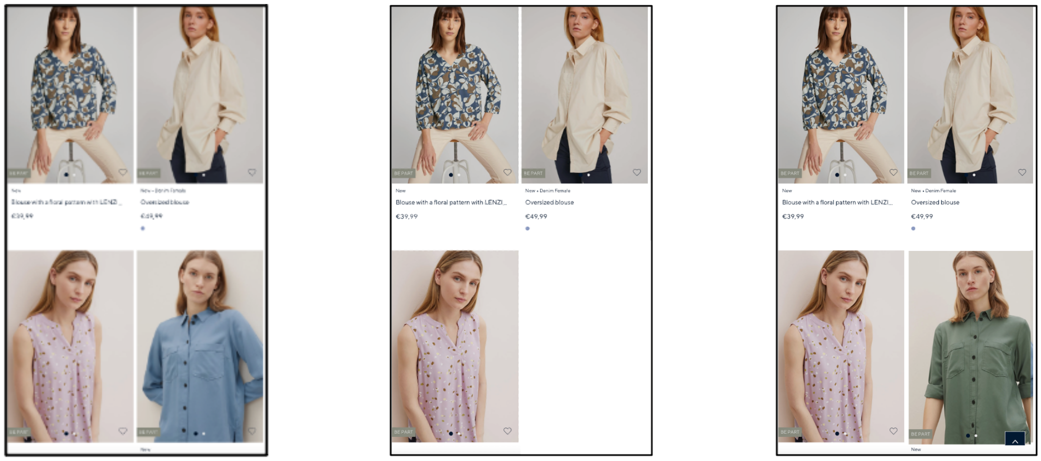

(a) Original website viewed by the customer **(b)** Forth product removed from the website **(c)** Empty position replaced with the random product

while in (6) for full assortment after the retailer launched some products. Intuitively, I want to understand with which products the retailer made a mistake, or which products have a negative impact on the retailer's well-being. The more practical approach would be to evaluate which products the retailer should remove from the website. In this case, the interpretation is more involved while the identified products are the same. I put the details in Appendix C.1.

Second, in each simulation, the customer sees the same number of products on the website and each simulation removes a small portion of the assortment (approximately 2% given 49 bins). This insures the change to the website is minor and is unlikely to cause structural changes in the customer journey. This is supported by the literature where Boatwright and Nunes (2001) found that reducing the assortment does not cause reduction in perceived variety (see also Broniarczyk et al. (1998)).

Third, the decision on which products to remove is based on the validation sample, making the simulation out-of-sample. This is done to reduce the impact of the size of the validation data on the quality of estimating the purchase and return rate. Moreover, with a small modification to the product embedding mechanism the retailer would be able to test how adding a new product to the assortment would impact future sales and returns. The criteria to remove the products does not depend on the product characteristics thus avoiding the problem of excluding the whole category (for example, dresses) from the website. Lastly, the choice to remove the products based on their predicted purchase/return rates is arbitrary and could be substituted by any other procedure (for example, if the retailer has access to costs it may simulate the exclusion of the least profitable products from the website). In Appendix C.2 I provide alternative results where instead of predicted value I use true purchase and return rate estimate on combination of train and validation data.

I plot the results of the simulation in Figure 3-8. Each point in the graph represents how adding the product group to the assortment changed the variables of interest[9].

In Figure 3-8, the products are split into four groups depending on the quadrant. Each of these groups has different practical implications for the retailer.

- ***Quadrants I and III, upper right and lower left,*** demonstrate that there is a tradeoff between sales and returns for some products. That is some products might be sold frequently, but at the cost of having high returns. Adding these products (lower left) means the remaining products are sold less, but also returned less. For these products, and for the products in the upper right there is no clear conclusion on whether they should be kept or removed from the website. The retailer needs further information such as costs of selling and costs of returns to make a decision.

[9] Notice that all are small (1%). Therefore, we can use the following approximation for the relative change: $\Delta x/x_1 \simeq \Delta x/x_2$. Thus, impact of adding product to the website is equivalent to a negative impact of removing product from the current version of the website.

Figure 3-8. Results of Assortment Simulation

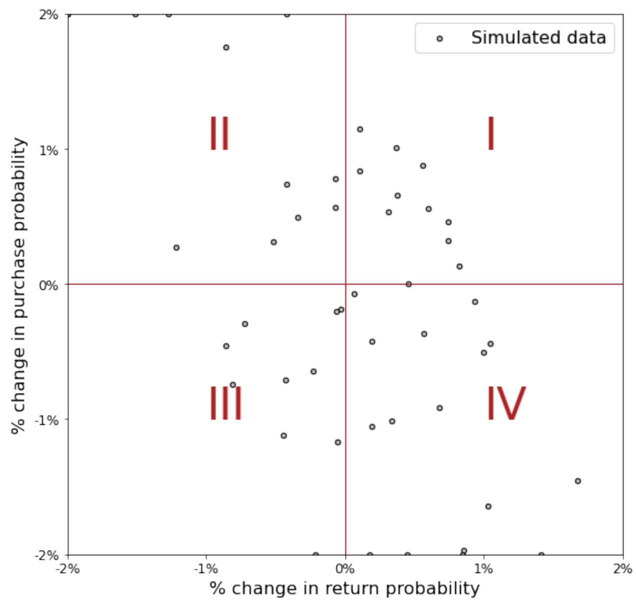

- ***Quadrant II, lower right***, represents the "best" products. These are the products that are sold frequently and are returned rarely. Adding them means higher sales and smaller return rates. The retailer should definitely keep those in the assortment.
- ***Quadrant IV, upper left,*** represents products that are clearly unprofitable for the retailer. Adding these products to the website made the retailer worse off. Specifically, the retailer suffers from lower sales and higher product returns. Intuitively these are the products that are cursed with a double whammy, they are not sufficiently attractive to be sold at high rates, but, when sold, they do not deliver on their promise. When these products are introduced to the website on the website they take the space on the webpage, thus increasing the search costs for other better products. It is recommended to avoid adding products from Quadrant IV to the website.

The results in Figure 3-8 demonstrate the three types of products, however, it does not provide intuition on the characteristics of these products. To answer this question, I consider an auxiliary prediction problem to visualize the characteristics of the products which lie within Quadrants II and IV. These segments represent two extremes and imply a clear planning decision.

Notice that in the simulation I split the product into 49 groups to reduce the computational burden. This implies that for each product within one group, the prediction and recommendation would be identical. Moreover, it is possible that some products appeared in the corresponding segment arbitrarily due to the noise in the data. This implies that it is infeasible to plot "Top-N" products from each segment as it won't be possible to distinguish between products within the group.

To produce different product level predictions I extract the product embedding from the model and predict the probability that the product belongs to each of these quadrants. Heterogeneity in product features (embeddings) would produce the variance for product-level predictions. Moreover, regularization would reduce the impact of "incorrectly" classified products. In Figure 3-9, I put examples of these products for top-9 categories based on total sales. Specifically, each row represents one of the nine product categories and the product most likely to be in Quadrant II/IV respectively (most likely is equivalent to having the highest of out-of-sample predicted probability). Further research could extend the analysis and provide more interpretable insight, for example, by using the interpretable characteristics (both quantitative and visual) and SHAP values.

Interestingly, even without precise quantitative analysis, it is clear from Figure 3-9 that products in Quadrant II and IV are substantially different. Specifically, products on the left include much more "non-standard" colors (pink, yellow, green). This may indicate that these colors draw attention from the customers and thus they are more likely to click on these

Figure 3-9. Top-5 Most Likely to be in Quadrant II and IV Products for Different Categories

(a) Quadrant IV **(b)** Quadrant II

products, however, customers are reluctant to purchase these high-risk products. Moreover, in the rare case of the purchase the customer is not satisfied with the product and eventually returns it back to the retailer. In Appendix C.3 I demonstrate that the products ranked solely on transaction data would be different indicating that the search component is important. Intuitively, the "worst" product is not only the product with low sales and high returns, but also the product is very noticeable by the customer and drags their attention from more successful products.

3.6 Conclusion

Assortment planning in online channels is a highly relevant but challenging problem. In comparison to classical retail, online retailers have more flexibility in designing the assortment due to the lack of physical constraints of the store. Moreover, most online retailers have access to customer search data and know precisely which product the customer was considering. However, online retailer faces different challenges – high product return rates. Therefore, the successful assortment should take into account the advantages (search data) and solve challenges (product returns) jointly.

In this paper, I propose a deep-learning solution to the assortment planning problem in the presence of frequent product returns. Specifically, I develop a model that approximates the customer journey from the moment he or she visits the website to the decision if the customer likes the product and wants to keep it.

I demonstrate that the model based on the transformer framework can extract relevant information about customer behavior. I use the model in counterfactual analysis to show that the retailer could identify both problematic and successful products. Finally, I unveil the black box by providing qualitative information on which characteristics are shared by the most/least successful products.

Appendix A

Appendix for Chapter 1

A.1 Tuning of Hyperparameters of the GBRT Model

We tuned the hyperparameters of the GBRT model with a grid search over the set of parameters presented in Table A.1. The criterion was predictive ability in the validation sample. The model was then tested using the held-out out-of-sample predictions. For each iteration of the grid search, we stopped adding additional regression trees after the accuracy on the validation sample did not improve for twenty-five consecutive trees.

A.2 Profit-Maximizing Policy

Result 1. Suppose (1) the firm's prior on the profitability of an item, π, is normally distributed, $\pi \sim \mathcal{N}(\mu_0, \sigma_0^2)$, (2) the firm observes an estimate of profitability $\widehat{\pi}|\pi \sim \mathcal{N}(\pi, \sigma_1^2)$, and (3) the firm seeks a policy to decide whether to put an item online or not. Then the profit maximizing policy, $\mathcal{P}(\phi)$, is a threshold policy:

Table A.1. Grid for the GBRT Hyperparameters

LightGBM parameter name	Set of tested values	Parameter Description
n_estimators	[3000]	Maximum number of boosting trees
learning_rate	[0.01,0.025,0.1]	Shrinkage rate
max_depth	[7,9,11]	Maximum depths of the regression tree
num_leaves	[32,48]	Maximum number of leaves in one regression tree
reg_lambda	[0,5]	Weight of L2 regularization
reg_alpha	[0,5]	Weight of L1 regularization
colsample_bytree	[0.5]	Random subset of features to be used in one regression tree

Note: Parameters not listed in the table take default values in LightGBM package.

$$\mathcal{P}(\phi) = \begin{cases} 1 & \text{if } \widehat{\pi} \geq \mu_0 \frac{\sigma_1^2}{\sigma_0^2} \\ 0 & \text{if } \widehat{\pi} < \mu_0 \frac{\sigma_1^2}{\sigma_0^2} \end{cases} \tag{A.1}$$

The policy in Equation (A.1) is intuitive. For example,

- If predictions are perfect, then $\sigma_1^2 = 0$ and the policy reverts to that of perfect predictions; launch those items for which $\widehat{\pi} \geq 0$
- If the model has no predictive ability, then $\sigma_1 \rightarrow \infty$ and the policy reverts to the prior mean, μ_0; launch all items if and only if the prior mean is positive.
- If there is no uncertainty in the prior, then $\sigma_1 \rightarrow 0$ and the policy again reverts to the prior mean; launch all items if and only if the prior mean is positive.
- For finite values of σ_1 and σ_0, the ratio, σ_1^2/σ_0^2, modifies the amount by which the

predicted profits must exceed prior beliefs in order to launch.

Proof of the threshold policy: The firm solves the following optimization problem:

$$\max_{\mathcal{P}(\phi)\in[0,1]} \mathbb{E}[\mathcal{P}(\phi)\cdot\pi + (1-\mathcal{P})(\phi)\cdot 0] = \max_{\mathcal{P}(\phi)(\phi)\in[0,1]} \mathbb{E}[\mathcal{P}(\phi)\cdot\pi] \tag{A.2}$$

where $\phi = (\widehat{\pi}, \pi, \sigma_1, \sigma_0)$ is the set of all known parameters; $\widehat{\pi}|\pi \sim \mathcal{N}(\pi, \sigma_1^2)$ and $\pi \sim \mathcal{N}(\mu_0, \sigma_0^2)$

Using the law of iterative expectations, we rewrite the initial maximization problem Equation (A.2) as:

$$\max_{\mathcal{P}(\phi)(\phi)\in[0,1]} \mathbb{E}[\mathcal{P}(\phi)\cdot\pi] = \max_{\mathcal{P}(\phi)(\phi)\in[0,1]} \mathbb{E}[\mathcal{P}(\phi)\cdot\mathbb{E}[\pi|\phi]] = \max_{\mathcal{P}(\phi)(\phi)\in[0,1]} \mathbb{E}[\mathcal{P}(\phi)\cdot\mathbb{E}[\pi|\widehat{\pi}]] \tag{A.3}$$

The last step relies on the assumption that $(\pi, \sigma_1, \sigma_0)$ are observable.

Because $\mathbb{E}[\pi|\phi]$ is a function of observables, ϕ, we can denote $\mathbb{E}[\pi|\phi] = f(\phi)$. Equation (A.3) is rewritten as:

$$\max_{\mathcal{P}(\phi)(\phi)\in[0,1]} \mathbb{E}[\mathcal{P}(\phi)f(\phi))] \tag{A.4}$$

Equation (A.4) implies that the optimal policy $\mathcal{P}^{\star}(\phi)$ has the following form ($\mathcal{I}(\cdot)$ is an indicator function):

$$\mathcal{P}^{\star}(\phi) = \mathcal{I}(f(\phi) \geq 0) = \mathbb{E}[\pi|\phi] \geq 0) \tag{A.5}$$

We show in the following that, for the case of normal priors, this policy would have a threshold form[1].

Because $\widehat{\pi}$ is normally distributed conditionally on π and since the prior is also normally

[1]Note that the optimal policy in Equation (A.5) does not depend on the normality assumption profitability; the policy is easily generalized to other distributions.

distributed, the posterior is normally distributed. Using standard formulae, we write:

$$\pi|\widehat{\pi} \sim \mathcal{N}\left(\frac{\widehat{\pi}\sigma_0^2 + \mu_0\sigma_1^2}{\sigma_0^2 + \sigma_1^2}, \frac{\sigma_0^2\sigma_1^2}{\sigma_0^2 + \sigma_1^2}\right) \text{ and } \widehat{\pi} \sim \mathcal{N}(\mu_0, \sigma_0^2 + \sigma_1^2) \tag{A.6}$$

From Equation (A.6), it follows that:

$$\mathbb{E}[\pi|\phi] = \frac{\widehat{\pi}\sigma_0^2 + \mu_0\sigma_1^2}{\sigma_0^2 + \sigma_1^2} \Rightarrow \mathcal{P}^\star(\phi) = \mathcal{I}\left(\widehat{\pi} \geq -\mu_0 \cdot \frac{\sigma_1^2}{\sigma_0^2}\right) \tag{A.7}$$

which is the threshold policy.

Result 2. Under the assumptions of Result 1, the optimal expected profit is:

$$\Pi^\star = (1 - \Phi(-\mu_0/\sigma_\nu)) \cdot \mu_0 + \sigma_\nu \cdot \varphi(-\mu_0/\sigma_\nu) \tag{A.8}$$

where $\Phi(\cdot)$ and $\varphi(\cdot)$ are the standard normal CDF and PDF respectively, and $\sigma_\nu = \sigma_0^2/\sqrt{\sigma_0^2 + \sigma_1^2}$

Proof: By substituting the optimal policy from Equation (A.7) and conditional expectation from Equation (A.7) to Equation (A.2), we rewrite the expected optimal profit as:

$$\begin{aligned}\Pi^\star &= \mathbb{E}\left[\mathcal{I}\left(\widehat{\pi} \geq -\mu_0 \cdot \frac{\sigma_1^2}{\sigma_0^2}\right) \cdot \frac{\widehat{\pi}\sigma_0^2 + \mu_0\sigma_1^2}{\sigma_0^2 + \sigma_1^2}\right] \\ &= \mathbb{E}[\mathcal{I}(\nu \geq 0)\nu] = \mathbb{P}[\nu \geq 0] \cdot \mathbb{E}[\nu|\nu \geq 0]\end{aligned} \tag{A.9}$$

where $\nu = \frac{\widehat{\pi}\sigma_0^2+\mu_0\sigma_1^2}{\sigma_0^2+\sigma_1^2} \sim \mathcal{N}\left(\frac{\widehat{\pi}\sigma_0^2+\mu_0\sigma_1^2}{\sigma_0^2+\sigma_1^2}; \frac{\sigma_0^4}{(\sigma_0^2+\sigma_1^2)^2}(\sigma_0^2 + \sigma_1^2)\right) \mathcal{N}(\mu_0, \sigma_0^4/(\sigma_0^2 + \sigma_1^2)) \sim \mathcal{N}(\mu_0, \sigma_\nu^2)$

Because ν is normally distributed, Equation (A.9) can be rewritten using the formula for the expectation of the truncated normal distribution:

$$\Pi^\star = (1 - \Phi(-\mu_0/\sigma_\nu)) \cdot \mu_0 + \sigma_\nu \cdot \varphi(-\mu_0/\sigma_\nu) \tag{A.10}$$

Result 3. The expected profit under the optimal policy is a decreasing function of σ_1^2.

Proof: Taking the derivative of Equation (A.10) with respect to σ_1^2:

$$\begin{aligned}
&- \mu_0 \cdot \phi(-\mu_0/\sigma_\nu) \left(-\frac{\mu_0}{2\sigma_0^2(\sigma_0^2+\sigma_1^2)^{1/2}}\right) - \frac{\sigma_0^2}{2(\sigma_0^2+\sigma_1^2)^{3/2}} \cdot \phi(-\mu_0/\sigma_\nu) \\
&+ \frac{\sigma_0^2}{(\sigma_0^2+\sigma_1^2)^{1/2}} \phi'(-\mu_0/\sigma_\nu) \left(-\frac{\mu_0}{2\sigma_0^2(\sigma_0^2+\sigma_1^2)^{1/2}}\right) \\
&= \left(\frac{\mu_0^2}{2\sigma_0^2(\sigma_0^2+\sigma_1^2)^{1/2}} - \frac{\sigma_0^2}{2(\sigma_0^2+\sigma_1^2)^{3/2}} + \left(\frac{\mu_0(\sigma_0^2+\sigma_1^2)^{1/2}}{\sigma_0^2}\right)\frac{-\mu_0}{2(\sigma_0^2+\sigma_1^2)}\right) \\
&\times \phi\left(-\frac{\mu_0}{\sigma_\nu}\right) \\
&= -\frac{\sigma_0^2}{2(\sigma_0^2+\sigma_1^2)^{3/2}} \phi\left(-\frac{\mu_0}{\sigma_\nu}\right)
\end{aligned} \tag{A.11}$$

Because $\phi(\cdot) > 0$ and $-\frac{\sigma_0^2}{2(\sigma_0^2+\sigma_1^2)^{3/2}}$ the expected profitability is decreasing function of σ_1^2 and therefore an increasing function of model accuracy.

A.3 Supporting Tables and Figures

Table A.2. Improvement in Predictive Accuracy Varying Minimum Threshold on Online Sales

Model	**Non-Image features**	**Image features**	$\mathbf{R^2}$	**Improvement over baseline**
CNN Features with 10 as threshold for online sales	Category, seasonality, price, color labels	Deep learning	43.14 (0.20)	+12.75%
CNN Features with 20 as threshold for online sales	Category, seasonality, price, color labels	Deep learning	46.88 (0.19)	+13.48%
CNN Features with 30 as threshold for online sales	Category, seasonality, price, color labels	Deep learning	51.23 (0.17)	+12.08%

Note: Improvements are calculated for baseline models estimated on the corresponding samples. Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.3. Tests of Uniqueness, Precision (variance of $N_i^{purchase}$), and Distance from Prior Collections

Model	**Non-Image features**	**Image features**	$\mathbf{R^2}$	**Improvement over baseline**
CNN Features	Category, seasonality, price, color labels	Deep learning	46.88 (0.19)	+0.00%
CNN Features (including uniqueness)	Category, seasonality, price, color labels image uniqueness	Deep learning	46.82 (0.23)	-0.13%
CNN Features (including vs. last year)	Category, seasonality, price, color labels image distance	Deep learning	44.55 (0.39)	-0.60%
CNN Features (precision weighting))	Category, seasonality, price, color labels variance weighting	Deep learning	46.77 (0.26)	-0.23%

Note: The sample of items included when estimating the last-year model exclude products sold only in the first year of the data. A GBRT/CNN model for the same items yields 44.85 (0.33). The –2.8% is relative to this model. Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.4. Improvement in Predictive Accuracy Using Alternative Prediction Models

Model	**Non-Image features**	**Image features**	R^2	**Improvement over baseline**
GBRT (CNN Features)	Category, seasonality, price, color labels	Deep-learning	46.88 (0.19)	+13.48%
Bagging Methods (CNN Features)	Category, seasonality, price, color labels	Deep-learning	45.35 (0.14)	+9.78%
LASSO (CNN Features)	Category, seasonality, price, color labels	Deep-learning	44.11 (0.32)	+6.78%

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.5. Predictions for the two Largest Categories (Dresses and Shirts)

Model	**Non-Image features**	**Image features**	R^2	**Improvement over baseline**
Non-Image Baseline	Category, seasonality, price	None	57.94 (0.27)	–
Color-labels	Category, seasonality, price, color labels	None	59.79 (0.24)	+3.19%
Automated Color Features	Category, seasonality, price, color labels	RGB	60.86 (0.22)	+5.04%
Automated Color and Patterns	Category, seasonality, price, color labels	RGB+ Gabor	61.91 (0.31)	+6.85%
Human-coded features	Category, seasonality, price, color labels	Human Coded	61.91 (0.27)	+6.85%
CNN Features	Category, seasonality, price, color labels	Deep learning	63.69 (0.19)	+9.92%

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.6. Improvement in Predictive Accuracy Using an Alternative CNN

Model	**Non-Image features**	**Image features**	R^2	**Improvement over baseline**
ResNet CNN (this paper)	Category, seasonality, price, color labels	Deep learning	46.88 (0.19)	+13.48%
VGG-19 CNN (this paper)	Category, seasonality, price, color labels	Deep learning	46.84 (0.18)	+13.39%

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.7. Improvement in Predictive Accuracy Using PCA (nonlinear and linear tested; linear shown)

Model	**Non-Image features**	**Image features**	$\mathbf{R^2}$	**Improvement over baseline**
Color Features	Category, seasonality, price, color labels	RGB	43.48 (0.18)	+5.25%
Color and Patterns	Category, seasonality, price, color labels	Gabor	41.37 (0.33)	+0.15%
CNN Features	Category, seasonality, price, color labels	Deep learning	46.55 (0.21)	+12.68%

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.8. Predictions Using Automated Pattern & Color Image-Processing Features

Model	**Non-Image features**	**Image features**	$\mathbf{R^2}$	**Improvement over baseline**
Non-Image Baseline	Category, seasonality, price	None	41.31 (0.18)	–
Color Features	Category, seasonality, price, color labels	RGB	44.06 (0.20)	+6.66%
Pattern Features	Category, seasonality, price, color labels	Gabor	44.34 (0.23)	+7.33%
Color and Patterns	Category, seasonality, price, color labels	RGB + Gabor	45.28 (0.18)	+9.61%
CNN Features	Category, seasonality, price, color labels	Deep learning	46.88 (0.19)	+13.48%
CNN Features (all images)	Category, seasonality, price, color labels	Deep learning	47.48 (0.22)	+14.93%

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.9. Improvement in Predictive Accuracy Using Alternative Image-Feature Extraction Methods

Model	**Non-Image features**	**Image features**	**R^2**	**Improvement over baseline**
RGB Features	Category, seasonality, price, color labels	RGB	44.04 (0.20)	+6.61%
HSV Features	Category, seasonality, price, color labels	HSV	44.30 (0.14)	+7.24%
ORB Features	Category, seasonality, price, color labels	ORB	43.53 (0.25)	+5.37%

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Figure A-1. Return Rate by Postlaunch Time of Purchase, Day of the Week and Month

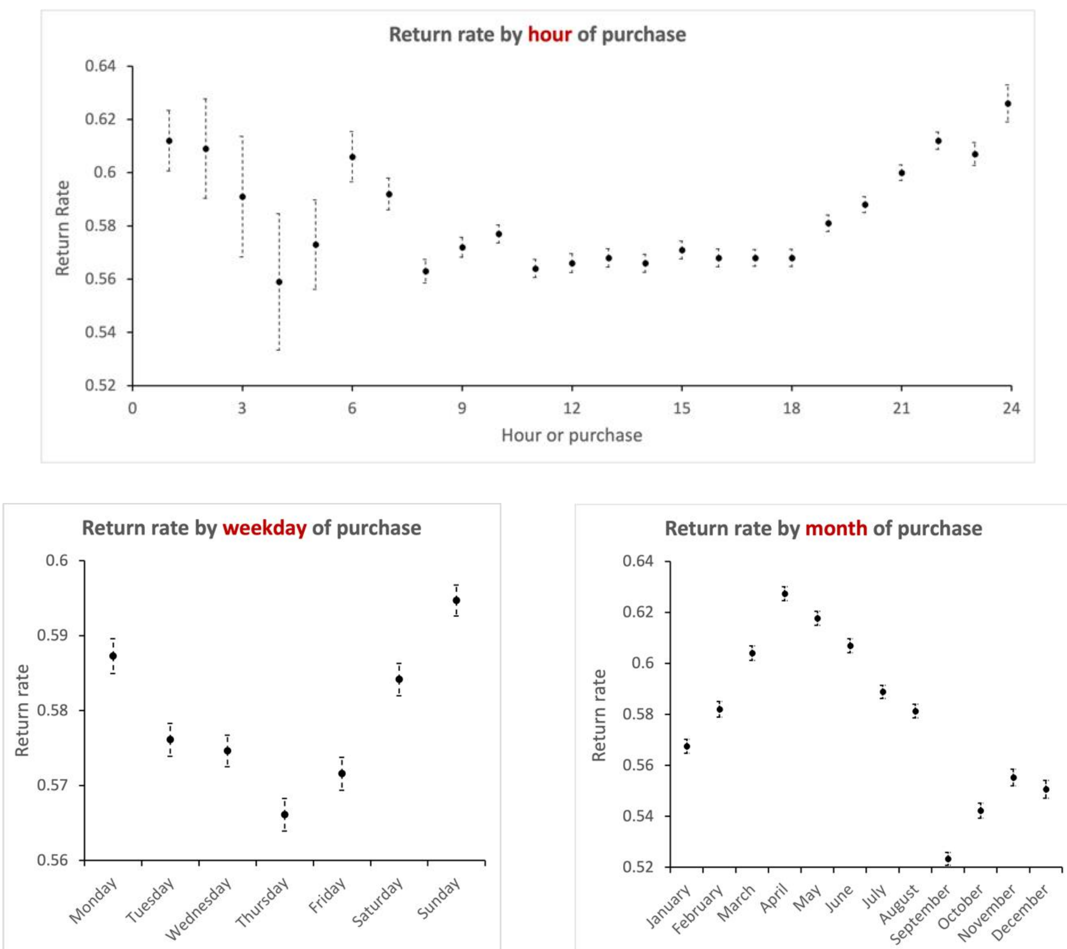

Table A.10. Online Sales and Return rates, Offline Sales and Return Rates, and Model-Predicted Online Return Rates by Product Category (based on all sales). Models estimated for color-label categories with at least 400 items with ≥20 sales.

Category	On. Sal.	On. Ret.	On. Ret. Rate	Off. Sal.	Off. Ret.	Off. Ret. Rate	# ≥ 20 Sales	R^2 (Base)	R^2 (Main)
Dresses	96,754	69,626	71.96%	45,923	1,615	3.52%	759	28.11 (0.65)	31.13 (0.83)
Shirts	80,586	39,379	48.87%	299,313	7,007	2.34%	1,213	14.78 (0.59)	24.08 (0.77)
Blouses	43,413	23,292	53.65%	104,778	2,667	2.55%	687	4.33 (0.96)	15.78 (1.00)
Pants	36,183	21,209	58.62%	103,353	3,264	3.16%	496	-1.10 (1.16)	-0.17 (1.25)
Knit	31,893	15,708	49.25%	137,227	3,889	2.83%	511	1.80 (1.24)	17.35 (1.13)
Jackets	21,304	12,228	57.40%	24,385	876	3.59%	302	– –	– –
Blazer	13,190	7,627	57.82%	27,748	993	3.58%	166	– –	– –
Cardigans	11,315	4,167	36.83%	16,462	507	3.08%	69	– –	– –
Skirts	9,252	5,259	56.84%	26,884	746	2.77%	135	– –	– –
Coats	5,170	3,238	62.63%	1,299	49	3.77%	88	– –	– –
Bolero	4,867	3,367	69.18%	0	0	0.00%	41	– –	– –
Sweatshirts	3,862	2,170	56.19%	6,126	191	3.12%	56	– –	– –
Jumpsuits	1,902	1,303	68.51%	462	10	2.16%	27	– –	– –
Top	1,543	728	47.18%	0	0	0.00%	27	– –	– –
Leather	614	287	46.74%	809	34	4.20%	8	– –	– –

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Table A.11. Online Sales and Return rates, Offline Sales and Return Rates, and Model-Predicted Online Return Rates by Color Labels. Models estimated for color-label categories with at least 400 items with ≥20 sales.

Category	**On. Sal.**	**On. Ret.**	**On. Ret. Rate**	**Off. Sal.**	**Off. Ret.**	**Off. Ret. Rate**	**# ≥ 20 Sales**	**R^2 (Base)**	**R^2 (Main)**
Blue	84,947	49,054	57.75%	217,457	5,658	2.60%	1,056	45.28 (0.28)	50.06 (0.47)
Grey	67,381	39,320	58.35%	92,119	2,787	3.02%	951	34.06 (0.58)	36.36 (0.54)
White	48,751	26,913	55.21%	118,060	3,154	2.67%	743	23.25 (0.6)	23.79 (0.87)
Red	42,708	25,749	60.29%	81,625	2,294	2.81%	542	45.80 (0.62)	50.90 (0.79)
Brown	41,540	23,557	56.71%	84,227	2,446	2.90%	590	19.09 (0.77)	24.72 (0.95)
Black	32,444	19,305	59.50%	67,663	2,028	3.00%	411	41.01 (0.53)	42.57 (0.83)
Green	10,550	6,581	62.38%	19,815	466	2.35%	144	– –	– –
Pink	3,539	2,118	59.85%	8,588	249	2.90%	53	– –	– –
Orange	3,054	1,865	61.07%	7,701	202	2.62%	53	– –	– –
Yellow	1,670	1,017	60.90%	2,356	64	2.72%	23	– –	– –
Violet	938	617	65.78%	2,273	66	2.90%	14	– –	– –
Several	318	151	47.48%	84,314	2,224	2.66%	5	– –	– –

Note: Standard deviations are reported in parentheses. Out-out-sample performance reported.

Figure A-2. Impact of Color Clusters on Return Rates

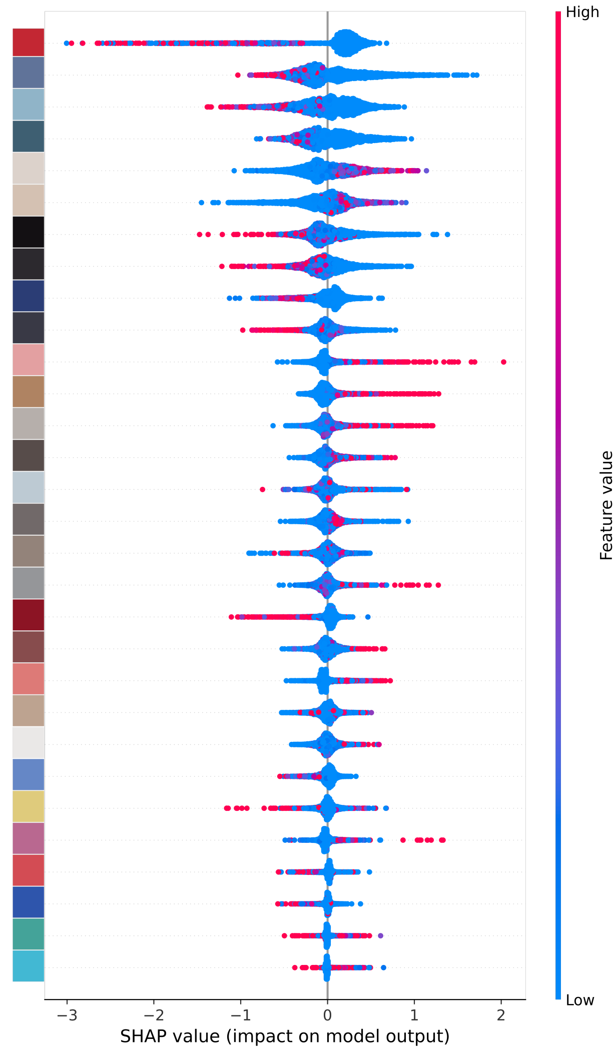

Figure A-3. The Impact of Non-Image Features

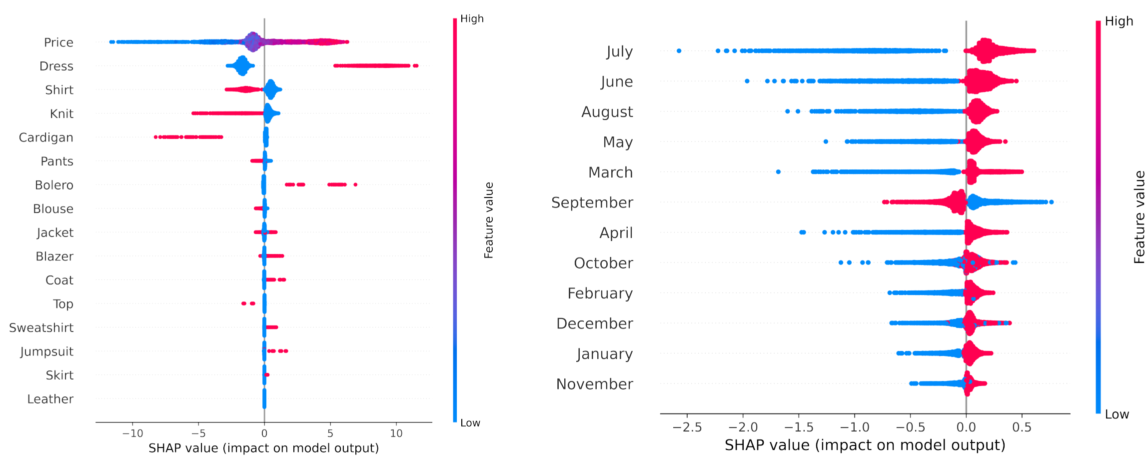

(a) Price and Category

(b) Seasonality

Table A.12. Product Return Rates and Price Discounts

	Model 1	**Model 2**	**Model 3**	**Model 4**
Proportion discounted	0.080***	0.086***	0.088***	0.078***
	(0.008)	(0.007)	(0.007)	(0.007)
Price (log10)	0.228***	0.221***	0.214***	0.279***
	(0.012)	(0.012)	(0.012)	(0.008)
Intercept	-0.156***	-0.113***	-0.083***	0.028***
	(0.044)	(0.041)	(0.041)	(0.015)
Category Controls	Yes	Yes	Yes	No
Color Controls	Yes	Yes	No	No
Seasonality Controls	Yes	No	No	No
Number observations	4585	4585	4585	4585
Adjusted R-squared	0.434	0.414	0.403	0.258

Note: Standard errors are heteroskedasticity robust ($\star p \leq 0.1, \star\star p \leq 0.05, \star\star\star p \leq 0.01$). R-squared is in-sample

Table A.13. Interpreting the Effect of Human-coded features (HCF) on Item Return Rates

	Regression Model	**SHAP Values**
Asymmetric	0.022***	0.92
	(0.008)	
Floral	-0.038***	-0.90
	(0.010)	
Striped	-0.063***	-0.95
	(0.009)	
Geometric/abstract	-0.020***	-0.89
	(0.007)	
Lace details	0.010	0.85
	(0.008)	
Metallic/sequin details	0.008	0.82
	(0.006)	
Graphic details	0.008	0.34
	(0.013)	
Text details	0.036*	0.69
	(0.021)	
Short Sleeves	-0.020***	-0.81
	(0.006)	
Medium Sleeves	-0.032***	-0.84
	(0.008)	
Long Sleeves	-0.032***	-0.87
	(0.007)	
Belt	0.025***	0.88
	(0.010)	
Zipper	-0.027**	-0.80
	(0.012)	
Intercept	0.039	–
	(0.026)	
Price (log10)	0.256***	–
	(0.014)	
Category Controls	Yes	Yes
Color Controls	Yes	Yes
Seasonality Controls	Yes	Yes
Number observations	1,972	1,972
Adjusted R-squared	0.628	–

Note: Standard errors are heteroskedasticity robust ($\star p \leq 0.1, \star\star p \leq 0.05, \star\star\star p \leq 0.01$). R-squared is in-sample

Appendix B

Appendix for Chapter 2

B.1 Data Preprocessing and Additional Summary Statistics

In the paper, I preprocessed the data to obtain better estimates of the model parameters. I took the following steps in the preprocessing:

1. Remove non-fashion products (e.g. linen, towels) and kid's apparel. These products constitute a small proportion of the data and are not the focus of the retailer.
2. Remove sessions without product page views. This could happen if the customer comes to the website from the third-party website and lands directly on the product page. These sessions do not represent the true search process and I are not able to recover the set of products from which the customer was choosing.
3. Remove sessions that do not have any clicked products after a page view and sessions which have clicked products before a page view. This implies that I keep only sessions

with the clean search process: the customer views the product page and selects the product to click. The alternative could happen if the customer found a product through an alternative means (from a third-party website) and in this case, it is impossible to infer the set of products from which he or she was choosing.

4. Remove 0.5% of longest/shortest sessions based on time spent. This includes accident clicks and long-tail outliers, for example, if a customer forgot to close the website.

B.2 Analysis with Deep Learning Embedding

Figure B-1. Breadth of Search and Product Returns

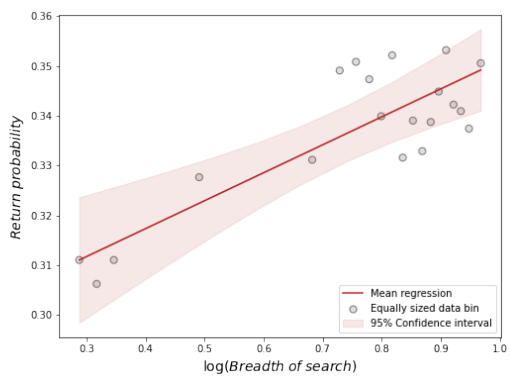

In Section 2.3.2 I provided a model free evidence that customer search and returns are related. In Figure B-1 I extend the analysis to demonstrate how variety of search and product returns are related. I implement the product embedding approach – by using the model from Chapter 3 I assign to each product a numerical vector representing this product. The property of these vectors is that similar product would have a similar embeddings, therefor by computing summary statistics of the embeddings I can quantify which products the customer was clicking. By taking the variance within the session I characterize the depth search search – low variance vs. high variance. The results are plotted in Figure B-1 and we can see that the customer with high variance are substantially more likely to return the product. Notably Figure B-1 captures similarities and differences between products on more granular level enhancing the results in Figure 2-2d and Figure 2-2e.

B.3 Derivation of Expected Purchase Utility

Without loss of generality, I drop all indices and subscripts in this section to preserve readability. Remember $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon), \eta \sim \mathcal{N}(0, \sigma_\eta)$ where ϵ and η are independent. Thus, ϵ and η have a joint normal distribution with diagonal variance-covariance matrix. From the properties of joint normal distribution, it follows that variable $\psi = \epsilon + \eta$ and ϵ also have a joint normal distribution. Hence, one can find the conditional distribution $\epsilon|\eta \sim \mathcal{N}(\frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\eta^2}\psi, \frac{\sigma_\epsilon^2\sigma_\eta^2}{\sigma_\epsilon^2+\sigma_\eta^2})$

In Section 2.4.1 I wanted to find the purchase expected utility as:

$$
\begin{aligned}
v &= \mathbb{E}_\epsilon\left[(\mu^u + \epsilon)\mathcal{I}(\mu^u + \epsilon \geq -R) + (-R) \cdot \mathcal{I}(\mu^u + \epsilon < -R)|\psi\right] = \\
&= \mathbb{E}_\epsilon\left[(\mu^u + \epsilon + R)\mathcal{I}(\mu^u + \epsilon + R \geq)|\psi\right] - R = \\
&= \mathbb{E}_\zeta[\zeta \cdot \mathcal{I}(\zeta \geq 0)|\psi] - R
\end{aligned}
\tag{B.1}
$$

where $\zeta = \mu^u + \epsilon + R$. Because $\epsilon|\psi$ is normally distributed then $\zeta|\psi$ is also normally distributed and thus I can compute the expectation above as:

$$
\begin{aligned}
v &= \frac{1}{\sqrt{2\pi}\sigma_\zeta}\int_0^\infty te^{-\frac{(t-\mu_\zeta)^2}{2\sigma_\zeta^2}}dt - R = \\
&= \frac{1}{\sqrt{2\pi}\sigma_\zeta}\int_0^\infty (t-\mu_\zeta)e^{-\frac{(t-\mu_\zeta)^2}{2\sigma_\zeta^2}}dt + \frac{\mu_\zeta}{\sqrt{2\pi}\sigma_\zeta}\int_0^\infty te^{-\frac{(t-\mu_\zeta)^2}{2\sigma_\zeta^2}}dt - R = \\
&= \sigma_\zeta\varphi\left(\frac{\mu_\zeta}{\sigma_\zeta}\right) + \mu_\zeta\Phi\left(\frac{\mu_\zeta}{\sigma_\zeta}\right) - R = \sigma_\zeta\left(\varphi\left(\frac{\mu_\zeta}{\sigma_\zeta}\right) + \frac{\mu_\zeta}{\sigma_\zeta}\Phi\left(\frac{\mu_\zeta}{\sigma_\zeta}\right)\right) - R \\
&= \sigma_\zeta T\left(\frac{\mu_\zeta}{\sigma_\zeta}\right) - R
\end{aligned}
\tag{B.2}
$$

where $\mu_\zeta = \mu^u + R + \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\eta^2}\psi$; $\sigma_\zeta = \sqrt{\frac{\sigma_\epsilon^2\sigma_\eta^2}{\sigma_\epsilon^2+\sigma_\eta^2}}$ and $T(x) = x\Phi(x) + \varphi(x)$

B.4 Derivation of Reservation Utilities for Model with Product Returns

In the original paper, Weitzman (1979) demonstrated that the reservation utility z for a product could be found from Equation (B.3) where I drop the individual (i) and product (j) indices for compactness

$$
c = \int_z^\infty (u-z)dF(u) \tag{B.3}
$$

I demonstrated in Section 2.4.1 that the return option changes the distribution of the reward and thus, in this case, I need to find the distribution of the expected purchase utility from Equation (B.2). Notice that the randomness comes from the signal of customer preference ψ while all other parameters are fixed and known to the customer.

$$
\begin{aligned}
F(u) &= \mathbb{P}[v(\psi) \leq u] = \mathbb{P}\left[\sigma_\zeta T(\mu_\zeta/\sigma_\zeta) - R \leq u\right] = \mathbb{P}\left[T(\mu_\zeta/\sigma_\zeta) \leq \frac{u+R}{\sigma_\zeta}\right] = \\
&= \mathbb{P}\left[\mu_\zeta/\sigma_\zeta \leq T^{-1}\left(\frac{u+R}{\sigma_\zeta}\right)\right] = \mathbb{P}\left[\mu_\zeta \leq \sigma_\zeta T^{-1}\left(\frac{u+R}{\sigma_\zeta}\right)\right] = \\
&= \mathbb{P}\left[\mu^u + R + \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\eta^2}\psi \leq \sigma_\zeta T^{-1}\left(\frac{u+R}{\sigma_\zeta}\right)\right] = \\
&= \mathbb{P}\left[\frac{\psi}{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}} \leq \frac{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}}{\sigma_\epsilon^2}\left(\sigma_\zeta T^{-1}\left(\frac{u+R}{\sigma_\zeta}\right) - \mu^u - R\right)\right] = \\
&= \Phi\left[\frac{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}}{\sigma_\epsilon^2}\left(\sigma_\zeta T^{-1}\left(\frac{u+R}{\sigma_\zeta}\right) - \mu^u - R\right)\right]
\end{aligned} \tag{B.4}
$$

Next, I plug-in the distribution from Equation (B.4) into Equation (B.3) and obtain:

$$
\begin{aligned}
c &= \int_{\frac{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}}{\sigma_\epsilon^2}\left(\sigma_\zeta T^{-1}\left(\frac{u+R}{\sigma_\zeta}\right)-\mu^u-R\right)}^{\infty} \sigma_\zeta T\left(\frac{\mu^u + R + \frac{\sigma_\epsilon^2}{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}}t}{\sigma_\zeta}\right) - R - z \, d\Phi(t) = \\
&= \sigma_\zeta \int_\theta^\infty T\left(\frac{\mu^u + R + \frac{\sigma_\epsilon^2}{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}}t}{\sigma_\zeta}\right) - T\left(\frac{\mu^u + R + \frac{\sigma_\epsilon^2}{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}}\theta}{\sigma_\zeta}\right) d\Phi(t)
\end{aligned} \tag{B.5}
$$

where I used the substitution $z = T\left(\frac{\mu^u+R+\frac{\sigma_\epsilon^2}{\sqrt{\sigma_\epsilon^2+\sigma_\eta^2}}\theta}{\sigma_\zeta}\right)$

B.5 Approximating the Solution to the Equation

In the paper, I made an identifying assumption that $\sigma_\epsilon = 1$. Thus, from Equation (B.5) it could be seen that the reservation utility is a function of three parameters: $z\star = f(\mu^u +$

$R, \sigma_\eta, c) = f(x_1, x_2, x_3)$. During the optimization algorithm – finding this function for each customer-product combination is not feasible as it involves many integration steps.

To circumvent the computational burden, I used trilinear interpolation technique. Specifically, for three-dimensional variable (x_1, x_2, x_3), I constructed a grid of values and computed the exact reservation utilities for each element of the grid. Notice that in this case, the space of possible values of (x_1, x_2, x_3) is divided into 3-dimensional cubes. For each of these cubes, I know the exact values of reservation utilities in eight vertices. For any vector within the cube, I approximate the reservation utility function $f(x_1, x_2, x_3)$ as:

$$\begin{aligned} & f_{true}(x_1, x_2, x_3) \simeq f_{approx}(x_1, x_2, x_3) \\ =& \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_3 + \alpha_4 x_1 x_2 + \alpha_5 x_2 x_3 + \alpha_6 x_1 x_3 + \alpha_7 x_1 x_2 x_3 \end{aligned} \tag{B.6}$$

where I require $f_{approx}(x_1, x_2, x_3) = f_{true}(x_1, x_2, x_3)$ at the grid (or cube vertices) points. Because $f_{approx}(x_1, x_2, x_3)$ has eight parameters and eight constraints, the linear system has a unique solution for each cell.

B.6 Derivation of Equivalent Set of Constraints on Model Parameters

After combining Equations (2.9) to (2.10), I can compute the variable W_i from Equation (2.16). For compactness and without loss of generality, I drop customer index i:

$$W_i = \left[\prod_{j=0}^{C-1} \mathcal{I}(\omega_{j+1} \geq \max_{s=j+2..S} \omega_s) \cdot \mathcal{I}(\max_{s=0..j} v_s \leq \max_{s=j+1..S} \omega_s)\right]$$
$$\times \mathcal{I}(\omega_C \geq \max_{s=C+1..S} \omega_s) \cdot \mathcal{I}(\max_{s=0..C} v_s \geq \max_{s=C+1..S} \omega_s) \tag{B.7}$$
$$\times \mathcal{I}(v_b \geq \max_{s=0..C} v_s) \cdot \mathcal{I}(x_b'\beta^u + \epsilon_b \geq R)$$

Consider the first part of the equation:

$$\begin{aligned} P_1 &= \left[\prod_{j=0}^{C-1} \mathcal{I}(\omega_{j+1} \geq \max_{s=j+2..S} \omega_s)\right] \cdot \mathcal{I}(\omega_C \geq \max_{s=C+1..S} \omega_s) \\ &= \left[\prod_{j=0}^{C-1} \left[\prod_{s=j+1}^{S} \mathcal{I}(\omega_{j+1} \geq \omega_s)\right]\right] \cdot \mathcal{I}(\omega_C \geq \max_{s=C+1..S} \omega_s) \\ &= \left[\prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1})\right] \cdot \mathcal{I}(\omega_C \geq \max_{j=C+1..S} \omega_j) \end{aligned} \tag{B.8}$$

Notice that Equation (B.7) is a necessary condition for $W = 1$. Thus, I can assume that these inequalities hold in further derivations. Specifically, it follows that $\max_{s=j+1..S} \omega_s = \omega_{j+1}$ and I can rewrite the part of the equation as:

$$\begin{aligned} P_2 &= \prod_{j=0}^{C-1} \mathcal{I}(\max_{s=0..j} v_s \leq \max_{s=j+1..S} \omega_s) = \prod_{j=0}^{C-1} \mathcal{I}(\max_{s=0..j} v_s \leq \omega_{j+1}) \\ &= \prod_{j=0}^{C-1} \left[\prod_{s=0}^{j} \mathcal{I}(v_s \leq \omega_{j+1})\right] = \prod_{j=0}^{C-1} \mathcal{I}(v_j \leq \omega_C) \end{aligned} \tag{B.9}$$

Similarly, I find that:

$$
\begin{aligned}
P_3 &= \mathcal{I}(\max_{s=0..C} v_s \geq \max_{s=C+1..S} \omega_s) \cdot \mathcal{I}(v_b \geq \max_{s=0..C} v_s) \\
&= \mathcal{I}(v_b \geq \max_{s=C+1..S} \omega_s) \cdot \prod_{j=0}^{C} \mathcal{I}(v_j \leq v_b)
\end{aligned}
\tag{B.10}
$$

Finally, combining all Equations (B.8) to (B.10) and adding the returns inequality, I find the equivalent simplified form of the inequality constraints:

$$
\begin{aligned}
W = &\left[\prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1})\right] \cdot \mathcal{I}(\omega_C \geq \max_{j=C+1..S} \omega_j) \\
&\times \left[\prod_{j=0}^{C-1} \mathcal{I}(v_j \leq \min\{\omega_C, v_b\})\right] \cdot \mathcal{I}(v_C \leq v_b) \cdot \mathcal{I}(v_b \geq \max_{j=C+1..S} \omega_j) \\
&\times \mathcal{I}(x_b' \beta^u + \epsilon_b \geq -R)
\end{aligned}
\tag{B.11}
$$

B.7 Derivation of Semi-Closed Form Likelihood

As in the previous sections, I drop the customer related index i for compactness. Recall the set of constraints in simplified form from Equation (2.13)

$$
\begin{aligned}
W = &\left[\prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1})\right] \cdot \mathcal{I}(\omega_C \geq \max_{j=C+1..S} \omega_j) \\
&\times \left[\prod_{j=0}^{C-1} \mathcal{I}(v_j \leq \min\{\omega_C, v_b\})\right] \cdot \mathcal{I}(v_C \leq v_b) \cdot \mathcal{I}(v_b \geq \max_{j=C+1..S} \omega_j) \\
&\times \mathcal{I}(x_b' \beta^u + \epsilon_b \geq -R)
\end{aligned}
\tag{B.12}
$$

where v_j is a function of unobservable shock ψ_j; ω_j is a function of unobserved shock ξ_j and unobserved shock ϵ_b. To compute the likelihood, I need to integrate out all these unobserved shocks. In Equation (B.13) I use the fact that all ψ_j and ξ_j are independent by assumption, while ϵ_b and ψ_b are dependent.

$$\begin{aligned}
&\int \cdots \int W dF \\
&= \int \cdots \int W \prod_{j=1, j\neq b}^{C} dF_{\psi_j}(\psi_j) \prod_{j=1}^{S} dF_{\xi_j}(\xi_j) dF_{\psi_b,\epsilon_b}(\psi_b, \epsilon_b) \\
&= \int \cdots \int W \prod_{j=1, j\neq b}^{C} dF_{\psi_j}(\psi_j) \prod_{j=1}^{S} dF_{\xi_j}(\xi_j) dF_{\epsilon_b|\psi_b}(\epsilon_b|\psi_b) dF_{\psi_b}(\psi_b)
\end{aligned} \tag{B.13}$$

The distribution $F_{\epsilon_b|\psi_b}(\epsilon_b|\psi_b)$ is known from Appendix B.4 and thus I can integrate out the variable ϵ_b as only one constraint depends on it.

$$\begin{aligned}
&\int \mathcal{I}(x_b'\beta^u + \epsilon_b \geq -R) dF_{\epsilon_b|\psi_b}(\epsilon_b|\psi_b) = \\
&1 - \Phi\left[-\frac{R + x_b'\beta^u + \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\eta^2}\psi_b}{\sqrt{\frac{\sigma_\epsilon^2\sigma_\eta^2}{\sigma_\epsilon^2+\sigma_\eta^2}}}\right] = RP_b(\psi_b)
\end{aligned} \tag{B.14}$$

Next, I notice that $\xi_j : j = C+1..S$ appear only in two constraints which could be simplified to:

$$
\begin{aligned}
&\int \cdots \int \mathcal{I}(\omega_C \geq \max_{j=C+1..S} \omega_j)\mathcal{I}(v_C \leq v_b) \cdot \mathcal{I}(v_b \geq \max_{j=C+1..S} \omega_j) \prod_{j=C+1}^{S} dF_{\xi_j}(\xi_j) \\
&= \int \cdots \int \mathcal{I}(\min\{v_b, \omega_C\} \geq \max_{j=C+1..S} \omega_j) \prod_{j=C+1}^{S} dF_{\xi_j}(\xi_j) \\
&= \int \cdots \int \prod_{j=C+1}^{S} \mathcal{I}(\min\{v_b, \omega_C\} \geq \omega_j) \prod_{j=C+1}^{S} dF_{\xi_j}(\xi_j) \\
&= \prod_{j=C+1}^{S} \int \mathcal{I}(\min\{v_b, \omega_C\} \geq \omega_j) dF_{\xi_j}(\xi_j)
\end{aligned}
\tag{B.15}
$$

Notice that $\omega_j(\xi_j)$ is an inversible function for each j. (Equation (2.5) implies that this function would be different depending on product characteristics and costs). Thus, I can compute:

$$
\begin{aligned}
&\prod_{j=C+1}^{S} \int \mathcal{I}(\min\{v_b, \omega_C\} \geq \omega_j(\xi_j)) dF_{\xi_j}(\xi_j) \\
&= \prod_{j=C+1}^{S} \int \mathcal{I}\left(\omega_j^{-1}(\min\{v_b, \omega_C\}) \leq \xi_j\right) dF_{\xi_j}(\xi_j) \\
&= \prod_{j=C+1}^{S} \left[1 - F_{\xi_j}\left(\omega_j^{-1}(\min\{v_b, \omega_C\})\right)\right]
\end{aligned}
\tag{B.16}
$$

Next, I modify constraints related to purchase decision. However, in this case I consider three separate cases: choosing the outside option, choosing the last searched option, all else:

$$
\left[\prod_{j=0}^{C-1}\mathcal{I}(v_j\le\min\{\omega_C,v_b\})\right]\cdot\mathcal{I}(v_C\le v_b)=
$$

$$
\begin{cases}
\left[\prod_{j=1}^{C-1}\mathcal{I}(v_j\le\min\{\omega_C,v_0\})\right]\cdot\mathcal{I}(v_C\le v_0)\cdot\mathcal{I}(v_0\le\omega_C) & b=0\\
\left[\prod_{j=1}^{C-1}\mathcal{I}(v_j\le\min\{\omega_C,v_C\})\right]\cdot\mathcal{I}(v_0\le\min\{\omega_C,v_C\}) & b=C\\
\left[\prod_{j=1,j\neq b}^{C-1}\mathcal{I}(v_j\le\min\{\omega_C,v_b\})\right]\mathcal{I}(v_C\le v_b\le\omega_C,v_0\le\min\{\omega_C,v_b\}) & \text{other}
\end{cases}
\tag{B.17}
$$

Notice that all other inequalities except those in Equation (B.17) do not depend on the unobserved shocks $\{\psi_j : j = 1..C, j \neq b\}$ and thus could be integrated out. Because $v_j(\psi_j)$ is an inversible function for each j (Equation (2.3) implies that this function would be different depending on product characteristics and costs). Consider integration of different cases from Equation (B.17) and remember that in case the customer chose the non-outside option, I should keep the result from Equation (B.14) as return probability $RP_b(\psi_b)$ depends on ψ_b.

Choice of outside option or $b = 0$

$$
\begin{aligned}
&\int\cdots\int\left[\prod_{j=1}^{C-1}\mathcal{I}(v_j\le\min\{\omega_C,v_0\})\right]\cdot\mathcal{I}(v_C\le v_0)\cdot\mathcal{I}(v_0\le\omega_C)\prod_{j=1}^{C}dF_{\psi_j}(\psi_j)\\
&=\int\cdots\int\left[\prod_{j=1}^{C-1}\mathcal{I}\left(\psi_j\le v_j^{-1}(\min\{\omega_C,v_0\})\right)\right]\cdot\mathcal{I}(\psi_C\le v_C^{-1}(v_0))\cdot\mathcal{I}(v_0\le\omega_C)\\
&\times\prod_{j=1}^{C}dF_{\psi_j}(\psi_j)=\mathcal{I}(v_0\le\omega_C)\left[\prod_{j=1}^{C-1}F_{\psi_j}\left(v_j^{-1}(\min\{\omega_C,v_0\})\right)\right]F_{\psi_C}\left(v_C^{-1}(v_0)\right)
\end{aligned}
\tag{B.18}
$$

Choice of last clicked option or $b = C$

$$\begin{aligned}
&\int \cdots \int RP_C(\psi_C) \left[\prod_{j=1}^{C-1} \mathcal{I}(v_j \leq \min\{\omega_C, v_C\})\right] \mathcal{I}(v_0 \leq \min\{\omega_C, v_C\}) \\
&\times \prod_{j=1}^{C} dF_{\psi_j}(\psi_j) = \int \cdots \int RP_C(\psi_C) \left[\prod_{j=1}^{C-1} \mathcal{I}\left(\psi_j \leq v_j^{-1}(\min\{\omega_C, v_C\})\right)\right] \\
&\times \mathcal{I}(v_0 \leq v_C)\mathcal{I}(v_0 \leq \omega_C) \prod_{j=1}^{C} dF_{\psi_j}(\psi_j) = \mathcal{I}(v_0 \leq \omega_C) \\
&\times \int RP_C(\psi_C) \left[\prod_{j=1}^{C-1} F_{\psi_j}\left(v_j^{-1}(\min\{\omega_C, v_C\})\right)\right] \mathcal{I}(v_0 \leq v_C) dF_{\psi_C}(\psi_C) \\
&= \mathcal{I}(v_0 \leq \omega_C) \int_{v_C^{-1}(v_0)}^{\infty} RP_C(\psi_C) \prod_{j=1}^{C-1} F_{\psi_j}\left(v_j^{-1}(\min\{\omega_C, v_C\})\right) dF_{\psi_C}(\psi_C)
\end{aligned} \tag{B.19}$$

Choice of other option or $0 < b < C$

$$\begin{aligned}
&\int \cdots \int RP_B(\psi_b) \left[\prod_{j=1, j\neq b}^{C-1} \mathcal{I}(v_j \leq \min\{\omega_C, v_b\})\right] \mathcal{I}(v_C \leq v_b \leq \omega_C) \\
&\times \mathcal{I}(v_0 \leq \min\{\omega_C, v_b\}) \prod_{j=1}^{C} dF_{\psi_j}(\psi_j) = \int \cdots \int RP_B(\psi_b) \\
&\times \left[\prod_{j=1, j\neq b}^{C-1} \mathcal{I}(v_j \leq \min\{\omega_C, v_b\})\right] \mathcal{I}(\psi_C \leq v_C^{-1}(v_b))(v_0 \leq v_b \leq \omega_C) \\
&\times \prod_{j=1}^{C} dF_{\psi_j}(\psi_j) = \mathcal{I}(v_0 \leq \omega_C) \\
&= \int_{v_b^{-1}(v_0)}^{v_b^{-1}(\omega_C)} RP_b(\psi_b) \left[\prod_{j=1, j\neq b}^{C-1} F_{\psi_j}\left(v_j^{-1}(\min\{\omega_C, v_b\})\right)\right] F_{\psi_C}(v_C^{-1}(v_0)) dF_{\psi_b}(\psi_b)
\end{aligned} \tag{B.20}$$

Notice that after combining Equations (B.18) to (B.20) and recalling Equation (B.14) and Equation (B.16), I can rewrite the original integral in Equation (B.13) as:

$$\int \cdots \int W dF(\epsilon, \psi, \xi)$$
$$\int \cdots \int \mathcal{I}(v_0 \leq \omega_C) \prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1}) \int_{\underline{\psi_b}}^{\overline{\psi_b}} B(\xi_C, \psi_b) dF_{\psi_b}(\psi_b) \prod_{j=1}^{C} dF_{\xi_j}(\xi_j) \tag{B.21}$$

where $B(\cdot, \cdot)$ is a function which depends only on two unobserved shocks ξ_C through ω_C and ψ_b through v_b.

Next, notice that only $\omega_C(\xi_C)$ depends on ξ_C, therefore, Equation (B.21) could be rewritten as:

$$\int \cdots \int W dF(\epsilon, \psi, \xi)$$
$$= \int_{-\infty}^{+\infty} \int_{\underline{\psi_b}}^{\overline{\psi_b}} \mathcal{I}(v_0 \leq \omega_C) B(\xi_C, \psi_b) D(\xi_C) dF_{\psi_b}(\psi_b) dF_{\xi_C}(\xi_C) \tag{B.22}$$
$$D(\xi_C) = \int \cdots \int \prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1}) \prod_{j=1}^{C-1} dF_{\xi_j}(\xi_j)$$

Notice that in Equation (B.22), I need to simulate only $C+1$ random shocks in comparison with $S + C + 1$ in Equation (B.12). Because typically $S >> C$ (number of viewed products is much higher than number of clicks), this is already a large improvement. However, $\prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1})$ could have quite many sharp constraints in longer sessions, which may result in a higher chance of having zero-valued integral approximation in Equation (B.22). Also, as discussed in the paper, the reservation utility $\omega_j \to \infty$ if search costs $c_j \to \overline{c_j}$, where $\overline{c_j}$ is an upper bound on costs and could be found from Equation (B.5) by making $\theta \to -\infty$. This implies that I can consider only values of the parameters which keep the search costs for clicked products lower than their corresponding upper bounds.

I notice that $\prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1})$ has a chain-like structure. Thus, I can sample random shocks iteratively. Let's assume I sampled some value of ξ_C with ξ_C^g being a realization of this random variable (thus, $\omega_C^g = \omega_C(\xi_C^g)$ also sampled). In this case, I may sample ξ_{C-1} in a way that $\omega_{C-1} \geq \omega_C$ (or $\xi_{C-1} \leq \min\{\omega_{C-1}^{-1}\left(\omega_C(\xi_C^g)\right), \bar{c}_{C-1}\}$) for random shock itself), however, I should adjust for probability of such event $F_{\xi_{C-1}}\left(\min\{\omega_{C-1}^{-1}\left(\omega_C(\xi_C^g)\right), \bar{c}_{C-1}\}\right)$. Notice that after generating ξ_{C-1}^g, I can repeat this procedure for ξ_{C-2} and so on.

After recursively applying of the procedure discussed in the previous paragraph for each j, I obtain the random sample $(\xi_1^g, .., \xi_C^g)$ such that $\prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1}) = 1$ but the probability needs to be adjusted by:

$$\prod_{j=1}^{C-1} F_{\xi_j}\left(\min\{\omega_j^{-1}\left(\omega_{j+1}(\xi_{j+1}^g)\right), \bar{c}_j\}\right) \tag{B.23}$$

Therefore, by using the recursive shock generation I can eliminate $\prod_{j=1}^{C-1} \mathcal{I}(\omega_j \geq \omega_{j+1})$ from Equation (B.21). Finally, notice that I can eliminate $\mathcal{I}(v_0 \leq \omega_C)$ from Equation (B.21) by sampling ξ_C from distribution such that $v_0 \leq \omega_C$ holds and adjust the probability by $F_{\xi_C}\left(\min\{\omega_C^{-1}(v_0), \bar{c}_C\}\right)$.

At the end, I summarize the procedure which I used to compute the objective function as follows:

1. Set the probability $1 \to W^g$
2. Generate random shock ξ_C^g from truncated $F_{\xi_C|\xi_C \leq \min\{\omega_C^{-1}(v_0), \bar{c}_C\}}$ and set $W^g \cdot F_{\xi_C|\xi_C \leq \min\{\omega_C^{-1}(v_0), \bar{c}_C\}} \to W^g$

3. If customer chose the non-outside option – generate random shock ψ_b^g from truncated $F_{\psi_b|\psi_b \geq \underline{\psi_b}}$ and set $W^g \cdot (1 - F_{\psi_b}(\underline{\psi_b})) \to W^g$, where $\underline{\psi_b}$ could be found from Equation (B.19) and Equation (B.20)

4. Compute $B(\xi_C^g, \psi_b^g)$ from Equations (B.14) to (B.20) and set $W^g \cdot B(\xi_C^g, \psi_b^g) \to W^G$

5. Generate recursively random shocks $(\xi_1^g, .., \xi_{C-1}^g)$ discussed in this subsection. Set $W^g \cdot \prod_{j=1}^{C-1} F_{\xi_j}\left(\min\{\omega_j^{-1}\left(\omega_{j+1}(\xi_{j+1}^g)\right), \bar{c}_j\}\right) \to W^g$

6. If customer did not purchase the last option clicked, the set

   $W^g \cdot \mathcal{I}\left[v_b(\psi_b^g) \leq \omega_C(\xi_C^g)\right] \to W^g$

7. Repeat steps (1-6) G times with different random seed. The estimate of likelihood for one individual would be $\widehat{P} = \sum_g W^g / G$

B.8 Identification

Consider the setting which closely follows the original model described in detail in Section 2.4.1. Assume the retailer has only one product in assortment. The customer i visiting the website receives the signal $\psi_i = \eta_i + \epsilon_i$ (η_i and ϵ_i have standard normal distribution and independent of each other) and has an option to buy the product or leave the website with utility 0. In case of a purchase the customer receives the unobserved preference shock ϵ_i and can return the product by paying R or keep it and receive utility $\beta_0^u + \epsilon_i$.

In this setting given the purchase and return data the retailer (or researcher) could identify parameters β_0^u and R. Firstly, notice that the expected purchase utility which takes into account the return option is simply Equation (2.3) with $\sigma_{\epsilon_{ij}} = \sigma_{\eta_{ij}} = 1$ and $\mu_{ij}^u = \beta_0^u$

$$v_{purchase}(\psi_i) = \frac{1}{\sqrt{2}} \cdot T\left(\sqrt{2}(\beta_0^u + R) + \frac{\psi_i}{\sqrt{2}}\right) - R \tag{B.24}$$

where $T(x) = \Phi(x) \cdot x + \varphi(x)$; $\Phi(x)$ and $\varphi(x)$ are standard normal cdf and pdf respectively and $\mathcal{I}(Condition) = 1$ if $Condition$ is $True$. $v_{purchase}(\psi_i)$ depends on ψ_i which is unknown to the researcher but known to the customer at the moment of purchase. Given the Equation (B.24) I can compute the probability that the customer chooses the outside option as:

$$\begin{aligned}
&\mathbb{P}[\text{choose outside}] = \mathbb{P}[v_{purchase}(\psi) < 0] = \\
=&\mathbb{P}\left[\frac{1}{\sqrt{2}} \cdot T\left(\sqrt{2}(\beta_0^u + R) + \frac{\psi_i}{\sqrt{2}} < 0\right) - R\right] \\
=&\mathbb{P}\left[\frac{\psi_i}{\sqrt{2}} < T^{-1}(\sqrt{2}R) - \sqrt{2}(\beta_0^u + R)\right] = \Phi\left[T^{-1}(\sqrt{2}R) - \sqrt{2}(\beta_0^u + R)\right]
\end{aligned} \tag{B.25}$$

Similarly, I can compute the probability of event when the customer makes a purchase and returns the product:

$$\begin{aligned}
&\mathbb{P}[\text{choose product, return it}] = \mathbb{P}[v_{purchase}(\psi) \geq 0, \beta_0^u + \epsilon_i \leq -R] = \\
=&\mathbb{E}_{\psi_i,\epsilon_i}\left[\mathcal{I}(v_{purchase}(\psi) \geq 0) \cdot \mathcal{I}(\beta_0^u + \epsilon_i \leq -R)\right] \\
=&\mathbb{E}_{\psi_i}\left[\mathcal{I}(v_{purchase}(\psi) \geq 0) \cdot \mathbb{E}_{\epsilon_i}\left[\mathcal{I}(\beta_0^u + \epsilon_i \leq -R)|\psi_i\right]\right]
\end{aligned} \tag{B.26}$$

Given the assumptions on ϵ_i and ψ I conclude that these variables have a join normal distribution and thus $\epsilon_i|\psi \sim \mathcal{N}(\frac{1}{2}\psi_i, \frac{1}{2})$ thus I can find:

$$
\begin{aligned}
&\mathbb{E}_{\epsilon_i}\left[\mathcal{I}(\beta_0^u + \epsilon_i \leq -R)|\psi_i\right] = \mathbb{P}[\epsilon_i \leq -R - \beta_0^u|\psi_i] \\
=&\mathbb{P}\left[\frac{\epsilon_i - \psi_i/2}{1/\sqrt{2}} \leq \frac{-R - \beta_0^u - \psi_i/2}{1/\sqrt{2}}|\psi_i\right] = \Phi\left(\frac{-R - \beta_0^u - \psi_i/2}{1/\sqrt{2}}\right)
\end{aligned}
\tag{B.27}
$$

Substituting Equation (B.27) in Equation (B.26) yields:

$$
\begin{aligned}
&\mathbb{P}[\text{choose product, return it}] \\
=&\mathbb{E}_{\psi_i}\left[\mathcal{I}(v_{purchase}(\psi) \geq 0) \cdot \Phi\left(\frac{-R - \beta_0^u - \psi_i/2}{1/\sqrt{2}}\right)\right] \\
=&\mathbb{E}_{\psi_i}\left[\mathcal{I}\left(\frac{\psi_i}{\sqrt{2}} < T^{-1}(\sqrt{2}R) - \sqrt{2}(\beta_0^u + R)\right) \cdot \Phi\left(\frac{-R - \beta_0^u - \psi_i/2}{1/\sqrt{2}}\right)\right]
\end{aligned}
\tag{B.28}
$$

In empirical setting the retailer could estimate on the data both $\mathbb{P}[\text{choose outside}]$ and $\mathbb{P}[\text{choose product, return it}]$ via event frequency ratios. I denote the estimates of these probabilities as $\widehat{P}_0$ and $\widehat{P}_1$ respectively. To find the estimates of the parameters β_0^u and R the retailer could solve the system of equations:

$$
\begin{aligned}
&\widehat{P}_0 = \Phi\left[T^{-1}(\sqrt{2}R) - \sqrt{2}(\beta_0^u + R)\right] \\
&\widehat{P}_1 = \mathbb{E}_{\psi_i}\left[\mathcal{I}\left(\frac{\psi_i}{\sqrt{2}} \geq T^{-1}(\sqrt{2}R) - \sqrt{2}(\beta_0^u + R)\right) \cdot \Phi\left(\frac{-R - \beta_0^u - \psi_i/2}{1/\sqrt{2}}\right)\right]
\end{aligned}
\tag{B.29}
$$

Notice that the first equation could be substituted into the second one and I get:

$$\widehat{P}_1 = \mathbb{E}_{\psi_i}\left[\mathcal{I}\left(\frac{\psi_i}{\sqrt{2}} \geq \Phi^{-1}(\widehat{P}_0)\right)\cdot\Phi\left(\frac{-R-\beta_0^u-\psi_i/2}{1/\sqrt{2}}\right)\right] \tag{B.30}$$

From Equation (B.30) if follows that I can make a substitution and solve the Equation (B.30) for parameter $\alpha = R + \beta_0^u$. Given that $\nu_i = \psi_i/\sqrt{2}$ has a standard normal distribution:

$$\begin{aligned}\widehat{P}_1 &= \mathbb{E}_{\nu_i}\left[\mathcal{I}\left(\nu_i \geq \Phi^{-1}(\widehat{P}_0)\right)\cdot\Phi\left(-\frac{\alpha}{1/\sqrt{2}}-\nu_i\right)\right] \\ &= \int_{\Phi^{-1}(\widehat{P}_0)}^{\infty}\Phi(-\alpha\sqrt{2}-\nu_i)\varphi(\nu_i)d\nu_i = B(\alpha)\end{aligned} \tag{B.31}$$

Notice that the integration range does not depend on parameters and because $\Phi(t)$ is an increasing function of t I have that $B(\alpha)$ is a decreasing function of α. Thus, I can conclude that the estimate of α could be found as $\widehat{\alpha} = B^{-1}(\widehat{P}_1) = Q(\widehat{P}_0, \widehat{P}_1)$, where I emphasize that it would depend on both estimates of probabilities $\widehat{P}_0$ and $\widehat{P}_1$. Given the estimate $\widehat{\alpha}$ I can use the first equation from Equation (B.29) to estimate R and β_0^u as:

$$\begin{aligned}\widehat{R} &= \frac{1}{\sqrt{2}}T\left(\Phi^{-1}(\widehat{P}_0)+\widehat{\alpha}\sqrt{2}\right) \\ \widehat{\beta_o^u} &= \widehat{\alpha}-\widehat{R}\end{aligned} \tag{B.32}$$

Equation (B.32) concludes the derivations and demonstrates that given the assumptions of the model the parameters of the model are identified.

Table B.1. Comparison with Alternative Estimation Methods

Variable	True parameter values	Paper approach	Simulated maximum likelihood	Smoothed simulated maximum likelihood	Maximizing the true likelihood
β_0^u	-1.4	-1.347	-0.285	-2.246	-1.267
β_1^u	-0.3	-0.280	0.204	1.988	-0.269
β_1^η	-0.5	-0.443	-0.093	2.365	-0.379
β_0^c	-4	-3.944	-0.155	-8.666	-3.841
β_1^c	0.3	0.312	0.047	0.261	0.311
R	-1.5	-1.427	-0.556	1.302	-1.091
Computation time, min	0	5	173	72	11
Relative memory usage	0	1	500	50	1

B.9 Additional Analysis on Synthetic Data

In Section 2.5 I discussed the estimation on simulated data. In this subsection, I compare the method with the methods that are widely used in literature discussed in Section 2.5.4. The results of the simulations are in Table B.1

Table B.1 demonstrates that the method provides the best accuracy and takes the least time to converge. The main efficiency gains come from the fact that the approach allows approximating the likelihood well with a very small number of random shocks (∼20) while the SML method fails to give a correct estimation even with 10,000 random shocks. The smoothed SML method provides an even worse quality for the reasons discussed in the main paper.

Next, I compare the two-stage approach in the paper with the direct maximization of the likelihood in Equation (B.3). First, I observe that the true likelihood maximization recovers parameters better than all other methods and the estimates are comparable with the true parameters. This supports the identifiability of the model. Second, I see that the parameters related to returns (3rd and 6th rows) have the worst accuracy and are shrunk towards zero. This observation supports the justification that the returns data is underweighted in the full maximization. Notice that this argument is even stronger on the real data which have longer sessions than in simulated data – this implies that the share of return constraints is even smaller. Finally, I note that by increasing the sample size, the true likelihood maximization improves the accuracy.

B.10 Description of the Greedy Algorithm for Ranking Optimization

Let's denote Ω a list of K ordered products. Assume this list ordered randomly and without loss of generality $\Omega = \{x_1, x_2, ...x_K\}$, where x_i characteristics of the i^{th} product in a list. Next, assume $\mathcal{V}(\cdot)$ is a function which takes as input an ordered list of products Ω and returns the number representing the "performance" of Ω. For example, in the paper $\mathcal{V}(\cdot)$ would be a procedure which simulates the behavior of customers who face some ordered list of products Ω on the website. In Section 2.8.2 the returning value of this function is the probability that the customer purchases a product and keeps it, however, any other value function is possible.

The algorithm which was used in Section 2.8.2 could be summarized as follows:

1. For each $k \in \{1..K\}$ repeat:

   For each $j \in \{k..K\}$ repeat:

Construct Ω_{jk} by switching positions of product j and k in a list Ω_k

Compute and store $\nu_{jk} = \mathcal{V}(\Omega_{jk})$

Choose the index $j^\star$ which maximizes ν_{jk}

Store $\Omega_k = \Omega_{j^\star k}$

2. Return Ω_K as the greedy-optimal ranking

Intuitively, I iteratively (one-by-one) optimize each position of the list while taking into account that products chosen as best on previous iterations remain on the same spot. It is straightforward to show that this algorithm requires $\sim K^2$ evaluations of function $\mathcal{V}(\cdot)$ which could be very slow to compute. Thus, it provides a feasible approximation to the complex problem of finding the best ranking.

B.11 Model Estimation Results

Table B.2. Results of Model Estimation on Real Data

Variable	Mean Utility	Signal Variance
Utility constant	-1.447	–
Relative Price	-0.003	–
Blouses	-0.191	-0.213
Pants	-0.046	0.012
Jackets and coats	0.014	0.105
Jeans	-0.054	-0.029
Dresses	0.080	0.242
Shoes	-0.166	-0.185
Polo Shirts	-0.376	-0.814
Tops	-0.319	-0.494
Knit	-0.223	-0.319
Sweat	-0.259	-0.416
Shorts	-0.293	-0.339
Skirts	-0.006	0.119
Blazer	-0.078	-0.026
Classic shirts	-0.272	-0.484
Brand: Denim	-0.009	0.019
Brand: Menplus	-0.025	0.012
Proportion of natural fabric	-0.08	-0.165
Color: Blue	-0.031	-0.057
Color: Gray	-0.020	-0.039
Color: White	0.047	0.123
Color: Green	0.019	-0.009
Color: Yellow	0.028	0.003
Color: Red	0.059	0.098
Search performed on a desktop (not mobile)	0.072	0.005
Search Costs Parameters		
Fixed Mean Search Costs	-4.092	–
Variable Position Search Costs (log)	0.303	–
Variable Page Search Costs	0.100	–
Search performed on a desktop (not mobile)	0.231	–
Return Costs Parameters		
Return Costs (log)	-1.3472	–

Appendix C

Appendix for Chapter 3

C.1 Simulation of Removing the Product from the Assortment

In Section 3.5 I discussed the results of "evaluating" the past decisions of the retailer or how the retailer's well-being would change if it would add some product to the assortment. However, a more practical problem is to evaluate the well-being after removing some product from the current assortment or, in other words, identify, which products are driving the profit/loss.

Figure Figure C-1 illustrates the results of the simulation of removing products from the assortment. In the figure, each point represents the change in the retailer's purchase and return rates after it would remove some products from the website.

The results look qualitatively similar to Figure 3-8, however, the segment interpretation would change. For example, *Quadrant IV, lower right*, represents the "best" products, as

Figure C-1. Replication of the main result by removing products from the assortment.

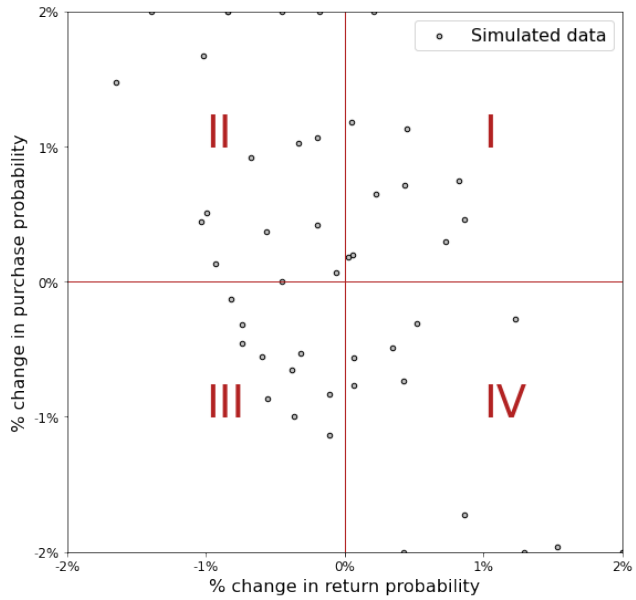

removing these products results in higher sales and lower return rates. Similarly, *Quadrant II, upper left*, would represent the products that induce losses.

C.2 Alternative Criteria to Remove Products from the Website

In Section 3.5 the products were grouped based on their predicted return and purchase rates. This was done to prevent the model to rely on the noisy estimates of return rate and allow it to be extended to adding a completely new product to the assortment (for example, a new fashion season). In this section, I split the products into groups based on their true purchase and return rate evaluated on the combination of train and validation data (to ensure reliable estimates). Figure C-2 replicates Figure 3-8 with alternative splitting criteria.

The results are qualitatively the same. Ultimately, the criterion is a less important step

Figure C-2. Alternative Product Splitting Criteria

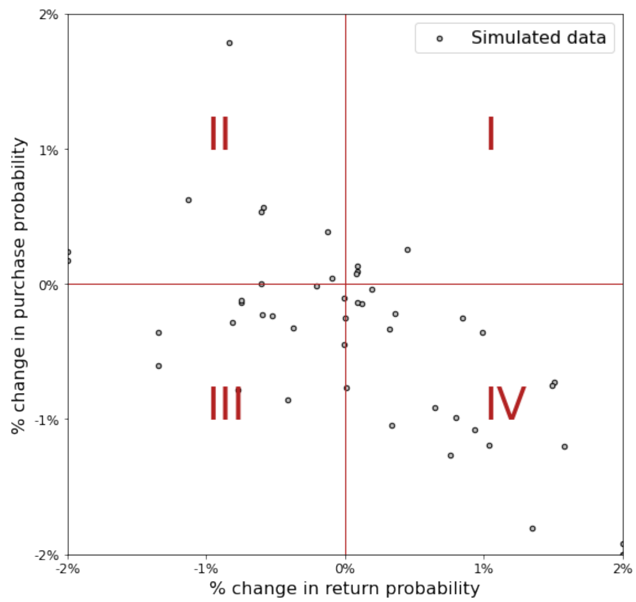

as the retailer potentially could run the simulation on the per product level – each product bin has exactly one product. For the paper, computing per product figure is infeasible as it would require simulating approximately 10,000 times. This would take approximately 138 days to compute. Moreover, due to the smaller changes to the website, it would require increasing the number of the simulated customer within a simulation further increasing the computation time. Practically, the criterion discussed in the paper could be used to rank or preselect products for per-product simulation. Further research would identify the optimal algorithm.

C.3 Alternative Scoring of the Products

Figure C-3 replicates the results in Figure 3-9 from Section 3.5. However, instead of using the results of the model, I rank the products based on the efficient purchase score $\mathbb{P}[purchase] \cdot (1 - \mathbb{P}[return])$. I define the efficient purchase score as and intuitively it demonstrates how likely the product would end up at customer's hands. Figure C-3 top-5 highest/lowest score products within each category.

Firstly, Figure C-3 demonstrate the visual differences between the best and worst products based on the score, which in turn supports the findings in the paper. Although, Figure C-3 and Figure 3-9 share some characteristics (left picture has more non-standard colors) they are noticeable different. This observation implies that the search step of the model is important as the "worst" product is not only the product with low sales and high returns, but also the product is very noticeable by the customer and drags their attention from more successful products.

Figure C-3. Highest and Lowest Ranked Products within Category Based on $\mathbb{P}[purchase] \cdot (1 - \mathbb{P}[return])$

(a) Low score.

(b) High score.

Bibliography

J. Abbey, M. Ketzenberg, and R. Metters. A more profitable approach to product returns. *MIT Sloan Management Review*, 60(1):1–6, 2018.

A. Alptekinoğlu and A. Grasas. When to carry eccentric products? optimal retail assortment under consumer returns. *Production and Operations Management*, 23(5):877–892, 2014. doi: https://doi.org/10.1111/poms.12081. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/poms.12081.

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

B. Asdecker. Returning mail-order goods: analyzing the relationship between the rate of returns and the associated costs. *Logistics Research*, 8:1–12, 2015.

N. N. Bechwati and W. S. Siegal. The impact of the prechoice process on product returns. *Journal of Marketing Research*, 42(3):358–367, 2005.

P. Boatwright and J. C. Nunes. Reducing assortment: An attribute-based approach. *Journal of Marketing*, 65(3):50–63, 2001. doi: 10.1509/jmkg.65.3.50.18330. URL https://doi.org/10.1509/jmkg.65.3.50.18330.

A. B. Bower and J. G. Maxham III. 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, 2012.

F. Branco, M. Sun, and J. M. Villas-Boas. Optimal search for product information. *Management Science*, 58(11):2037–2056, 2012.

S. M. Broniarczyk, W. D. Hoyer, and L. McAlister. Consumers' perceptions of the assortment offered in a grocery category: The impact of item reduction. *Journal of Marketing Research*, 35(2):166–176, 1998. ISSN 00222437. URL http://www.jstor.org/stable/3151845.

B. J. Bronnenberg, J. B. Kim, and C. F. Mela. Zooming in on choice: How do consumers search for cameras online? *Marketing science*, 35(5):693–712, 2016.

G. P. Cachon, C. Terwiesch, and Y. Xu. Retail assortment planning in the presence of consumer search. *Manufacturing & Service Operations Management*, 7(4):330–346, 2005. doi: 10.1287/msom.1050.0088. URL https://doi.org/10.1287/msom.1050.0088.

F. Caro and J. Gallien. Dynamic assortment with demand learning for seasonal consumer goods. *Management Science*, 53(2):276–292, 2007. doi: 10.1287/mnsc.1060.0613. URL https://doi.org/10.1287/mnsc.1060.0613.

Y.-K. Che. Customer return policies for experience goods. *The Journal of industrial economics*, pages 17–24, 1996.

X. Chen, Y. Wang, and Y. Zhou. Optimal policy for dynamic assortment planning under multinomial logit models. *Mathematics of Operations Research*, 46(4):1639–1657, 2021. doi: 10.1287/moor.2021.1133. URL https://doi.org/10.1287/moor.2021.1133.

Y. Chen and S. Yao. Sequential search with refinement: Model and application with click-stream data. *Management Science*, 63(12):4345–4365, 2017.

Z.-Y. Chen, Z.-P. Fan, and M. Sun. Machine learning methods for data-driven demand estimation and assortment planning considering cross-selling and substitutions. *INFORMS Journal on Computing*, 35(1):158–177, 2023. doi: 10.1287/ijoc.2022.1251. URL https://doi.org/10.1287/ijoc.2022.1251.

M. Conlin, T. O'Donoghue, and T. J. Vogelsang. Projection bias in catalog orders. *American Economic Review*, 97(4):1217–1249, 2007.

H. Cui, S. Rajagopalan, and A. R. Ward. Predicting product return volume using machine learning methods. *European Journal of Operational Research*, 281(3):612–627, 2020.

S. Davis, M. Hagerty, and E. Gerstner. Return policies and the optimal level of "hassle". *Journal of Economics and Business*, 50(5):445–460, 1998.

R. Dew, A. Ansari, and O. Toubia. Letting logos speak: Leveraging multiview representation learning for data-driven branding and logo design. *Marketing Science*, 41(2):401–425, 2022.

R. O. Duda and P. E. Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972.

A. Dukes and L. Liu. Online shopping intermediaries: The strategic design of search environments. *Management Science*, 62(4):1064–1077, 2016.

D. Dzyabura and J. R. Hauser. Recommending products when consumers learn their preference weights. *Marketing Science*, 38(3):417–441, 2019.

D. Dzyabura, S. El Kihal, J. R. Hauser, and M. Ibragimov. Leveraging the power of images in managing product return rates. *Available at SSRN 3209307*, 2019.

D. Dzyabura, S. El Kihal, and R. Peres. Image analytics in marketing. In *Handbook of Market Research*, pages 665–692. Springer, 2021.

S. El Kihal and E. Shehu. It's not only what they buy, it's also what they keep: Linking marketing instruments to product returns. *Journal of Retailing*, 98(3):558–571, 2022.

S. El Kihal, N. Nurullayev, C. Schulze, and B. Skiera. A comparison of return rate calculation methods: Evidence from 16 retailers. *Journal of Retailing*, 97(4):676–696, 2021.

Emma Thomasson. Online retailers go hi-tech to size up shoppers and cut returns, 2013. URL https://www.reuters.com/article/net-us-retail-online-returns-idUSBRE98Q0GS20131002. *Reuters*.

R. Gonzales and R. Woods. Digital image processing 4th edition, 2018.

J. Hartmann, M. Heitmann, C. Schamp, and O. Netzer. The power of brand selfies. *Journal of Marketing Research*, 58(6):1159–1177, 2021.

J. R. Hauser. Testing the accuracy, usefulness, and significance of probabilistic choice models: An information-theoretic approach. *Operations Research*, 26(3):406–421, 1978.

K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

R. He and J. McAuley. Vbpr: visual bayesian personalized ranking from implicit feedback. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.

J. Hernández-Orallo. Roc curves for regression. *Pattern Recognition*, 46(12):3395–3411, 2013.

Y. Hong and P. A. Pavlou. Product fit uncertainty in online markets: Nature, effects, and antecedents. *Information Systems Research*, 25(2):328–344, 2014.

E. Honka and P. Chintagunta. Simultaneous or sequential? search strategies in the us auto insurance industry. *Marketing Science*, 36(1):21–42, 2017.

M. Ibragimov. Customer and product returns. 2023.

iBusiness. Wie Shopbetreiber das Retourenproblem wirklich lösen, 2016. URL https://www.ibusiness.de/aktuell/db/858729veg.html. Accessed on Day Month Year.

N. Janakiraman, H. A. Syrdal, and R. Freling. The effect of return policy leniency on consumer purchase and return decisions: A meta-analytic review. *Journal of retailing*, 92 (2):226–235, 2016.

Jinjoo Lee. Retailers' many unhappy returns, 2022. URL https://www.wsj.com/articles/retailers-many-unhappy-returns-11641387605. *The Wall Street Journal*.

T. T. Ke, Z.-J. M. Shen, and J. M. Villas-Boas. Search for information on multiple products. *Management Science*, 62(12):3576–3603, 2016.

J. B. Kim, P. Albuquerque, and B. J. Bronnenberg. Online demand under limited consumer search. *Marketing science*, 29(6):1001–1023, 2010.

J. Klostermann, A. Plumeyer, D. Böger, and R. Decker. Extracting brand information from social networks: Integrating image, text, and social tagging data. *International Journal of Research in Marketing*, 35(4):538–556, 2018.

A. G. Kök and Y. Xu. Optimal and competitive assortments with endogenous pricing under hierarchical consumer choice models. *Management Science*, 57(9):1546–1563, 2011. doi: 10.1287/mnsc.1110.1402. URL https://doi.org/10.1287/mnsc.1110.1402.

A. G. Kök, M. L. Fisher, and R. Vaidyanathan. Assortment planning: Review of literature and industry practice, 2009. URL https://doi.org/10.1007/978-0-387-78902-6_6.

J. Liu and O. Toubia. A semantic approach for estimating consumer content preferences from online search queries. *Marketing Science*, 37(6):930–952, 2018.

L. Liu, D. Dzyabura, and N. Mizik. Visual listening in: Extracting brand image portrayed on social media. *Marketing Science*, 39(4):669–686, 2020.

S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

C. Lynch, K. Aryafar, and J. Attenberg. Images don't lie: Transferring deep visual semantic features to large-scale multimodal learning to rank. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 541–548, 2016.

S. Mahajan and G. J. van Ryzin. Retail inventories and consumer choice, 1999. URL https://doi.org/10.1007/978-1-4615-4949-9_17.

B. S. Manjunath and W.-Y. Ma. Texture features for browsing and retrieval of image data. *IEEE Transactions on pattern analysis and machine intelligence*, 18(8):837–842, 1996.

M. K. Mantrala, M. Levy, B. E. Kahn, E. J. Fox, P. Gaidarev, B. Dankworth, and D. Shah. Why is assortment planning so difficult for retailers? a framework and research agenda. *Journal of Retailing*, 85(1):71–83, 2009. ISSN 0022-4359. doi: https://doi.org/10.1016/j.jretai.2008.11.006. URL https://www.sciencedirect.com/science/article/pii/S0022435908000882. Enhancing the Retail Customer Experience.

J. McAuley, C. Targett, Q. Shi, and A. Van Den Hengel. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 43–52, 2015.

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

I. Morozov, S. Seiler, X. Dong, and L. Hou. Estimation of preference heterogeneity in markets with costly search. *Marketing Science*, 40(5):871–899, 2021.

U. Narang and V. Shankar. Mobile app introduction and online and offline purchases and product returns. *Marketing Science*, 38(5):756–772, 2019.

Noor Nanji. Zara starts charging shoppers for online returns, 2022. URL https://www-bbc-com.cdn.ampproject.org/c/s/www.bbc.com/news/business-61423753.amp. Accessed on Day Month Year.

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

J. A. Petersen and V. Kumar. Can product returns make you money? *MIT Sloan Management Review*, 2010.

J. A. Petersen and V. Kumar. Perceived risk, product returns, and optimal resource allocation: Evidence from a field experiment. *Journal of Marketing Research*, 52(2):268–285, 2015.

J. A. Quelch and D. Kenny. Extend profits, not product lines. *Harvard Business Review*, 72 (5):153–160, 1994.

R. P. Rooderkerk and A. G. Kök. Omnichannel assortment planning, 2019. URL https://doi.org/10.1007/978-3-030-20119-7_4.

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

D. Sauré and A. Zeevi. Optimal dynamic assortment planning with demand learning. *Manufacturing & Service Operations Management*, 15(3):387–404, 2013. doi: 10.1287/msom.2013.0429. URL https://doi.org/10.1287/msom.2013.0429.

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

M. Shi, C. Chussid, P. Yang, M. Jia, V. Dyk Lewis, and W. Cao. The exploration of artificial intelligence application in fashion trend forecasting. *Textile Research Journal*, 91(19-20): 2357–2386, 2021.

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

G. J. Stigler. The economics of information. *Journal of political economy*, 69(3):213–225, 1961.

R. M. Ursu. The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4):530–552, 2018.

R. M. Ursu, Q. Wang, and P. K. Chintagunta. Search duration. *Marketing Science*, 39(5): 849–871, 2020.

R. M. Ursu, Q. Zhang, and E. Honka. Search gaps and consumer fatigue. *Marketing Science*, 42(1):110–136, 2023.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need, 2017.

R. Wang and O. Sahin. The impact of consumer search cost on assortment planning and pricing. *Management Science*, 64(8):3649–3666, 2018. doi: 10.1287/mnsc.2017.2790. URL https://doi.org/10.1287/mnsc.2017.2790.

M. L. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

S. L. Wood. Remote purchase environments: The influence of return policy leniency on two-stage decision processes. *Journal of Marketing Research*, 38(2):157–169, 2001.

M. Zhang and L. Luo. Can consumer-posted photos serve as a leading indicator of restaurant survival? evidence from yelp. *Management Science*, 69(1):25–50, 2023.

S. Zhang, D. Lee, P. V. Signgh, and K. Srinivasan. What makes a good image? airbnb demand analytics leveraging interpretable image features. *Management science*, 68(8): 5644–5666, 2021.