Return-Aware Search: Jointly Modeling Clicks, Purchases, and Returns in Online Retail*

Marat Ibragimov † Siham El Kihal‡ John Hauser§ Raluca Ursu¶

February 6, 2026

Abstract

The increasing volume of product returns has become a substantial burden for online retailers. While existing research focuses on managing returns through a purchase/return framework, this paper argues that prepurchase/search actions provide an early warning of return behavior. Using data from a large European apparel retailer, we propose a joint model of customer clicks, purchases, and returns. First, a reduced-form model demonstrates the empirical link between search and returns: purchasing the last-clicked product, browsing fewer items, and lower price variety predict a lower return probability. Next, we extend the Weitzman search model to include returns, demonstrating that omitting the return decision results in biased estimates of customer preferences. Finally, we show that by reordering products on the website according to the proposed model, retailers can effectively manage product return rates and improve overall profit.

Keywords: consumer search, product returns, discrete choices, structural modeling, e-retailing

*This paper has benefited from feedback obtained at the 44th INFORMS Marketing Science Conference 2022, 11th Workshop on Consumer Search and Switching Costs (NYU Stern), ISPOC 2022 Series, 45th INFORMS Marketing Science Conference 2023, and the seminars at MIT Sloan, Frankfurt School of Finance & Management, Boston College, UT Dallas, University of Utah, Emory University, University of Kansas, University of Delaware, CU Boulder, Purdue University, Rutgers University. We wish to thank Douglas Bowman, Giovanni Compiani, Daria Dzyabura, Zhengrong Gu, Tetyana Kosyakova, Ilya Morozov, Artem Timoshenko, David Schweidel, Duncan Simester, and Jeremy Yang for their valuable comments and suggestions.

†Goizueta Business School, Emory University

‡Vienna University of Economics and Business

§Sloan School of Management, Massachusetts Institute of Technology

¶Stern School of Business, New York University

1 Introduction

Understanding and managing product returns is an important challenge for online retailers (The New Yorker 2023). Returns significantly reduce profits by simultaneously reducing revenue (through refunds) and increasing costs (e.g., reverse logistics, dry cleaning). For example, L.L. Bean has spent $50 million annually on return costs during 2012-2017, amounting to about 30% of the retailer's annual profits (Abbey et al. 2018). Return costs are often so high that major online retailers, such as Amazon and Walmart, have begun to allow customers to keep the product, because extracting benefits from the returned product is less than the return costs (The Wall Street Journal 2022). Similarly, Zara recently began charging online shoppers for returns unless the products are returned to a physical store (BBC 2022). Overall, about three-quarters of retailers charge for at least one return method (National Retail Foundation 2025).

Product returns are typically studied within a purchase-to-return framework. In this framework, research has established that product characteristics jointly affect the probability of purchase and return because the option to return a product is valuable and it impacts the purchase decision. From a managerial perspective, research suggests that changes in policies aimed at reducing returns (e.g., toward a stricter policy) must be evaluated relative to the potential negative effects on customers' purchase behavior.

Although frequently used, the purchase-to-return framework overlooks an essential component of the customer journey – customer search. Before making a purchase, customers spend significant time exploring the retailer's website. They review different product characteristics, compare alternatives, and click on products they like. Only after gathering sufficient information do they make a purchase. The existing purchase-to-return framework overlooks search and could substantially limit retailers' options to better manage frequent product returns, for example, by modifying the search environment (e.g., the ranking of the products on the website).

By contrast, traditional search-to-purchase models focus extensively on how the search environment shapes consumer behavior, yet they treat the purchase decision as the final point of the customer journey. By neglecting post-purchase actions such as returns, these models fail to account for the potential interdependence between search patterns and return outcomes. This theoretical gap can lead to flawed inferences regarding consumer preferences and sub-optimal managerial decisions. For instance, a search-to-purchase only framework might identify a highly popular product as a prime candidate for promotion to drive revenue; however, if that same product exhibits a disproportionately high return rate, such a strategy would ultimately erode the firm's profitability rather than enhance it.

In this paper, we demonstrate empirically that customer actions during search inform retailers about potential returns. We also show that retailers can effectively mitigate high return rates by strategically optimizing the customer search environment. Our analysis uses a unique dataset from a major European apparel retailer that captures all customer actions from opening the retailer's website to the final return decision. In particular, we observe the products consumers were shown in response to their queries, the products they clicked and purchased, as well as whether they returned any of the purchased items. We document that specific customer search patterns, such as the average price clicked or the number of clicks made, correlate with the probability of returning the product, even after controlling for product characteristics.

These findings motivate the development of an integrated structural model that unifies search, purchase, and return decisions within a single framework. We demonstrate that this joint model can be successfully estimated using standard industry data and show that ignoring the return stage leads to biased inferences regarding consumer preferences. Finally, we show that leveraging our model to inform product rankings is an effective, low-cost strategy for retailers to manage and reduce the volume of product returns in the online channel.

The remainder of the paper is organized as follows. Section 2 reviews relevant literature on product returns and customer search. Section 3 describes the data used in the analysis and documents model-free evidence that motivates the joint model. Section 4 develops our model and Section 5 shows that the proposed model can be estimated, and that omitting returns can lead to a substantial bias in the estimated parameters. Section 6 presents model results using our data and Section 7 performs a counterfactual where we re-rank products with the goal of minimizing returns. Section 8 provides a summary and suggestions for future research.

2 Related Research

Our paper contributes to the literature on product returns and customer search. Existing research on product returns spans both theoretical and empirical approaches. Theoretically, researchers have focused on return policies to demonstrate that the option to return products serves as a risk-reducing mechanism that encourages customers to experience the product (Che 1996) or as a signal of product quality (Moorthy and Srinivasan 1995).

Empirical research has focused on firms' optimization of return policies. In an attempt to identify the optimal return policy, researchers recognize the trade-off 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) extensively reviews the effect of return policy leniency on purchases and returns. Anderson et al. (2009) propose a structural model in which the option to return is embedded in a customer's purchase decision: the customer learns private information only after purchasing the product. Other empirical studies demonstrate that a variety of policy 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, 2015, Sahoo et al. 2018, Shehu et al. 2020, El Kihal and Shehu 2022). Empirical studies suggest prescriptive instruments, such as visualization systems, to decrease return rates. These instruments reduce return rates by reducing uncertainty about the product match with the customer (Hong and Pavlou 2014). Other researchers use machine learning to accurately predict returns and identify product-related features that enable the firm to better select and design fashion products for the retailer's website (Cui et al. 2020, Dzyabura et al. 2021). Overall, existing empirical research on product returns focuses on post-purchase customer behavior; in our paper, we demonstrate that events preceding the purchase can inform the customer return decision.

The other field of research we contribute to – customer search – is an established, mature field. The literature typically follows either a sequential (Weitzman 1979) or a simultaneous (Stigler 1961) approach. Both approaches assume the customer knows the distribution of the rewards and seeks to resolve uncertainty about products through search before making a final choice.

The availability of click-stream data has enabled researchers to empirically study customer search behavior (Bronnenberg et al. 2016, Chen and Yao 2017, Ursu et al. 2020) and to provide detailed insights into the search-to-purchase component of the customer journey. For example, Bronnenberg et al. (2016) examines customer search behavior for cameras and shows that early search is highly predictive of customer purchase and that the first-time discovery of the purchased alternative happens towards the end of the search. Chen and Yao (2017) shows that refinement tools significantly impact customer behavior and market structure. Ursu (2018) shows that product ranked lists affect search more than purchase decisions. Jiang et al. (2021) demonstrate the value of search data in firms' retargeting efforts. 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 (Gardete and Hunter 2024), intermediaries (Dukes and Liu 2016), search duration (Ursu et al. 2020), and search fatigue (Ursu et al. 2023).

Within the search literature, our paper is most closely related to research on customer learning and utility uncertainty. In this body of work, authors typically assume that a single product inspection yields only a noisy signal rather than precise information about the product's utility; subsequent searches are then required to increase the precision of that signal (Ke et al. 2016, Branco et al. 2012, Ke and Villas-Boas 2019, Gardete and Hunter 2024, Chick and Frazier 2012, Ursu et al. 2020). Our model shares the assumption that the customer possesses only partial information upon the initial click and has to make a purchase decision under uncertainty. However, in our framework, the final utility is revealed (or the learning occurs) after the purchase, and the customer can return the product for a full refund.

To date, researchers have focused primarily on the purchase-to-return sub-journey (returns literature) or the search-to-purchase sub-journey (search literature). Research on the search-to-purchase-to-return is scarce and uses a theoretical lens (Jerath and Ren 2024, Janssen and Williams 2024). We expand these research streams to focus on the entire search-to-purchase-to-returns journey in the empirical setting. Our research provides complementary insights to the returns literature (search predicts returns) and to the search literature (the possibility of returning a product changes a customer's optimal sequential search strategy). We demonstrate that by focusing on the entire customer journey, we gain additional insight into customer behavior and explore when existing models may fail. We also demonstrate the practical value of the proposed model by improving firms' profits through changes to product rankings on the website.

3 Data

3.1 Data Summary

We use individual-level data from the online channel of a large European apparel retailer. We focus on the online channel because (1) the online channel is substantially more susceptible to product returns (Dzyabura et al. 2023), and (2) the online channel is an ideal setting in which to observe all customer actions, including clicks, purchases, and returns.

The retailer sells medium-priced fashion products for women, men, and children. Its main product is adult clothing, which accounts for 95% of purchases. As is typical for Europe, the retailer has a generous return policy. Products can be returned for free within 60 days after purchase for a full refund, with or without a reason for returning. Prior to our analysis, the retailer did not use specific policies to discourage customer returns.

Standard to the search literature, we observe detailed information about customer behavior on the retailer's website. Specifically, we capture all customer actions from their visit to the website to the termination of their session. Data include the set of products presented to the customer, as well as the sequence of customer clicks on product pages (if any). Consistent with prior work, we interpret clicks as search decisions (Chen and Yao 2017, Ursu 2018, Morozov et al. 2021). For all customer visits to the website, we observe the outcome – whether the customer purchased a product (and which one) or left without purchasing anything.

The unique aspect of our data is that, after a customer purchases a product, we additionally observe whether the product was kept or returned. This contrasts with the classical search settings (Bronnenberg et al. 2016, Honka and Chintagunta 2017, Ursu et al. 2025), where the observation period ends with the customer's purchase decision. In industries where product returns (or cancellations) are frequent, omitting this final stage may result in a suboptimal policy and potentially hurt retailers' profits and customer satisfaction.

Our observation period is October 1, 2019 to February 28, 2020[1]. The retailer tracks individual user actions on the website via a third-party company that employs various technologies to capture detailed customer behavior during a browsing session. We provide further details on the construction of sessions in Appendix Appendix A. To better isolate the relationship between search and returns, we focus on sessions in which at most one item was purchased.[2] Extending our analysis to incorporate multiple items is an interesting but challenging task that lies beyond the scope of the present work and is left for future research.

Customers can access the retailer's website through a desktop or mobile device (54.8% accessed through a mobile device in our data). On the website, the customer can browse all products and choose to observe only products from the preferred category. Customers can inspect products from a specific category either by navigating through the menu provided by the retailer (for example, by choosing "Women" and then "Jeans" from the drop-down menu) or through the search bar (for example, by typing in "women's jeans"). This step takes the consumer to a product list, which displays 48 items per page (across several pages), each with a small image of the product, its price, and name. When the customer clicks on a specific product, further information is revealed on the product page, such as more (and higher quality) product images, information on sizes and availability, and detailed product descriptions. To illustrate the information available to the customer, we provide the retailer's website screenshot in Appendix Appendix A.

After all preprocessing steps, our data contain 756,612 browsing sessions.[3] Of these, 45,990 (6.1%) resulted in a purchase; notably, 41.1% of these purchases were subsequently

[1]We have access to data until May 15, 2020. However, we exclude the months when the COVID-related restrictions took place in the country where our retailer primarily operates.

[2]37.4% of sessions with a purchase involve only a single item. We randomly subsampled sessions without a purchase to ensure the purchase rate of 37.4% remains preserved.

[3]We outline the primary preprocessing steps in Appendix Appendix A.

Table 1: Category-level Summary Statistics

Category	# Sessions	Average # Clicks	Purchase Rate	Return Rate
Jackets & Coats	107,753	2.5	0.064	0.454
Knitwear	46,066	3.1	0.044	0.350
Jeans	29,206	2.4	0.083	0.369
Dresses	27,854	2.3	0.044	0.560
Shoes	26,263	1.8	0.067	0.370
Shirts & Tops	15,461	3.0	0.033	0.313
Blouses	14,743	2.7	0.038	0.443
Trousers	14,659	2.4	0.056	0.396
Sweatshirts	11,304	2.3	0.056	0.290
Shirts & Polos	7,930	3.1	0.059	0.099
Shirts	6,914	3.2	0.050	0.277
Skirts	6,293	1.8	0.041	0.452
Blazers	5,139	2.2	0.064	0.482
Jumpsuits	575	1.4	0.087	0.620
Shorts	250	2.7	0.080	0.150

returned. In 42% of sessions, customers searched products only within one category. Table 1 summarizes customer search, purchase, and return decisions across the retailer's product categories for those sessions where consumers searched only one category. There is substantial heterogeneity in how customers search for different types of products. Average clicks range from 1.4 to 3.2 across categories, and purchase rates can be as high as 8.7%. Notably, most categories exhibit return rates above 30%, underscoring the importance of understanding return behavior. Because of the high heterogeneity across categories, in our modeling section, we estimate the model one category at a time. There are no theoretical constraints preventing estimation on pooled cross-category data; however, doing so would require a substantially expanded vector of product attributes that includes category indicators and their interactions with product characteristics.

The retailer offers approximately 10,000 products in its assortment, which makes it infeasible to include product-level fixed effects in the structural model. However, certain product attributes may differentially affect customer valuations. To capture these effects, we will use a set of interpretable characteristics extracted from product images, following (Dzyabura et al. 2023), in addition to product prices, material, and color information provided by the retailer. A detailed list of these characteristics can be found in Table 2. For brevity, we focus on interpretable features; however, if the retailer's primary goal is predictive performance, they could additionally incorporate deep-learning-based image embeddings that capture more nuanced characteristics of the product.

Table 2: Product Characteristics Extracted and Used in the Analysis

Feature	Description
Price	Price of the product listed on the retailer's website
Natural material share	Share of composition that is natural materials (e.g., cotton, wool)
Color	Dummy-coded color categories provided by the retailer (e.g., red, blue, black)
Asymmetry	Degree to which the left and right halves of image differ in structure or shape
Brightness (mean)	Average luminance of the product image
Brightness (variance)	Variation in luminance capturing contrast between bright and dark regions
Color asymmetry (horizontal)	Difference in color distribution between the left and right halves of image
Color asymmetry (vertical)	Difference in color distribution between the top and bottom halves of image
Horizontal lines	Strength or count of horizontal edges detected in image
Vertical lines	Strength or count of vertical edges detected in image
Complexity of pattern	Measure of the amount of texture or pattern richness in image
Number of straight lines	Number of straight edge segments (Hough-transform based)

3.2 Model-Free Evidence

Prior work has shown how consumer search decisions relate to their purchase decisions (Honka et al. 2024, Bronnenberg et al. 2016). Here, we additionally provide evidence on the relation between search and return probabilities of the purchased products.

Towards this end, we estimate the following linear specification:

$$Return_i = \alpha_{j(i)} + \gamma_{t(i)} + X_i^{search}\beta + \epsilon_i, \tag{1}$$

where $Return_i$ is a dummy for whether the purchased product was returned or not, X_i^{search} is a vector summarizing the customer's search behavior (summarized in Table 3), $\alpha_{j(i)}$ represents the regression-model fixed effects of the product j purchased by customer i, and $\gamma_{t(i)}$ represents time fixed effects. For this analysis, we focus only on sessions that ended with a purchase, as we cannot observe the return outcome for consumers who did not make a purchase on the website. Including product fixed effects in the regression is important since some products may have intrinsically higher return rates (Table 1). These fixed effects allow us to explore how differences in customer return decisions vary depending on their search behavior, after controlling for product-specific characteristics. Finally, time fixed effects control for seasonality or weekday-specific behavior.

Table 3: Search Variables Definitions

Feature	Description
First/Last clicked product purchase	Indicator for whether the purchased product was the first/last clicked
Log(Number of clicks)	Number of clicked products during the session
Average price	Average price of clicked products
Price variability (std)	Standard deviation of prices of clicked products
Number of colors	Number of unique colors among clicked products
Log(Session length, sec)	Length of browsing session (seconds)
Log(Position on page)	On-page rank of the purchased product (1–48)
Low/High Returns	Share of clicked products with low/high historical return rates

Our results can be found in Table 4. Customers who click on many products are more likely to return the purchased item, consistent with the idea that extensive search reflects greater uncertainty or difficulty in finding a product matching consumer preferences. Conversely, purchasing the last-clicked product is negatively associated with returns, which may suggest that customers settle on an item that best fits their needs and therefore require less exploration (i.e., they do not need to look at other items or inspect previously searched items). Additionally, customers are more likely to return expensive products. Relatedly, customers who clicked on products spanning a wider range of prices are more likely to return the item. This pattern is consistent with preference uncertainty - customers who are unsure about what they want, evaluate products across broad price tiers, increasing the chance of a mismatch and, ultimately, a return.

These results are descriptive and should not be interpreted causally. For example, forcing customers to click fewer products would not necessarily reduce return rates. Instead, these reduced-form findings highlight that returns are systematically related to customer search behavior. Such evidence can be used to detect early signals of high return risk or to guide experimental designs aimed at identifying causal mechanisms.

In the next sections, we develop a structural model that rationalizes these relationships and helps us understand how search and returns are related and how retailers can leverage this understanding to improve ranking algorithms.

4 Model Development

Customer search is frequently examined through the lens of the Weitzman framework (Weitzman 1979). In this classic model, a customer faces multiple "boxes" – search options with uncertain rewards – and can sequentially open these boxes to resolve uncertainty. The

Table 4: Reduced-form Evidence: Relationship Between the Return and Search Behavior

	Return (dummy)
First clicked product purchase	-0.0015 (0.0075)
Last clicked product purchase	-0.0238^{***} (0.0080)
Log(Number of clicks)	0.0276^{***} (0.0085)
Average price	0.0005^{***} (0.0002)
Price variability (std)	0.0008^{***} (0.0002)
Number of colors	-0.0081^{**} (0.0032)
Log(Session length, sec)	0.0005 (0.0022)
Log(Position on page)	-0.0022 (0.0022)
Low returns	0.0176 (0.0142)
High returns	-0.0014 (0.0147)
Session with only one click	0.0101 (0.0102)
Product Fixed Effects	Yes
Weekday Fixed Effects	Yes
Month Fixed Effects	Yes
Observations	45,990
F-statistic	10.358

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

customer knows the reward distribution for each box and incurs a search cost to open it. Weitzman showed that the optimal search strategy is characterized by a set of reservation utilities that determine (i) the optimal order in which boxes should be opened, (ii) when the search process should stop, and (iii) which product the customer should choose to purchase once they stop searching.

The Weitzman search framework has been widely adopted in empirical work on customer search. In these applications, each product – for example, a hotel on Booking.com – is treated as a box, and clicking to reveal additional product information is interpreted as opening the box (see Section 2 for details).

A typical assumption in the Weitzman framework is that opening a box fully reveals the true value of an option. Consequently, a customer who chooses to purchase a product is assumed to have perfect information and therefore has no reason to return it. This assumption is violated in many contemporary retail settings. As shown in Section 3.2, product returns are prevalent, and there is a systematic relationship between the attributes customers search for and their subsequent likelihood of returning the product.

To address this limitation, we develop an extension of the Weitzman model designed for environments with frequent returns. In our model, opening a "box" does not fully resolve uncertainty. For example, information available after search – such as product descriptions or images – may be noisy or incomplete. Only after receiving the product at home does the customer fully observe its realized utility. At that point, the customer may choose to return the product, incurring a fixed cost (which may include monetary, time, or psychological components) or keep it.

In summary, we conceptualize customer decision-making as a sequential, three-stage process: (1) deciding on which products to click (and in what order); (2) selecting a product to purchase based on available (but imperfect) pre-purchase information; and (3) upon receiving the product, deciding whether to keep or return it, based on physical product inspection. For clarity, we refer to the classic empirical Weitzman model as the *No-returns Model* and our proposed extension as the *Joint Model*.

4.1 Model Overview

Figure 1 illustrates the *Joint Model*. Consider a customer who visits the retailer's website and observes the list of products $\mathcal{V}_i$. By observing a set of available alternatives, the customer gathers initial information about the product, which consists of product-related characteristics x_{ij} (price, category, color, etc.) and an individual pre-click preference shock ξ_{ij}. The customer can click on any of these products to reveal additional post-click information ϵ_{ij}. However, each click incurs a click cost of c_{ij}, for example, the customer may need to move the mouse to a certain part of the webpage, click, and process the information revealed). After clicking on a product, the customer either continues clicking on other options or stops to purchase any of the products clicked so far. After deciding to terminate the search, the customer purchases the best product among those clicked or leaves the website without a purchase (with an outside option). If the purchase is made, the customer receives the product and inspects it in more detail at home (e.g., tries it on, holds it up, feels the material, and compares its fashion match to the customer's other apparel products). Inspection reveals additional information (for example, fit with the body type). All revealed information is denoted as the ψ_{ij}. Based on all information accumulated (online clicks and offline inspection), the customer decides whether to keep the product or to incur a return cost R (e.g., return label, travel time, etc.) to return the product to the retailer for a full refund.

Figure 1: Overview of the Joint Model.

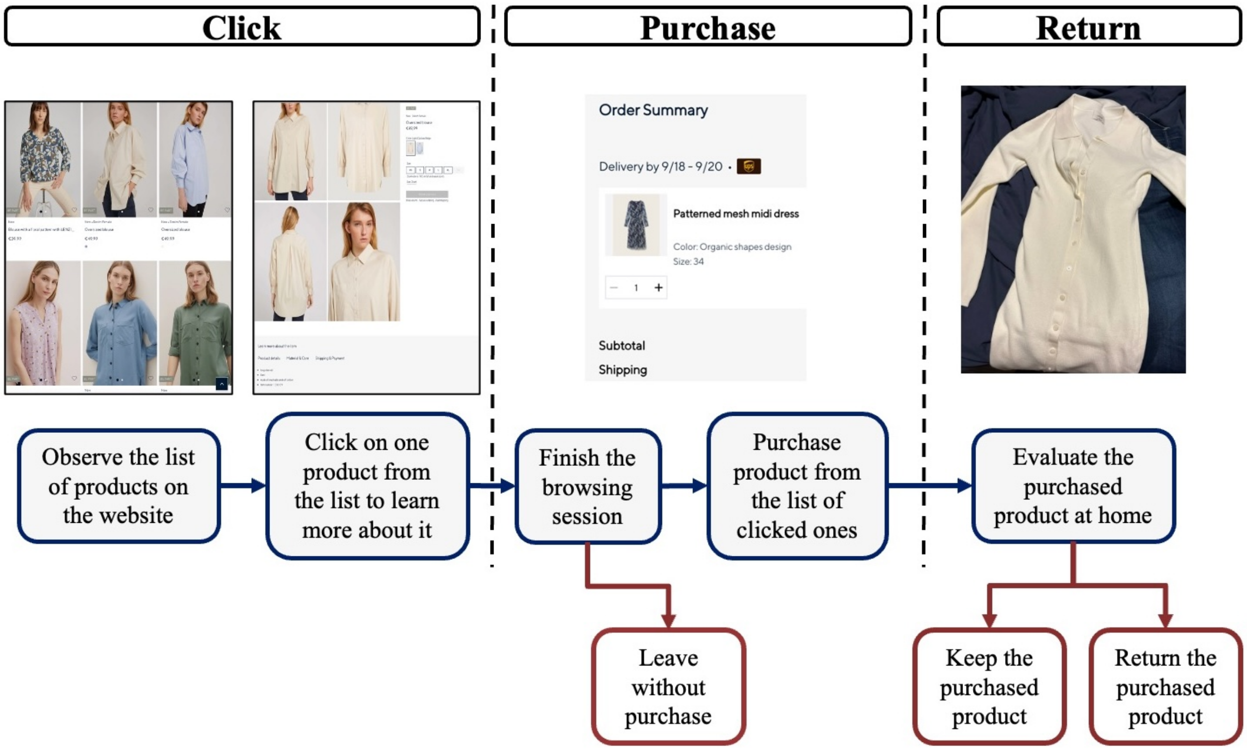

The main divergence from the Weitzman framework is ψ_{ij}, which is not revealed upon paying cost c_{ij} to click on the product. Specifically, the *Joint Model* assumes that part of the final utility ψ_{ij} is revealed only after inspection at home, thus the customer would have to make a click and purchase decision with limited information about this part of the utility.

4.2 Click Costs, Utility, and Returns

For ease of notation and without loss of generality, we number products such that j indexes the order in which customer i clicks on products (for example, $j = 2$ implies the second clicked product, while $j = 0$ implies the outside option, which is always available). The customer's final utility could take one of three possible forms (click costs are paid before the realization of this utility and thus not included in the equation):

$$u_{ij} = \begin{cases} \mu_{ij} + \xi_{ij} + \epsilon_{ij} + \psi_{ij} & \text{purchased and kept a product} j \neq 0 \\ -R & \text{purchased and returned a product} j \neq 0 \\ 0 & \text{chose outside option} j = 0 \end{cases} \tag{2}$$

where μ_{ij} is the customer's preference for the attributes x_{ij} of product j.

To make the estimation and identification of the model feasible, we impose additional assumptions on the structure of preference shocks. We follow the empirical Weitzman search framework (Ursu 2018, Honka and Chintagunta 2017), and assume that individual preference shocks $(\xi_{ij}, \epsilon_{ij}, \psi_{ij})$ conditionally on observed product characteristics x_{ij} are independent and normally distributed. Because the utility is invariant to a positive linear transformation, we normalize the variance of after-click and before-click preference shocks to 1. This normalization helps us focus on product returns and on comparisons to reduced models.

Expected purchase utility. The customer does not observe the "at-home-inspection" shock ψ_{ij} until after the purchase decision. However, the customer must make all click and purchase decisions using the available information and thus would calculate the expected purchase utility by integrating out the unobserved shock ψ_{ij}. In Appendix Appendix B, we demonstrate that under the assumptions discussed previously, the expected purchase utility takes a simplified form:

$$\omega_{ij} = \mathbb{E}_{\psi_{ij}}[u_{ij}|x_{ij}, \xi_{ij}, \epsilon_{ij}] = \sigma_{\psi_{ij}} \mathcal{T}\left(\frac{R + \mu_{ij} + \xi_{ij} + \epsilon_{ij}}{\sigma_{\psi_{ij}}}\right) - R \tag{3}$$

where $\mathcal{T}(\kappa) = \kappa\Phi(\kappa) + \phi(\kappa)$ and $\Phi(\kappa)$ and $\phi(\kappa)$ are the cumulative distribution and probability density functions of the standard normal distribution.

The quantity in Equation (3) represents the customer's expected utility from purchasing a given product, taking into account the possibility of returning the product in the future. This formulation highlights how the return option alters the structure of the purchase utility. In contrast, under the *No-returns Model*, the expected purchase utility simplifies to $\mu_{ij}+\xi_{ij}+\epsilon_{ij}$, which coincides with the final realized utility. If the variance of ψ_{ij} is close to zero, there would be almost no discrepancy between the look of the product on the website (online) and at home (offline), so the customer would make a purchase decision with (almost) full information and our *Joint Model* would converge to the *No-returns Model.* In contrast, when most of the customer's learning occurs after receiving the product (high variance of ψ_{ij}), we expect returns to occur and to affect the purchase utility.

Also, without the option to return (e.g., because returns are prohibitively expensive), we again see how the customer's optimal decision converges to that in the *No-returns Model*:

Lemma 1. *When the return cost satisfies $R \to \infty$, the Joint Model converges to the No-returns Model.*

We can rewrite Equation (3) as

$$\omega_{ij} = \mu_{ij} + \xi_{ij} + \epsilon_{ij} + \sigma_{\psi_{ij}}(\mathcal{T}(x) - x)\,, \tag{4}$$

where

$$x = \frac{R + \mu_{ij} + \xi_{ij} + \epsilon_{ij}}{\sigma_{\psi_{ij}}}, \qquad \mathcal{T}(x) = x\Phi(x) + \phi(x).$$

Letting $R \to +\infty$ implies $x \to +\infty$ for any finite $\mu_{ij} + \xi_{ij} + \epsilon_{ij}$. Consider the limit

$$\lim_{x\to+\infty} \big(x\Phi(x) + \phi(x) - x\big) = \lim_{x\to+\infty} x(\Phi(x) - 1) + \lim_{x\to+\infty} \phi(x) = 0, \tag{5}$$

using $x(1 - \Phi(x)) \to 0$ and $\phi(x) \to 0$ as $x \to +\infty$.

Substituting this into Equation (4) gives

$$\lim_{R\to+\infty} \omega_{ij} = \mu_{ij} + \xi_{ij} + \epsilon_{ij}. \tag{6}$$

Not having the option to return means that the customer's expected utility declines, so the customer is less likely to purchase, as we show next.

Lemma 2. *Having a return option strictly increases the customer's expected utility.*

From Equation (4), the incremental value of the return option is

$$g(x) = \sigma_{\psi_{ij}} (x\Phi(x) + \phi(x) - x)\,.$$

Differentiating yields

$$g'(x) = \sigma_{\psi_{ij}} (\Phi(x) - 1) \leq 0,$$

so $g(x)$ is weakly decreasing. At $x = 0$,

$$g(0) = \sigma_{\psi_{ij}} \phi(0) > 0.$$

Since $g(x)$ decreases monotonically and satisfies

$$\lim_{x\to+\infty} g(x) = 0,$$

it follows that $g(x) \geq 0$ for all x, with $g(x) > 0$ for all finite x.

Our results in this section are intuitive: returns are costly, but they increase the value of purchasing, by increasing the purchase utility. Also, we showed under what conditions our model converges back to the Weitzman model, providing internal validity.

4.3 Optimal Click Strategy when Return is an Option

According to assumptions of our model, ψ_{ij} is independent from other shocks and thus the decision to return the product j would not depend on the realization of ψ_{ik} for other products k. This implies that the return option does not change the assumption of the Weitzman search framework, but only changes the distribution of the reward. Therefore, we can use the standard structure of customer click rules (Ursu et al. 2025). However, the form of the click rules must be updated to account for the changed structure of the reward function. Conceptually, the continuation, stopping, and purchase rules would have the same structure based on reservation utilities, however, the functional form of these rules is different and an additional return decision rule is added. We summarize the revised decision rules below and provide more detailed equations in the next section. We provide derivations in Appendix Appendix C.

Selection rule. The customer choosing to click will select the product with the highest reservation utility z_{ij} derived from the system in Equation (7).

$$\begin{aligned} c_{ij} &= \sigma_{\psi_{ij}} \int_{\theta}^{\infty} \left[\mathcal{T}\left(\frac{R + \mu_{ij} + \xi_{ij} + t}{\sigma_{\psi_{ij}}} \right) - \mathcal{T}\left(\frac{R + \mu_{ij} + \xi_{ij} + \theta}{\sigma_{\psi_{ij}}} \right) \right] d\Phi_{\epsilon_{ij}}(t) \\ z_{ij} &= \sigma_{\psi_{ij}} \mathcal{T}\left(\frac{R + \mu_{ij} + \xi_{ij} + \theta}{\sigma_{\psi_{ij}}} \right) - R \end{aligned} \tag{7}$$

The second equation is a 1-to-1 mapping $\theta \to z_{ij}$, but to find z_{ij}, we must first solve the first implicit equation for θ. Intuitively, before the click, customers do not know the value of ϵ_{ij}; thus, their reservation utilities cannot depend on it. By computing the integral in Equation (7), customers estimate the expected difference between expected purchase utilities in Equation (3) with and without a new click. Importantly, in Appendix Appendix C we demonstrate that the reservation utility could be represented as a function of only three arguments $\mu_{ij} + R + \xi_{ij}$, $\sigma_{\psi_{ij}}$, and c_{ij} which substantially reduces the computational burden and makes the estimation feasible.

Stopping rule. The customer continues to click until the customer's maximal expected utility of clicked options from Equation (3) exceeds the maximal reservation utilities of not-clicked options in Equation (7). This stopping rule is conceptually similar to the standard search framework.

Purchase rule. When the stopping rule is reached, the customer purchases either the highest-expected utility product (taking into account return option) from the set of all clicked products or chooses the outside option.

Return rule. If a product (not the outside option) is purchased, the customer keeps (does not return) the product if the customer's utility for the chosen-and-inspected product is larger than the negative return costs, $-R$.

4.4 Likelihood

Let M_i denote the number of products presented to the customer (for example, the number of products shown on the list page). The customer clicks on C_i products from this set of products according to the optimal search rules discussed in the previous section. Recall that the index j represents the order in which the customer clicks on the product (e.g., $j = 2$ denotes the second clicked product, and $j = C_i$ denotes the last clicked product). This notation implies that the customer does not click on products with $j > C_i$, enabling us to enumerate the order of non-clicked products randomly.

Consider the customer who was presented with M_i products, clicked on C_i products, purchased a product with index b, and decided to return it. This sequence implies the following constraints, where $\mathcal{I}$[constraint] is the indicator function that takes on a value of 1 if the constraint is satisfied:

Click continuation. After clicking on option j, the customer continues clicking on other options if the value of searching is larger than the value of the best option in hand:

$$\forall j < C_i \quad \mathcal{I}\left[z_{i(j+1)} \geq \max_{s=j+2..M_i} z_{is}\right] \mathcal{I}\left[\max_{s=0..j} \omega_{is} < \max_{s=j+1..M_i} z_{is}\right] = 1 \tag{8}$$

Click stopping. The customer stops clicking when the maximum expected utility of clicked options is higher than the value of searching the remaining options:

$$\mathcal{I}\left[\max_{s=0..C_i} \omega_{is} \geq \max_{s=C_i+1..M_i} z_{is}\right] = 1 \tag{9}$$

Purchase. Given that the customer has clicked C_i products and decided to stop clicking, the customer purchases a product if the expected utility of the purchased product is greater than the expected utility of all other clicked products, including the outside option.

$$\mathcal{I}\left[\omega_{ib} \geq \max_{s=0..C_i} \omega_{is}\right] = 1 \tag{10}$$

Return. Given the customer bought product b, the customer returns the product if the product utility is lower than $-R$.

$$\mathcal{I}\left[\mu_{ib} + \xi_{ib} + \epsilon_{ib} + \psi_{ib} \leq -R\right] = 1 \tag{11}$$

Equations (8) to (11) define the set of constraints that must be satisfied to observe the given browsing session. Multiplication of the indicator functions for these conditions is the same as requiring all conditions to hold and produces a binary variable W_i that takes 1 if and only if all constraints are satisfied. The case when the customer decides to keep the product or chooses the outside options closely follows the derivations in Equations (8) to (11). In Appendix Appendix E, we demonstrate that the set of Equations (8) to (11) can be rewritten in the more compact form in Equation (12):

$$
\begin{aligned}
1 = W_i = &\left[\prod_{j=1}^{C_i-1} \mathcal{I}[z_{ij} \geq z_{i(j+1)}]\right] \mathcal{I}\left[\min\{\omega_{ib}, z_{iC_i}\} \geq \max_{s=C_i+1..M_i} z_{is}\right] \\
&\times \left[\prod_{j=0}^{C_i-1} \mathcal{I}[\omega_{ij} \leq \min\{z_{iC_i}, \omega_{ib}\}]\right] \mathcal{I}[\omega_{iC_i} \leq \omega_{ib}] \mathcal{I}[\mu_{ib} + \xi_{ib} + \epsilon_{ib} + \psi_{ib} \leq -R]
\end{aligned}
\tag{12}
$$

Because the researcher does not observe individual shocks $\xi_{ij}, \epsilon_{ij}, \psi_{ij}$ we obtain the probability of observing the given click sequence of customer i by integrating out these variables to determine the probability that all constraints are satisfied. This integration produces the log-likelihood function:

$$
LL(\beta) = \sum_{i=1}^{N} \log \int W(\xi_i, \epsilon_i, \psi_i, \beta)\, dF(\xi_i, \epsilon_i, \psi_i) \tag{13}
$$

where $F(\xi_i, \epsilon_i, \psi_i)$ represents the joint distribution of unobservable shocks given the parameter value β, and $\beta = (\beta^u, \beta^c, \beta^R, \beta^\psi)$ is the parameter vector to be estimated.

If computations were feasible, we could find the estimates of the parameters by maximizing the log-likelihood function in Equation (13). Unfortunately, Equation (13) highlights two complications prohibiting direct maximization.

First, to compute the reservation utility z_{ij}, we need to solve Equation (7), which includes integration. Because z_{ij} depends on the model's parameters and needs to be computed for each iteration of the optimization algorithm, it is infeasible to compute the exact value of the integral for each customer-product pair. This is a known issue with Weitzman-like click models; however, in the standard case, the integral could be replaced with a function of the form $c_{ij} = h(z_{ij})$ where h is an invertible function (Kim et al. 2010). In our case, the argument in the integral is a non-linear function, and there is no known closed form for the integral in Equation (13). To make the estimation feasible and find z_{ij}, we use a trilinear interpolation described in Appendix Appendix D. Intuitively, we compute the exact integral for a predefined fine grid of points and approximate the values between the grid points by a continuous function, thereby reducing the computation time.

Second, no known closed-form solution exists to the integral in Equation (13). The inner integral involves integration over $\xi_i, \epsilon_i, \psi_i$ that contains $M_i + C_i + 1$ variables (M_i shocks are realized before clicks, C_i shocks are realized from the set of clicked products, and one shock is realized from the purchased product). In our applications, we found that existing methods like accept-reject simulator (Chen and Yao 2017) and its smoothed version (Honka and Chintagunta 2017, Ursu 2018) could not be directly applied to our model. The main reason is that in many retail settings M_i is very large making the likelihood very close to 0, which requires a very large number of simulation variables, which makes the estimation time too large even for modern machines.

To reduce the number of required simulated variables, we recognize that some shocks could be integrated out from Equation (13) and need not be simulated. For example, from Equation (12), it follows that only the return constraint $\mathcal{I}[\mu_{ib}+\xi_{ib}+\epsilon_{ib}+\psi_{ib} \leq -R]$ contains the value of the shock ψ_{ib} and according to assumptions of the model, ψ_{ib} is independent of all other shocks. Therefore, the integral in likelihood Equation (13) could be equivalently replaced with

$$\int W(\xi_i, \epsilon_i, \psi_i, \beta)\, dF(\xi_i, \epsilon_i, \psi_i) = \int W_2(\xi_i, \epsilon_i, \beta)\Phi\left(\frac{-R-\mu_{ib}-\xi_{ib}-\epsilon_{ib}}{\sigma_{\psi_{ib}}}\right) dF(\xi_i, \epsilon_i) \quad (14)$$

where $W_2(\xi_i, \epsilon_i, \beta)$ and $F(\xi_i, \epsilon_i)$ do not depend on ψ_{ij}

Equation (14) only requires the simulation of ξ_i, ϵ_i. In Appendix Appendix F, we show that sufficiently many constraints can be integrated out for the inner integral in Equation (13). Only one remaining constraint needs to be replaced with a smoothed version. Partially closed-form integration substantially reduces the required number of draws for simulated variables and allows maximization with a gradient-based algorithm.

5 Model Parameters and Identification

Implementing the model requires specifying functional forms for its key components. In this section, we discuss the functional and distributional assumptions that ensure parameter identification and demonstrate using Monte Carlo simulation the feasibility of estimation using the proposed approach.

5.1 Model Parameters

Mean utility μ_{ij}. We model the mean utility of product j for customer i as a function of observed product characteristics x_{ij}. Customers weight these characteristics according to a preference vector β^u, so that $\mu_{ij} = x'_{ij}\beta^u$. The intercept in β^u captures the (negative) utility of the outside option. The specific variables included in x_{ij} are described in the empirical section.

Click costs c_{ij}. Click costs depend on the browsing environment; for example, clicking on a product at the top of the webpage may require less effort (Ursu 2018). Because j indexes the click order, we write $\log c_{ij} = position'_{ij}\beta^c$.

Return costs R. To ensure that returns costs are positive, we assume that $\log R = \beta^R$. We also assume that return costs are fixed and do not vary across products, as our estimation focuses on one category at a time and on single-item orders, i.e., we assume the customer will not find returning one pair of jeans to be more costly than another pair sold by the same retailer.

Post-purchase shock ψ_{ij}. This unobserved shock captures the difference between the information available to the customer on the website and the information learned once the product is received. We parameterize the log-variance of this unobserved shock as a linear function of product characteristics: $\log \sigma_\psi = x'_{ij}\beta^\psi$. For instance, a simple white t-shirt would rarely surprise a customer upon delivery (low σ_ψ), whereas an elegant evening dress could have many subtle features not fully captured by online images, leading to a high degree of uncertainty (high σ_ψ) that is only resolved post-purchase.

5.2 Parameter Identification

The parameters to be estimated include the customer preference vector β^u, the heterogeneity of unobserved fit β^ψ, search costs β^c, and return costs β^R. In our model, we assume the return cost is constant across all customers; therefore, it is identified from the variation in observed return decisions. Specifically, for two customers who examine an identical set of products and purchase the same item, the return cost $R = \exp(\beta^R)$ specifies the threshold for unobserved shocks as defined in Equation (11) and, thus, higher overall return probability would imply lower value of β^R.

Identification of β^u and β^ψ is derived from variation in customer product choices. Consider customers who click on the same set of products but select different items for purchase. Following Equation (10), a product attribute that appears in more frequent purchases suggests either a higher value for the corresponding preference parameter β^u or a higher variance in the unobserved preference fit β^ψ. However, if products with that same attribute are simultaneously returned less frequently, it implies that β^u is higher or β^ψ is lower. Thus, Equations (10) and (11) jointly identify the preference vector and the unobserved preference-fit heterogeneity. Consistent with Anderson et al. (2009), the intercept in β^ψ is not identified; we therefore normalize it to zero. The intercept in β^u is identified by the outside option, representing the customer's choice to exit the session without a purchase. Finally, search costs β^c are identified in a manner similar to prior work (see argument in Ursu et al. (2025)).

Overall, the model contains more constraints than necessary for identification; we therefore rely on a conservative set of assumptions. Future research may relax these assumptions to identify additional parameters. Notably, the return-related parameters rely heavily on the single return constraint in Equation (11), which only exists in sessions with a purchase. This implies that estimation precision depends not only on total session volume but also on the conversion rate and the observed variation in returns. This dependency is critical across different product categories. For example, in Table 1, while "Knitwear" represents a more popular search category than "Jeans," the substantially lower purchase rate for "Knitwear" suggests that the "Jeans" category may provide more reliable estimates for return-related parameters. Consequently, a category with higher conversion rates can achieve superior estimation precision even with a smaller total sample size.

The assumptions of the model allow us to identify all parameters. In the next section, we use Monte Carlo simulations to (1) confirm identification, (2) demonstrate the feasibility of estimation, and (3) compare the results to those from the *No-returns Model*.

5.3 Monte Carlo Simulations

To assess parameter recovery and highlight the implications of modeling the complete customer journey, we conduct Monte Carlo simulations using synthetic data. Specifically, we assume that customers behave according to the proposed framework and simulate the click, purchase, and return decisions of 10,000 synthetic customers. The simulation parameters are chosen such that the moments of the simulated data closely resemble those in the actual dataset from Table 1. Specifically, in our simulated data the average number of clicks is 2.3, purchase rate is 0.05 and return rate is 0.37.

We assume that the retailer offers an assortment of 100 products. Each product has two observable characteristics: price and a binary attribute (e.g., artificial vs. natural material). Utility depends on both of these characteristics, while the variance of the unobserved post-purchase preference fit, σ_ψ, depends only on the binary attribute.[4] Search costs vary by product position on the website, where lower positions are assumed to be closer to the top of the page and thus entail lower click costs. Each customer observes a random subset of 50 products.

We estimate both the *No-returns Model* and *Joint Model* on the synthetic data and compare the ability of each model to recover the "true" parameters. The results of the estimation are summarized in Table 5.

The estimates indicate that the proposed approach accurately recovers the true model parameters, making it suitable for empirical analysis. In contrast, the *No-returns Model*

[4] For the sake of simplicity we chose the attribute to be binary, however, the results do not depend on whether the attribute is binary or not

Table 5: Model Estimation based on Synthetic Data

	True	Estimated (Joint)	Estimated (No-returns)
Utility (μ)			
Price (β_1^u)	−0.8	−0.85 (0.03)	−0.85 (0.04)
Artificial fabric (β_2^u)	−0.1	−0.08 (0.04)	0.13 (0.01)
Constant (β_0^u)	−4.0	−3.96 (0.02)	−3.63 (0.03)
Click cost ($\log c$)			
Position (β_1^c)	0.7	0.70 (0.02)	0.70 (0.02)
Constant (β_0^c)	−9.0	−9.00 (0.06)	−8.53 (0.08)
Variance of unobserved fit ($\log \sigma_\psi$)			
Binary attribute (β_1^ψ)	0.3	0.29 (0.04)	–
Return cost ($\log R$)			
Constant (β^R)	−0.8	−0.77 (0.02)	–
Log-likelihood		-8.549	-8.653
Observations		10,000	10,000

Standard errors computed via bootstrapping with 10 repetitions.

delivers mixed performance, and, as expected, it cannot recover return costs or the variance of the unobserved fit.

Parameters correctly recovered by the No-returns Model. The coefficient β_1^u on price and position effects β_1^c are recovered precisely. This is not surprising, since the variance of the unobserved preference fit σ_ψ does not depend on price. From the structure of the model in Equation (3), it follows that the expected purchase utility is increasing in μ for fixed σ_ψ. Thus, if two products have different prices and fixed other attributes, their relative utility rankings by μ and by ω are identical. Therefore, the sensitivity to price would be inferred similarly in both models, and the coefficient values are not biased when using the *No-returns Model*. The same logic applies to position-related search cost effects.

Parameters estimated with moderate bias. For the constant terms (utility and search costs), the *No-returns Model* produces systematic bias: it overestimates the utility intercept (or equivalently underestimates the value of the outside option) and overestimates search costs. Intuitively, the *No-returns Model* attributes higher observed clicks (i.e., higher reservation utility) to higher baseline utility, rather than to lower search costs. This is intuitive because, according to Equation (3), the return option increases the customer's expected purchase utility; thus, a model that ignores it will overestimate customer utility. These biases shrink as return costs R increase. In the extreme case where returns are prohibitively costly ($R \to \infty$), the *Joint Model* converges to the *No-returns Model*, and returns no longer affect inference.

Parameters incorrectly inferred by the No-returns Model. The most striking bias arises for the preference to *Artificial Fabric* (β_2^u), where even the sign of the coefficient is estimated incorrectly. Because $\beta_1^\psi > 0$, products made with artificial materials are riskier: customers expect greater post-purchase variance/opportunity. They may nevertheless purchase these products, when they know they can return them at a fixed cost R. Consequently, such products generate relatively high sales and high return rates. The *No-returns Model*, which ignores post-purchase resolution of uncertainty, interprets the attribute as being positively valued and thus estimates $\beta_2^u > 0$. In contrast, the *Joint Model* uses return behavior to correctly infer that the attribute is not actually preferred by customers.

In practice, incorrect inference about the sign of a coefficient can have serious managerial implications. For example, a firm might mistakenly promote or design products with undesirable features if sales data alone are used to guide recommendations. A product with high sales but high return rates may look attractive in a model that ignores returns, whereas the *Joint Model* distinguishes between genuinely valued and "risky" attributes.

In summary, these simulations demonstrate that the proposed estimation procedure accurately recovers the true parameters and is well-suited for empirical applications. The simulations also underscore that modeling the complete customer journey matters: ignoring the return decision does more than add noise – it can lead to systematically incorrect inferences and misguide managerial decisions. In the next section, we estimate our model using retailer data and show that accounting for returns when ranking products can substantially reduce return rates.

6 Estimation Results

We estimate the *Joint Model* using the simulated likelihood approach discussed in Section 5.3. Specifically, we draw 100 sets of random shocks for each customer and repeat the estimation 10 times to compute bootstrap standard errors. To make the estimation feasible, we focus on the *Jeans* category by restricting our sample to customers browsing only products from this category. This choice is motivated by several factors. First, it is one of the three largest categories based on the number of browsing sessions, providing a robust dataset for analysis. Second, compared to the *Jackets & Coats* category, jeans are more homogeneous. This homogeneity simplifies the definition of product characteristics; for instance, *Jackets & Coats* includes highly diverse items such as winter parkas and spring windbreakers, where customer perceptions of attributes like color may vary significantly across sub-categories. Modeling such a heterogeneous category would require a substantially larger set of characteristics, thereby increasing computational burden. Lastly, Table 1 indicates that *Jeans* has a relatively high purchase rate, which facilitates a more precise identification of return-related parameters – the primary focus of this paper (see Section 5.2 for details). It should be noted that these justifications are not prohibitive; a retailer could estimate the model for any category or even pool all products using appropriate control variables – the model would still produce reliable estimates. The only meaningful constraints are the computational resources and the time required to achieve model convergence, especially when running 10 bootstrap repetitions.

Because the retailer's assortment is large (1,200 products even within the *Jeans* category), it is not feasible to estimate the structural model using product-specific fixed effects. Instead, we rely on automatically generated, interpretable characteristics to describe each product from Section 3. In total, the dataset contains 29,206 browsing sessions, of which 2,428 resulted in a transaction, with 896 items returned after purchase. Across these sessions, customers were exposed to 269,159 products and clicked on 12,811. To highlight the value of estimating the *Joint Model*, we also estimate the *No-returns Model*.

Table 6 reports the parameter estimates from the *No-returns* and *Joint Models*. For compactness, we split the estimated parameters of each model into two columns: the first column reports the estimates of the utility parameters, β^u, while the second column reports the estimates of the post-purchase preference fit variance β^ψ (note that the *No-returns Model* does not incorporate the return decision and thus cannot estimate the variance of the unobserved preference fit, and thus no values are reported in that column). For example, the estimate of 0.129 for *Proportion Natural Material* indicates that products with a higher share of natural materials yield higher mean utility, while the estimate of -0.650 shows that these products also have lower variance of preference fit (or less discrepancy between pre-purchase and post-purchase utility).

The results based on empirical data support the insights from our synthetic data exercise. In particular, the constant terms in both the utility and search cost equations differ substantially across models: the no-returns specification overestimates search costs by about 18% relative to the joint model[5]. This large discrepancy in the estimated cost could negatively

[5]Because the estimates represent the log-search costs, it follows that $\exp(-8.593 + 8.756) = \exp(0.163) \approx 1.18$.

Table 6: Model Estimation Based on the Empirical Data in the Jeans Category

	Joint Model		No Returns Model	
Mean Utility / Variance of Unobserved Fit	μ	σ_ψ	μ	σ_ψ
Price	-0.007^{***}	–	-0.007^{***}	–
	(0.000)		(0.000)	
Proportion Natural Material	0.129^{**}	-0.650^{***}	-0.196^{***}	–
	(0.054)	(0.078)	(0.023)	
Gender: Men	0.275^{***}	-0.328^{***}	0.081^{***}	–
	(0.029)	(0.030)	(0.010)	
Subcategory: Alexa	0.316^{***}	-0.378^{***}	0.045^{***}	–
	(0.023)	(0.041)	(0.007)	
Subcategory: Loose	-0.427^{***}	0.645^{***}	0.154^{***}	–
	(0.107)	(0.083)	(0.022)	
Subcategory: Skinny	0.056^{*}	-0.159^{***}	-0.094^{***}	–
	(0.029)	(0.055)	(0.006)	
Subcategory: Slim	0.100^{***}	-0.206^{***}	-0.067^{***}	–
	(0.023)	(0.042)	(0.007)	
Subcategory: Tapered	-0.047	-0.066	-0.140^{***}	–
	(0.082)	(0.174)	(0.022)	
Subcategory: Other	0.002	0.164^{***}	0.102^{***}	–
	(0.034)	(0.040)	(0.005)	
Color: Black	-0.031	-0.267^{***}	-0.141^{***}	–
	(0.023)	(0.040)	(0.008)	
Color: Blue	-0.014	-0.067	-0.066^{***}	–
	(0.032)	(0.055)	(0.005)	
Color: Gray	0.048^{***}	-0.406^{***}	-0.128^{***}	–
	(0.016)	(0.050)	(0.009)	
Color: Other	-0.086^{***}	0.228^{***}	-0.011^{***}	–
	(0.027)	(0.064)	(0.007)	
Type: Menplus	0.108^{***}	-0.405^{***}	-0.081^{***}	–
	(0.018)	(0.046)	(0.006)	
Asymmetry	-0.199^{*}	1.545^{***}	0.555^{***}	–
	(0.103)	(0.351)	(0.067)	
Brightness (mean)	-0.046^{**}	-0.031	-0.062^{***}	–
	(0.022)	(0.035)	(0.002)	
Brightness (variance)	0.022	0.021	0.032^{***}	–
	(0.014)	(0.023)	(0.001)	
Color asymmetry (horizontal)	-0.299^{***}	0.100	-0.166^{***}	–
	(0.092)	(0.126)	(0.033)	
Color asymmetry (vertical)	-0.250^{***}	$0.139*$	-0.116^{***}	–
	(0.061)	(0.082)	(0.008)	
Horizontal lines	-0.018^{***}	0.013^{***}	-0.010^{***}	–
	(0.003)	(0.005)	(0.001)	
Vertical lines	0.034^{*}	0.012	0.040^{***}	–
	(0.018)	(0.030)	(0.003)	
Complexity of pattern	-0.287^{***}	0.334^{***}	-0.063^{***}	–
	(0.034)	(0.077)	(0.016)	
Number of straight lines	-0.005^{***}	0.001	-0.004^{***}	–
	(0.001)	(0.002)	(0.000)	
Constant	-4.498^{***}	–	-4.312^{***}	–
	(0.010)		(0.014)	
Search Cost ($\log c$)				
Position	0.982^{***}		1.013^{***}	
	(0.008)		(0.009)	
Constant	-8.593^{***}		-8.756^{***}	
	(0.033)		(0.052)	
Return Cost ($\log R$)				
Constant	-0.621^{***}		–	
	(0.027)			
Log-likelihood	-11.310		-11.617	
Observations	29,206		29,206	

$^{***}\ p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

impact retailers' understanding of the value of a click on the website.

More strikingly, the *No-returns Model* often produces substantially different preference estimates, sometimes even reversing signs. For example, the *No-returns Model* suggests that customers dislike natural materials in jeans (-0.196), which may lead managers to shift the assortment toward synthetic products and underprice products made from natural materials. In reality, the *Joint Model* shows that customers prefer jeans made of natural materials ($+0.129$). The key is that natural materials also reduce the variance of unobserved fit (σ_ψ), meaning that customers expect fewer surprises and are less likely to return such products after purchase. The *No-returns Model* cannot capture this mechanism and therefore misinterprets the value of artificial fabrics. Similar logic applies to other characteristics.

7 Counterfactuals: Mitigating Return Rates Through Optimized Product Rankings

In the previous section, we demonstrate the value of the model for studying product returns. Although these findings can help the retailer better infer customer preferences, they do not provide a direct managerial application. To illustrate the practical value of our model, we consider the product-ranking problem faced by the retailer. As shown in Table 6, customers are substantially more likely to click on products displayed at the top of the page. This suggests that the retailer can strategically use these high-visibility positions to affect both purchases and returns.

The optimal ranking problem is extremely complex. For example, if the retailer displays 10 products on a page, then there are $10! \approx 3 \cdot 10^6$ possible orderings. Given that most retailers have substantially more products in a category,[6] exhaustively evaluating all permutations is computationally infeasible. As a result, most retailers rely on *score-based* ranking rules, where products are sorted according to a scalar measure of "usefulness," which in turn typically depends on that product's characteristics, click through rates, and purchase rates (Ursu 2018).

What we want to show is that a model that does not take into account return decisions may lead to a ranking algorithm that may inadvertently promote products with high return rates, which can significantly reduce retailer overall profits. In this section, we use our *Joint Model* - which jointly captures search, purchase, and returns - to evaluate how different ranking strategies influence retailer welfare from both the sales and return perspectives. Specifically, we treat the *Joint Model* as the data-generating process and compare six ranking

[6] For example, our retailer has 48 products on each page (see Figure A.1) and 1,200 in the *Jeans* category alone.

rules, consistent with prior work (Ursu 2018, Ghose et al. 2012, 2014):

- *Default* – product positions remain unchanged from the retailer's original ordering.
- *Joint Model* – products are ranked according to the mean utility μ_{ij}^{F} estimated from the *Joint Model*.
- *No-returns Model* – products are ranked according to the mean utility μ_{ij}^{NR} estimated from the *No-returns Model*.
- *Click Rate* – products are ranked according to their historical click through rates.
- *Purchase Rate* – products are ranked according to their historical purchase rates.
- *Return Rate* – products are ranked according to their historical return rates.

Our ultimate goal is to mitigate product returns; however, examining return rates alone is insufficient. Product returns occur only conditional on the customer making a purchase; therefore, all changes to the website must be evaluated from the perspective how these changes impact both returns and purchases. A trivial way to reduce returns to zero is to reduce sales to zero – an obviously undesirable policy. The higher sales and higher returns tradeoff could vary across product categories. For example, if a returned product has a very low resale value, the retailer may prefer lower sales with substantially lower returns.

We do not have access to the cost structure of the retailer. Instead, to illustrate the potential of the *Joint Model* we use the industry average to approximate the impact on the profit. We assume the retailer's profit margin is 50% of the original price for kept items and that return processing costs equal 30% of price (Shopify 2025, CBRE Research 2023). During our counterfactual simulations for each customer, we simulate 1,000 different sets of random shocks and record their decisions according to our estimated model. We report average outcomes (profit improvement, purchase, and return rate) across these repetitions and customers.

The counterfactual simulation results are summarized in Table 7, where we report the relative (%) changes in comparison to the *Default Ranking*. Interestingly, all ranking methods that are not based on a model (those based on click through rates, purchase rates, or return rates) produce worse outcomes than those based on a model (either the *No-returns Model* and the *Joint Model*). This is not surprising: the *No-returns Model* takes into account both clicks and purchase decisions and therefore accurately captures characteristics of popular products, allowing customers to find these items more easily when they are placed higher on the page. However, because the *No-returns Model* overlooks the fact that many popular products are also more likely to be returned, the return rate increases substantially. As a result, the net profit improvement remains moderate. The *return-rate* ranking performs the worst and does not meaningfully reduce the return rate for the retailer. This result is not surprising given the large assortment size. A return can be observed only if a product is purchased, meaning that product-level return-rate estimates are extremely noisy (many products sell only 3–5 units per month). This noise makes the *return-rate* ranking perform nearly as poorly as a random ranking and, therefore, is unable to improve profits.[7]

Overall, the *Joint Model* ranking reduces return rates without sacrificing purchases much, generating substantially higher profit gains.

Table 7: Comparison of Model Predictions (Counterfactual Rankings)

	Joint Model	No-returns Model	Click Rate	Purchase Rate	Return Rate
Purchase Prob. Change	-0.05%	+4.64%	+2.31%	+0.88%	-0.10%
Return Prob. Change	-3.44%	+2.79%	+3.44%	+1.04%	+0.09%
Estimated Profit Change	+2.69%	+1.16%	+0.75%	-0.41%	-0.25%

These results highlight two important points. First, incorrect inference of customer preferences can lead to suboptimal product rankings and missed profit opportunities. Second, a simple, low-cost change to the website can substantially reduce return rates without hurting sales significantly. Industry reports emphasize that returns are costly (Reuters 2023), and even marginal improvements can translate into sizable profit gains. While different retailers face varying cost structures, having a more-complete *Joint Model* enables them to make more informed and ultimately more profitable decisions.

8 Conclusions and Future Research

Product returns pose a significant financial challenge for e-commerce companies, incurring substantial costs through refunds, complex reverse logistics, and high processing overhead. This paper establishes an empirical link between customer search behavior and the subsequent decision to return a product. Our evidence indicates that decisions by customers during the search phase are connected to return outcomes, necessitating a joint modeling approach.

We introduce a joint model that extends the common empirical Weitzman search framework to simultaneously model search, purchase, and return decisions. We demonstrate that this proposed model can be efficiently estimated using standard datasets available to most online retailers. Critically, we find that neglecting to account for returns can severely bias consumer preference estimates – even reversing the sign of crucial relationships – potentially leading to sub-optimal management and reduced profitability.

Finally, we highlight the practical application of our framework in the context of product ranking algorithms. By leveraging rankings built on estimates from our model, retailers can

[7]The performance of the purely random ranking model was not substantially different from *Default ranking*, and we excluded it from Table 7 for compactness.

effectively manage return rates and increase overall profits. Given that these optimizations are not expensive to implement, our model provides an economical and effective tool for mitigating the impact of frequent product returns.

There are several promising directions for future research. First, this study focused exclusively on sessions involving at most one purchased product. In practice, customers frequently purchase multiple items. While retailers could circumvent this constraint by duplicating sessions – treating a multi-item purchase as two separate "twin" sessions – this approach is a simplification. It ignores the strategic relationships within bundled purchases, such as a customer ordering two similar T-shirts with the intent to keep only one. Incorporating these strategic, multi-item behaviors into a rational choice model represents a high-value direction for both academic theory and practical application.

Second, data limitations prevented the estimation of a heterogeneous model. For retailers with extensive purchase histories and longitudinal data across multiple sessions, estimating personalized parameter sets for individual customers is a feasible next step. Such models would allow retailers to significantly increase profitability through highly tailored, personalized rankings for existing customers.

Lastly, it would be empirically valuable to explore alternative manipulations of the search environment and their impact on product returns. For example, the retailer could consider changing the granularity of filters, the criteria for product sorting, the logic of recommendation engines, or the introduction of AI-enhanced customer experiences. Understanding the effect of such changes on return probability offers a direct path to optimizing both customer satisfaction and retailer profitability.

References

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

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

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

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

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

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

CBRE Research (2023). Reverse logistics revs up as 2023 holiday sales rise.

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

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

Chick, S. E. and Frazier, P. (2012). Sequential sampling with economics of selection procedures. Management Science, 58(3):550–569.

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

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

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

Dzyabura, D., El Kihal, S., Hauser, J. R., and Ibragimov, M. (2023). Leveraging the power of images in managing product return rates. Marketing Science, 42(6):1125–1142.

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

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

El Kihal, S. and Shehu, E. (2022). 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.

Gardete, P. M. and Hunter, M. (2024). Multiattribute search: Empirical evidence and information design. Marketing Science, 43(5):1052–1080.

Ghose, A., Ipeirotis, P. G., and Li, B. (2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. Marketing Science, 31(3):493–520.

Ghose, A., Ipeirotis, P. G., and Li, B. (2014). Examining the impact of ranking on consumer behavior and search engine revenue. Management Science, 60(7):1632–1654.

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

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

Honka, E., Seiler, S., and Ursu, R. (2024). Consumer search: What can we learn from pre-purchase data? Journal of Retailing, 100(1):114–129.

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

Janssen, M. and Williams, C. (2024). Consumer search and product returns in e-commerce. American Economic Journal: Microeconomics, 16(2):387–419.

Jerath, K. and Ren, Q. (2024). Consumer search and product returns. Marketing Science.

Jiang, Z., Chan, T., Che, H., and Wang, Y. (2021). Consumer search and purchase: An empirical investigation of retargeting based on consumer online behaviors. Marketing Science, 40(2):219–240.

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

Ke, T. T. and Villas-Boas, J. M. (2019). Optimal learning before choice. Journal of Economic Theory, 180(C):383–437.

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

Lee, Jinjoo (2022). Retailers' many unhappy returns. The Wall Street Journal.

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

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

National Retail Federation and Happy Returns (2025). 2025 retail returns landscape. Technical report, National Retail Federation.

Noor Nanji (2022). Zara starts charging shoppers for online returns. Accessed on Day Month Year.

Owen, David (2023). hat Happens to All the Stuff We Return? The New Yorker.

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

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

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

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

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

Shopify (2025). Ecommerce returns: Statistics, trends, and strategy for 2025.

Stigler, G. J. (1961). The economics of information. Journal of Political Economy, 69(3):213–225.

Ursu, R., Seiler, S., and Honka, E. (2025). The sequential search model: A framework for empirical research. Quantitative Marketing and Economics, 23(1):165–213.

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

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

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

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

Appendix A   Data Cleaning and Retailer's Website

Appendix A.1   Data Cleaning

This section details the preprocessing pipeline used to arrive at our final estimation sample. For implementation details, please refer to the accompanying Python scripts. We applied the following filters to the raw session data:

- Remove sessions that started after February 28th, 2020, to account for the potential impact of COVID-related restrictions taking place in the country of the retailer.
- Remove sessions that are shorter than 5 seconds and longer than 4 hours.
- Remove sessions without product page views. This could happen if the customer arrives on the website from a third-party site and lands directly on the product page. These sessions do not reflect the true search process, and we are unable to recover the set of products from which the customer was choosing.
- Remove sessions where customer used sorting tool on the website (for example, by price). This customer action changes the product's order and could bias estimates of search costs.
- 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 retailer's focus.

Appendix A.2   Retailer's Website

The customer browsing the retailer's website during our observational period would observe the information in Figure A.1. By 2024, our retailer had updated its website design. To provide the most accurate information, we use the Wayback Machine (https://web.archive.org/) to capture the website's version as of 2020. Unfortunately, some product pictures were not stored by the platform.

Appendix B   Derivation of Expected Purchase Utility

In Section 4.2, we introduced the expected purchase utility, which represents the utility a customer expects to receive after purchasing a product. Because the purchase decision is made only after clicking on the product page, the customer observes x, ξ, and ϵ; the only remaining unknown component is ψ.[8] Hence, the customer would compute the expectation in Equation (B.1)

$$\omega(x,\xi,\epsilon) = \mathbb{E}_\psi\left[(\mu(x)+\xi+\epsilon+\psi+R)\cdot\mathcal{I}(\mu(x)+\xi+\epsilon+\psi\geq 0)+\mathcal{I}(\mu(x)+\xi+\epsilon+\psi\leq 0)\cdot(-R)\mid x,\xi,\epsilon\right] \tag{B.1}$$

where x is set of observable product characteristics and $\mu(x)$ determenistic known function of average utility given product characteristics x.

The expression in Equation (B.1) could be simplified:

$$\omega = \mathbb{E}_\psi\left[(\mu(x)+\xi+\epsilon+\psi+R)\cdot\mathcal{I}(\mu(x)+\xi+\epsilon+\psi+R\geq 0)-R\mid x,\xi,\epsilon\right] = \mathbb{E}_\psi\left[\zeta\cdot\mathcal{I}(\zeta\geq 0)\mid x,\xi,\epsilon\right]-R \tag{B.2}$$

where $\zeta = \mu(x)+\xi+\epsilon+\psi+R$

Notice that $\zeta \mid x,\xi,\epsilon \sim \mathcal{N}\left(\mu(x)+\xi+\epsilon+R;\sigma_\psi(x)\right)$ because $\psi \mid x \sim \mathcal{N}\left(0;\sigma_\psi(x)\right)$ and it is independent of all other shocks. The expectation has a known expression (see truncated normal distribution) and thus:

$$\omega(x,\xi,\epsilon) = \sigma_\psi(x)\cdot\mathcal{T}\left(\frac{\mu(x)+\xi+\epsilon+R}{\sigma_\psi(x)}\right)-R \tag{B.3}$$

where $\mathcal{T}(\kappa) = \kappa\Phi(\kappa)+\varphi(\kappa)$ with $\Phi(\kappa)$ and $\varphi(\kappa)$ being CDF and PDF of standard normal distribution respectively.

[8] Without loss of generality, we drop all indices and subscripts in this section to preserve readability.

Figure A.1: Screenshot of Retailer's Website

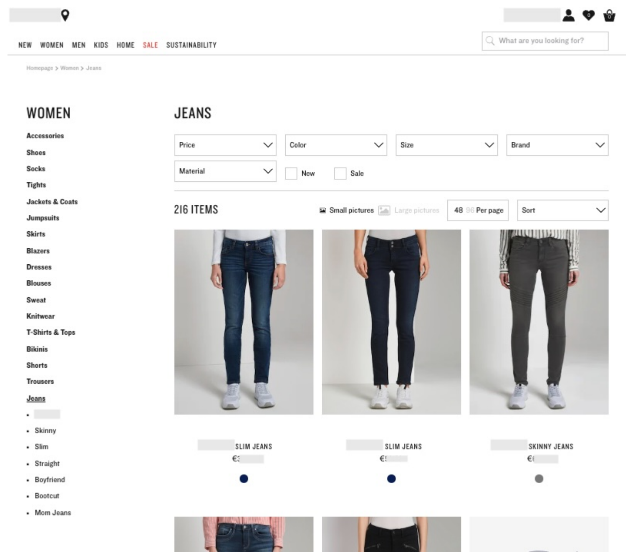

Appendix C Derivation of Reservation Utilities for Model with Product Returns

In the original paper, Weitzman (1979) demonstrated that the reservation utility z for a product could be found from Equation (C.1) where we drop the individual i and product j indices for compactness:

$$c = \int_{z}^{\infty} (u - z) dF(u) \tag{C.1}$$

where c is the search cost; $F(u)$ is the distribution of the reward after opening the box (or clicking on product page).

To adapt the expression from Equation (C.1) we need to find the distribution of the reward for clicking on the product page. Recall that before clicking on the product page the customer observes for all products: characteristics x and pre-click shock ξ, and the click reveals shock ϵ. Because post-purchase shock ψ is independent between products we can integrate it out for each product and obtain the expected purchase utility $\omega(x, \xi, \epsilon)$ from Equation (B.3) for each product.

Notice that $\omega(x, \xi, \epsilon)$ depends on observed to customer variable x, ξ and unobserved shock ϵ. Therefore, the customer knowledge about the product would be described by the distribution of $\omega(x, \xi, \epsilon)$ where $\epsilon \sim \mathcal{N}(0, 1)$. Thus, we can find the CDF for the distribution of the reward ω as:

$$\begin{aligned}
F(u) &= \mathbb{P}[\omega(x, \xi, \epsilon) \leq u \mid x, \xi] \\
&= \mathbb{P}\left[\sigma_\psi(x) \cdot \mathcal{T}\left(\frac{\mu(x) + \xi + \epsilon + R}{\sigma_\psi(x)}\right) - R \leq u \mid x, \xi\right] \\
&= \mathbb{P}\left[\mathcal{T}\left(\frac{\mu(x) + \xi + \epsilon + R}{\sigma_\psi(x)}\right) \leq \frac{u + R}{\sigma_\psi(x)} \mid x, \xi\right] \\
&= \mathbb{P}\left[\mu(x) + \xi + \epsilon + R \leq \sigma_\psi(x)\mathcal{T}^{-1}\left(\frac{u + R}{\sigma_\psi(x)}\right) \mid x, \xi\right] \\
&= \Phi\left[\sigma_\psi(x)\mathcal{T}^{-1}\left(\frac{R + u}{\sigma_\psi(x)}\right) - \mu(x) - \xi - R\right]
\end{aligned} \tag{C.2}$$

where assumption that $\epsilon \sim \mathcal{N}(0; 1)$ was used.

Formally, we can substitute the distribution into Equation (C.1) and find reservation utilities z. However, the expression could be simplified if we use the following substitutions:

$$
\begin{aligned}
u &= s_1(t) = \sigma_\psi(x)\mathcal{T}\left(\frac{\mu(x)+\xi+R+t}{\sigma_\psi(x)}\right) - R \\
z &= s_2(\theta) = \sigma_\psi(x)\mathcal{T}\left(\frac{\mu(x)+\xi+R+\theta}{\sigma_\psi(x)}\right) - R
\end{aligned} \tag{C.3}
$$

Substitution simplifies the CDF from Equation (C.2)

$$
\begin{aligned}
F(u) = F(s_1(t)) &= \Phi\left[\sigma_\psi(x)\mathcal{T}^{-1}\left(\frac{R+s_1(t)}{\sigma_\psi(x)}\right) - \mu(x) - \xi - R\right] \\
&= \Phi\left[\sigma_\psi(x)\mathcal{T}^{-1}\left(\frac{R+\sigma_\psi(x)\mathcal{T}\left(\frac{\mu(x)+\xi+R+t}{\sigma_\psi(x)}\right) - R}{\sigma_\psi(x)}\right) - \mu(x) - \xi - R\right] \\
&= \Phi(t)
\end{aligned} \tag{C.4}
$$

Therefore, we can rewrite the original integral as:

$$
\begin{aligned}
c &= \int_z^\infty (u-z)dF(u) = \int_z^\infty (u-z)f(u)du \\
&= \int_{s_1^{-1}(z)}^\infty (s_1(t)-z)f(s_1(t))ds_1(t) = \int_{s_1^{-1}(z)}^\infty (s_1(t)-z)dF(s_1(t)) \\
&= \int_{s_1^{-1}(z)}^\infty (s_1(t)-z)d\Phi(t) \\
&= \int_{s_1^{-1}(s_2(\theta))}^\infty (s_1(t)-s_2(\theta))d\Phi(t) \\
&= \sigma_\psi(x)\int_\theta^\infty \left[\mathcal{T}\left(\frac{\mu(x)+\xi+R+t}{\sigma_\psi(x)}\right) - \mathcal{T}\left(\frac{\mu(x)+\xi+R+\theta}{\sigma_\psi(x)}\right)\right] d\Phi(t)
\end{aligned} \tag{C.5}
$$

In Equation (C.5), all variables c, x, ξ are known to the customer for each product before click, and thus this equation defines the $\theta^*(c, x, R, \xi)$ as a solution. Finally, to find reservation utility z^* we simply substitute θ^* to Equation (C.3):

$$
z^*(c, x, R, \xi) = \sigma_\psi(x)\mathcal{T}\left(\frac{\mu(x)+\xi+R+\theta^*(c,x,R,\xi)}{\sigma_\psi(x)}\right) - R \tag{C.6}
$$

The dimension of product characteristics x could be high, which makes the reservation utility less tractable. However, because $\mu(x)$ and $\sigma_\psi(x)$ are deterministic function the reservation utility could be re-parameterized as $\theta^*(c,x,R,\xi) = \theta^*(\mu(x)+\xi+R, \sigma_\psi(x), c) \Rightarrow z^*(c,x,R,\xi) = z^*(\mu(x)+\xi+R, \sigma_\psi(x), c)$. Because $\mu(x)+\xi+R, \sigma_\psi(x), c$ are all scalar values, the reservation utility could be represented as a function of only three arguments, which substantially simplifies computational burden and approximations.

Appendix D Approximating the Solution to the Equation

In Appendix Appendix C, we derived the equation for reservation utilities. Unfortunately, no known closed-form solution exists for this problem; therefore, numerical approximation is required. Because a typical retailer deals with thousands of sessions and hundreds of viewed products per session, solving the equation explicitly for each set of realized values within the SMLE algorithm would be computationally infeasible. Luckily, we showed that it is possible to represent the reservation utility as a function of three scalar variables: $z^* = z^*(\mu(x)+\xi+R, \sigma_\psi(x), c) = f(x_1, x_2, x_3)$, where we replaced the parameters with x_1, x_2, x_3 for compactness. Therefore, we can precompute the reservation utility for a fixed set of x_1, x_2, x_3 and then approximate the remaining values using interpolation techniques.

In the paper, we utilize a trilinear interpolation technique. Specifically, for three-dimensional variables (x_1, x_2, x_3), we 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, we know the exact values of reservation utilities in eight vertices. For any vector within the cube, we approximate the reservation utility function $f(x_1, x_2, x_3)$ as:

$$\begin{aligned} f_{\text{true}}(x_1, x_2, x_3) &\simeq f_{\text{approx}}(x_1, x_2, x_3) \\ &\simeq \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_3 + \alpha_4 x_1 x_2 + \alpha_5 x_2 x_3 + \alpha_6 x_1 x_3 + \alpha_7 x_1 x_2 x_3 \end{aligned} \tag{D.1}$$

where we require $f_{\text{true}}(x_1, x_2, x_3) = f_{\text{approx}}(x_1, x_2, x_3)$ at the grid (or cube vertices) points. Because $f_{\text{approx}}(x_1, x_2, x_3)$ has eight parameters and eight constraints, the linear system has a unique solution for each cell. Notice that f_{approx} needs to be computed only once; thus, it avoids scaling the computational time with the number of iterations in the SMLE algorithm.

Appendix E Derivation of Equivalent Set of Constraints on Model Parameters

In Section 4.4 we derived a set of constraints on model parameters. However, simply combining all sets of constraints is complicated by a large number of min and max operators, which generate redundant constraints. For example, one such constraint $\mathcal{I}\left[z_1 \geq \max_{s=2..M} z_s\right]$ is equivalent to $M-2$ constraints of the form $\mathcal{I}\left[z_1 \geq z_s\right]$.

Recall that C is the number of clicked products and M is the number of products presented to the customer. The index j implies the order of clicked products, and $j = 0$ is the outside option. Index b represents a purchased product. For compactness and without loss of generality, we drop the customer index i. We consider the case when customer purchases some product b and returns it – derivation to other cases closely repeat steps outlined in this section.

Click Continuation. Equation (8) could be rewritten as:

$$\begin{aligned} 1 &= \prod_{j=0}^{C-1} \left[\mathcal{I}\left[\max_{s=0..j} \omega_s \leq \max_{s=j+1..M} z_s \right] \mathcal{I}\left[z_{j+1} \geq \max_{s=j+2..M} z_s \right] \right] \\ &= \left[\prod_{j=0}^{C-1} \mathcal{I}\left[\max_{s=0..j} \omega_s \leq \max_{s=j+1..M} z_s \right] \right] \left[\prod_{j=0}^{C-1} \mathcal{I}\left[z_{j+1} \geq \max_{s=j+2..M} z_s \right] \right] \\ &= \left[P_1\right]\left[P_2\right] \end{aligned} \tag{E.1}$$

Consider the second part P_2 of the Equation (E.4):

$$\begin{aligned} P_2 &= \prod_{j=0}^{C-1} \mathcal{I}\left[z_{j+1} \geq \max_{s=j+2..M} z_s \right] = \prod_{j=0}^{C-1} \prod_{s=j+2..M} \mathcal{I}\left[z_{j+1} \geq z_s\right] \\ &= \mathcal{I}\left[z_C \geq \max_{s=C+1..M} z_s \right] \prod_{j=1}^{C-1} \mathcal{I}\left[z_j \geq z_{j+1}\right] \end{aligned} \tag{E.2}$$

Notice that Equation (E.2) implies that the lower index j corresponds to larger reservation utility z_j for clicked options. This is intuitive, as according to the Weitzman model, the customer would click on options in the descending order of their reservation utility. Formally, we have that $\forall j \leq C : \max_{s=j+1..M} z_s = z_{j+1}$ and the first part P_1 of the Equation (E.4) could be simplified to:

$$\begin{aligned} P_1 &= \prod_{j=0}^{C-1} \mathcal{I}\left[\max_{s=0..j} \omega_s \leq \max_{s=j+1..M} z_s \right] = \prod_{j=0}^{C-1} \mathcal{I}\left[\max_{s=0..j} \omega_s \leq z_{j+1} \right] \\ &= \prod_{j=0}^{C-1} \left[\prod_{s=0..j} \mathcal{I}\left[\omega_s \leq z_{j+1}\right] \right] = \prod_{j=0}^{C-1} \mathcal{I}\left[\omega_j \leq z_C\right] \end{aligned} \tag{E.3}$$

Finally, click continuation constraint in Equation (E.4) could be replaced with:

$$1 = \mathcal{I}\left[z_C \geq \max_{s=C+1..M} z_s\right]\left[\prod_{j=1}^{C-1}\mathcal{I}\left[z_j \geq z_{j+1}\right]\right]\left[\prod_{j=0}^{C-1}\mathcal{I}\left[\omega_j \leq z_C\right]\right] \tag{E.4}$$

Click Stopping and Purchase. It is convenient to combine Equations (9) and (10) into one:

$$\begin{aligned}
1 &= \mathcal{I}\left[\max_{s=0..C}\omega_s \geq \max_{s=C+1..M} z_s\right]\mathcal{I}\left[\omega_b \geq \max_{s=0..C}\omega_s\right] \\
&= \mathcal{I}\left[\max_{s=0..C}\omega_s \geq \max_{s=C+1..M} z_s\right]\left[\prod_{j=0}^{C}\mathcal{I}\left[\omega_b \geq \omega_j\right]\right] \\
&= \mathcal{I}\left[\omega_b \geq \max_{s=C+1..M} z_s\right]\left[\prod_{j=0}^{C}\mathcal{I}\left[\omega_b \geq \omega_j\right]\right]
\end{aligned} \tag{E.5}$$

Finally, we can derive Equation (12) by combining Equations (E.4) and (E.5) and return constraint in Equation (11):

$$\begin{aligned}
1 &= \mathcal{I}\left[z_C \geq \max_{s=C+1..M} z_s\right]\left[\prod_{j=1}^{C-1}\mathcal{I}\left[z_j \geq z_{j+1}\right]\right]\left[\prod_{j=0}^{C-1}\mathcal{I}\left[\omega_j \leq z_C\right]\right] \\
&\times \mathcal{I}\left[\omega_b \geq \max_{s=C+1..M} z_s\right]\left[\prod_{j=0}^{C}\mathcal{I}\left[\omega_b \geq \omega_j\right]\right] \\
&\times \mathcal{I}\left[\mu_b + \xi_b + \epsilon_b + \psi_b \leq -R\right] \\
&= \left[\prod_{j=1}^{C-1}\mathcal{I}\left[z_j \geq z_{j+1}\right]\right]\mathcal{I}\left[\mu_b + \xi_b + \epsilon_b + \psi_b \leq -R\right] \\
&\times \mathcal{I}\left[z_C \geq \max_{s=C+1..M} z_s\right]\mathcal{I}\left[\omega_b \geq \max_{s=C+1..M} z_s\right] \\
&\times \left[\prod_{j=0}^{C-1}\mathcal{I}\left[\omega_j \leq z_C\right]\right]\left[\prod_{j=0}^{C}\mathcal{I}\left[\omega_b \geq \omega_j\right]\right] \\
&= \left[\prod_{j=1}^{C-1}\mathcal{I}\left[z_j \geq z_{j+1}\right]\right]\mathcal{I}\left[\mu_b + \xi_b + \epsilon_b + \psi_b \leq -R\right] \\
&\times \mathcal{I}\left[\min\{z_C, \omega_b\} \geq \max_{s=C+1..M} z_s\right] \\
&\times \left[\prod_{j=0}^{C-1}\mathcal{I}\left[\omega_j \leq \min\{z_C, \omega_b\}\right]\right]\mathcal{I}\left[\omega_C \leq \omega_b\right]
\end{aligned} \tag{E.6}$$

Finally, we obtained the simplified form of the set of constraints that must be satisfied in order to sequence of customer actions exist:

$$\begin{aligned}
1 = &\left[\prod_{j=1}^{C-1}\mathcal{I}[z_j \geq z_{j+1}]\right]\mathcal{I}\left[\min\{\omega_b, z_C\} \geq \max_{s=C+1..M} z_s\right] \\
\times &\left[\prod_{j=0}^{C-1}\mathcal{I}[\omega_j \leq \min\{z_C, \omega_b\}]\right]\mathcal{I}[\omega_C \leq \omega_b]\mathcal{I}[\mu_b + \xi_b + \epsilon_b + \psi_b \leq -R]
\end{aligned} \tag{E.7}$$

The derivation for cases, when a customer purchases a product and keeps it as well as when the customer not purchasing anything are straightforward repetitions of the steps outlined in this section.

Appendix F Derivation of Semi-Closed Form Likelihood

Notice that we can define a variable $W(\xi,\epsilon,\psi)$ that represents whether the set of parameters β and shocks ξ,ϵ,ψ satisfy as set of constraints in Equation (E.7):

$$\begin{aligned} W(\xi,\epsilon,\psi) = &\left[\prod_{j=1}^{C-1} \mathcal{I}[z_j \geq z_{j+1}]\right] \mathcal{I}\left[\min\{\omega_{ib}, z_C\} \geq \max_{s=C+1..M_i} z_s\right] \\ &\times \left[\prod_{j=0}^{C-1} \mathcal{I}[\omega_j \leq \min\{z_C, \omega_b\}]\right] \mathcal{I}[\omega_C \leq \omega_b]\mathcal{I}[\mu_b + \xi_b + \epsilon_b + \psi_b \leq -R] \end{aligned} \tag{F.1}$$

where ω_j is a function of unobserved to researcher shocks ξ_j and ϵ_j; z_j is a function of unobserved to researcher shock ξ_j.

To obtain the likelihood function, one would integrate out all unobserved shocks from the variable $W(\xi,\epsilon,\psi)$

$$\begin{aligned} \int \cdots \int W(\xi,\epsilon,\psi) dF(\xi,\epsilon,\psi) &= \int \cdots \int W(\xi,\epsilon,\psi) \left[\prod_{j=1}^{M} dF_{\xi_j}(\xi_j)\right] \left[\prod_{j=1}^{C} dF_{\epsilon_j}(\epsilon_j)\right] dF_{\psi_b}(\psi_b) \\ &= \int \cdots \int W(\xi,\epsilon,\psi) \left[\prod_{j=1}^{M} d\Phi(\xi_j)\right] \left[\prod_{j=1}^{C} d\Phi(\epsilon_j)\right] d\Phi\left(\frac{\psi_b}{\sigma_{\psi_b}}\right) \end{aligned} \tag{F.2}$$

where we used the assumption that all shocks are independent and normally distributed

One approach to approximating the integral in Equation (F.2) is to employ a Monte Carlo simulator: generating multiple sets of independent random shocks (ξ,ϵ,ψ), computing the binary indicator $W(\xi,\epsilon,\psi)$ for each, and averaging the results. However, this method is computationally prohibitive; the number of shocks $(M+C+1)$ can be very large in online retail contexts, therefore, the indicator function is often zero in high-dimensional spaces. A simulator would require us to sample hundreds of thousands of draws. To make estimation feasible, our goal is two-fold: (1) reduce the number of shocks that must be sampled, and (2) simplify Equation (F.2) to require fewer realizations for a stable probability estimate.

Elimination of return related constraint. Notice that only the last inequality in Equation (F.1) depends on ψ_b and we can replace the integral with the expression. Therefore, Equation (F.2) could be represented as:

$$\begin{aligned} LL &= \int \cdots \int W(\xi,\epsilon,\psi) \left[\prod_{j=1}^{M} d\Phi(\xi_j)\right] \left[\prod_{j=1}^{C} d\Phi(\epsilon_j)\right] d\Phi\left(\frac{\psi_b}{\sigma_{\psi_b}}\right) \\ &= \int \cdots \int W_1(\xi,\epsilon) W_2(\xi_b,\epsilon_b,\psi_b) \left[\prod_{j=1}^{M} d\Phi(\xi_j)\right] \left[\prod_{j=1}^{C} d\Phi(\epsilon_j)\right] d\Phi\left(\frac{\psi_b}{\sigma_{\psi_b}}\right) \\ &= \int \cdots \int W_1(\xi,\epsilon) \int W_2(\xi_b,\epsilon_b,\psi_b) d\Phi\left(\frac{\psi_b}{\sigma_{\psi_b}}\right) \left[\prod_{j=1}^{M} d\Phi(\xi_j)\right] \left[\prod_{j=1}^{C} d\Phi(\epsilon_j)\right] \end{aligned} \tag{F.3}$$

where the last component could be simplified as:

$$\begin{aligned}\mathcal{R}(\xi_b,\epsilon_b) &= \int W_2(\xi_b,\epsilon_b,\psi_b)d\Phi\left(\frac{\psi_b}{\sigma_{\psi_b}}\right)\\ &= \int \mathcal{I}\left[\mu_b+\xi_b+\epsilon_b+\psi_b \le -R\right]d\Phi\left(\frac{\psi_b}{\sigma_{\psi_b}}\right) = 1-\Phi\left(\frac{R+\mu_b+\xi_b+\epsilon_b}{\sigma_{\psi_b}}\right)\end{aligned} \tag{F.4}$$

where the continuous function $\mathcal{R}(\xi_b,\epsilon_b)$ explicitly reflects that it depends on two unobserved shocks ξ_b and ϵ_b. Importantly, after the derivations above we reduce the number of sampled shocks by 1 as we dont need to sample ψ_b.

Elimination of constraints related to not clicked options. Next, the set of shocks $\{\xi_j\}_{j=C+1}^{M}$ appear only in one constraint and could be simplified to:

$$\begin{aligned}LL &= \int\cdots\int W_1(\xi,\epsilon)\mathcal{R}(\xi_b,\epsilon_b)\left[\prod_{j=1}^{M} d\Phi(\xi_j)\right]\left[\prod_{j=1}^{C} d\Phi(\epsilon_j)\right]\\ &= \int\cdots\int W_{11}(\xi_{1..C},\epsilon)\mathcal{R}(\xi_b,\epsilon_b)W_{12}(\epsilon_b,\xi_b,\xi_{C..M})\left[\prod_{j=C+1}^{M} d\Phi(\xi_j)\right]\left[\prod_{j=1}^{C} d\Phi(\xi_j)\right]\left[\prod_{j=1}^{C} d\Phi(\epsilon_j)\right]\\ &= \int\cdots\int W_{11}(\xi_{1..C},\epsilon)\mathcal{R}(\xi_b,\epsilon_b)\left[\int\cdots\int W_{12}(\epsilon_b,\xi_b,\xi_{C..M})\prod_{j=C+1}^{M} d\Phi(\xi_j)\right]\left[\prod_{j=1}^{C} d\Phi(\xi_j)d\Phi(\epsilon_j)\right]\\ &= \int\cdots\int W_{11}(\xi_{1..C},\epsilon)\mathcal{R}(\xi_b,\epsilon_b)\mathcal{NC}(\epsilon_b,\xi_b,\xi_C)\left[\prod_{j=1}^{C} d\Phi(\xi_j)d\Phi(\epsilon_j)\right]\end{aligned} \tag{F.5}$$

where

$$\begin{aligned}W_{11}(\xi_{1..C},\epsilon) &= \left[\prod_{j=1}^{C-1}\mathcal{I}[z_j \ge z_{j+1}]\right]\left[\prod_{j=0}^{C-1}\mathcal{I}[\omega_j \le \min\{z_C,\omega_b\}]\right]\mathcal{I}[\omega_C \le \omega_b]\\ W_{12}(\epsilon_b,\xi_b,\xi_{C..M}) &= \mathcal{I}\left[\min\{\omega_b(\xi_b,\epsilon_b),z_C(\xi_C)\} \ge \max_{s=C+1..M} z_s(\xi_s)\right]\end{aligned} \tag{F.6}$$

Next, we find $\mathcal{NC}(\epsilon_b,\xi_b,\xi_C)$ as:

$$\begin{aligned}\mathcal{NC}(\epsilon_b,\xi_b,\xi_C) &= \int\cdots\int \mathcal{I}\left[\min\{\omega_b(\xi_b,\epsilon_b),z_C(\xi_C)\} \ge \max_{s=C+1..M} z_s(\xi_s)\right]\prod_{j=C+1}^{M} d\Phi(\xi_j)\\ &= \int\cdots\int \prod_{s=C+1}^{M}\mathcal{I}\left[\min\{\omega_b(\xi_b,\epsilon_b),z_C(\xi_C)\} \ge z_s(\xi_s)\right]\prod_{j=C+1}^{M} d\Phi(\xi_j)\\ &= \prod_{j=C+1}^{M}\int \mathcal{I}\left[\min\{\omega_b(\xi_b,\epsilon_b),z_C(\xi_C)\} \ge z_j(\xi_j)\right]d\Phi(\xi_j)\\ &= \prod_{j=C+1}^{M}\int \mathcal{I}\left[z_j^{-1}\left(\min\{\omega_b(\xi_b,\epsilon_b),z_C(\xi_C)\}\right) \ge \xi_j\right]d\Phi(\xi_j)\\ &= \prod_{j=C+1}^{M}\left[1-\Phi\left(z_j^{-1}\left(\min\{\omega_b(\xi_b,\epsilon_b),z_C(\xi_C)\}\right)\right)\right]\end{aligned} \tag{F.7}$$

where we used the fact that $z_j(\xi_j)$ is an invertible function for each j, and index j reflects the fact that the function would depend on estimable parameters of the model.

Notice that the simplified form depends only on three unobserved shocks: ξ_C through z_C, ξ_b and ϵ_b

through ω_b (if $b = C$ or $b = 0$ it is only two shocks). This implies that we eliminated $M - C - 1$ sampled shocks, which, given that $M >> C$ would by itself lead to substantial improvement in both computational accuracy and time.

Elimination of purchase-related constraints. Equation (F.5) could be represented as:

$$\begin{aligned} LL &= \int \cdots \int \left[\int \cdots \int W_{11}(\xi_{1..C}, \epsilon) \mathcal{R}(\xi_b, \epsilon_b) \mathcal{NC}(\epsilon_b, \xi_b, \xi_C) \prod_{j=1}^{C} d\Phi(\epsilon_j) \right] \prod_{j=1}^{C} d\Phi(\xi_j) \\ &= \int \cdots \int \left[W_{111}(\xi_{1..C}) \int \cdots \int W_{112}(\xi_C, \epsilon) \mathcal{R}(\xi_b, \epsilon_b) \mathcal{NC}(\epsilon_b, \xi_b, \xi_C) \prod_{j=1}^{C} d\Phi(\epsilon_j) \right] \prod_{j=1}^{C} d\Phi(\xi_j) \\ &= \int \cdots \int W_{111}(\xi_{1..C}) \mathcal{PC}(\xi_b, \xi_C) \prod_{j=1}^{C} d\Phi(\xi_j) \end{aligned} \tag{F.8}$$

where

$$\begin{aligned} W_{111}(\xi_{1..C}) &= \prod_{j=1}^{C-1} \mathcal{I}[z_j \geq z_{j+1}] \\ W_{112}(\xi_C, \epsilon) &= \left[\prod_{j=0}^{C-1} \mathcal{I}[\omega_j \leq \min\{z_C, \omega_b\}] \right] \mathcal{I}[\omega_C \leq \omega_b] \end{aligned} \tag{F.9}$$

Now we show that $\mathcal{PC}(\xi_b, \xi_C)$ could have a simplified form if $b \neq C$:

$$\begin{aligned}
\mathcal{PC}(\xi_b, \xi_C) &= \int \cdots \int W_{112}(\xi_C, \epsilon)\mathcal{R}(\xi_b, \epsilon_b)\mathcal{NC}(\epsilon_b, \xi_b, \xi_C) \prod_{j=1}^{C} d\Phi(\epsilon_j) \\
&= \int \cdots \int \mathcal{I}[\omega_C \leq \omega_b] \prod_{j=0}^{C-1} \mathcal{I}[\omega_j \leq \min\{z_C, \omega_b\}]\mathcal{R}(\xi_b, \epsilon_b)\mathcal{NC}(\epsilon_b, \xi_b, \xi_C) \prod_{j=1}^{C} d\Phi(\epsilon_j) \\
&= \int \cdots \int \mathcal{I}[\omega_C \leq \omega_b] \prod_{j=0}^{C-1} \mathcal{I}[\omega_j \leq \min\{z_C, \omega_b\}]\mathcal{R}(\xi_b, \epsilon_b)\mathcal{NC}(\epsilon_b, \xi_b, \xi_C) \prod_{j=1}^{C} d\Phi(\epsilon_j) \\
&= \iint \mathcal{R}(\xi_b, \epsilon_b)\mathcal{NC}(\epsilon_b, \xi_b, \xi_C)\mathcal{I}[\omega_C \leq \omega_b]\mathcal{I}[\omega_0 \leq \min\{z_C, \omega_b\}]\mathcal{I}[\omega_b \leq \min\{z_C, \omega_b\}] \\
&\times \left[\int \cdots \int \prod_{j=1, j\neq b}^{C-1} \mathcal{I}[\omega_j \leq \min\{z_C, \omega_b\}] \prod_{j=1, j\neq b}^{C-1} d\Phi(\epsilon_j)\right] d\Phi(\epsilon_b)d\Phi(\epsilon_C) \\
&= \iint \mathcal{R}(\xi_b, \epsilon_b)\mathcal{NC}(\epsilon_b, \xi_b, \xi_C)\mathcal{I}[\omega_C \leq \omega_b]\mathcal{I}[\omega_0 \leq \min\{z_C, \omega_b\}]\mathcal{I}[\omega_b \leq z_C] \\
&\times \left[\prod_{j=1, j\neq b}^{C-1} \int \mathcal{I}[\omega_j \leq \min\{z_C, \omega_b\}]d\Phi(\epsilon_j)\right] d\Phi(\epsilon_b)d\Phi(\epsilon_C) \\
&= \iint \mathcal{R}(\xi_b, \epsilon_b)\mathcal{NC}(\epsilon_b, \xi_b, \xi_C)\mathcal{I}[\omega_C \leq \omega_b]\mathcal{I}[\omega_0 \leq \min\{z_C, \omega_b\}]\mathcal{I}[\omega_b \leq z_C] \\
&\times \left[\prod_{j=1, j\neq b}^{C-1} \int \mathcal{I}[\epsilon_j \leq \omega_j^{-1}(\min\{z_C, \omega_b\})]d\Phi(\epsilon_j)\right] d\Phi(\epsilon_b)d\Phi(\epsilon_C) \\
&= \mathcal{I}[\omega_0 \leq z_C] \int \mathcal{R}(\xi_b, \epsilon_b)\mathcal{NC}(\epsilon_b, \xi_b, \xi_C)\mathcal{I}[\omega_0 \leq \omega_b]\mathcal{I}[\omega_b \leq z_C] \\
&\times \left[\prod_{j=1, j\neq b}^{C-1} \Phi\left(\omega_j^{-1}(\min\{z_C, \omega_b\})\right)\right] \Phi\left(\omega_C^{-1}(\omega_b)\right) d\Phi(\epsilon_b)
\end{aligned} \tag{F.10}$$

The derivations for $b = 0$ (outside option) and $b = C$ (last clicked option) follow exactly the same steps and are presented below:

$$\begin{aligned}
b = 0 &: \mathcal{I}[\omega_0 \leq z_C]\mathcal{NC}(\xi_C) \left[\prod_{j=1, j\neq b}^{C-1} \Phi\left(\omega_j^{-1}(\min\{z_C, \omega_0\})\right)\right] \Phi\left(\omega_C^{-1}(\omega_0)\right) \\
b = C &: \mathcal{I}[\omega_0 \leq z_C] \int \mathcal{R}(\xi_C, \epsilon_C)\mathcal{NC}(\epsilon_C, \xi_C)\mathcal{I}[\omega_0 \leq \omega_C] \left[\prod_{j=1}^{C-1} \Phi\left(\omega_j^{-1}(\min\{z_C, \omega_C\})\right)\right] d\Phi(\epsilon_b)
\end{aligned} \tag{F.11}$$

Notice that for the case $b = 0$, the function $\mathcal{PC}(\xi_b, \xi_C)$ has a closed form solution, and thus the sampling of ϵ_b is not needed. For $b > 0$ the integral has no known closed form solution and should be approximated using sampled values of ϵ_b. However, because of the constraint $\mathcal{I}[\omega_0 \leq \omega_b]\mathcal{I}[\omega_b \leq z_C]$ the sampling could be more efficient if we sample $\epsilon_b \in (\omega_b^{-1}(\omega_0), \omega_b^{-1}(z_C))$ and adjust the likelihood function with corresponding probability.

After combining the equation, we can rewrite the original Equation (F.2) as:

$$LL = \int \cdots \int \prod_{j=1}^{C-1} \mathcal{I}[z_j \geq z_{j+1}]\mathcal{PC}(\xi_b, \xi_C) \prod_{j=1}^{C} d\Phi(\xi_j) \tag{F.12}$$

Equation (F.12) has C binary indicators. This still could result in inefficient estimation, as for the list of generated random variables $\{\xi_j\}_{j=1}^C$, a substantial proportion of the realization of $W(\xi, \epsilon, \psi)$ would still be equal to 0. To circumvent this problem, we notice that $\left[\prod_{j=1}^{C-1} \mathcal{I}\left[z_j \geq z_{j+1}\right]\right] \mathcal{I}[\omega_0 \leq z_C]$ has a chain like structure, thus, we can sample random variables in a more efficient way:

1. Sample random shock ξ_C such that $\omega_0 \leq z_C(\xi_C)$ and denote it as ξ_C^g. Store the probability of the event $\mathbb{P}\left(\omega_0 \leq z_C(\xi_C)\right) = \mathbb{P}\left(\xi_C \leq z_C^{-1}(\omega_0)\right) = \Phi\left(z_C^{-1}(\omega_0)\right)$
2. Sample random shock ξ_{C-1} such that $z_{C-1}(\xi_{C-1}) \geq z_C(\xi_C^g)$ and denote it as ξ_{C-1}^g. Store the probability of the event $\Phi\left(z_{C-1}^{-1}\left(z_C(\xi_C^g)\right)\right)$
3. Repeat Step 2 until random shock ξ_1^g is generated

This procedure allows sampling random shocks more efficiently from the corresponding truncated distributions. This ensures that $\{\xi_j^g\}_{j=1}^C$ are such that $\left[\prod_{j=1}^{C-1} \mathcal{I}\left[z_j \geq z_{j+1}\right]\right] \mathcal{I}[\omega_0 \leq z_C] = 1$. Therefore, Equation (F.12) could be approximated by:

$$\int \cdots \int W(\xi, \epsilon, \psi) dF(\xi, \epsilon, \psi) = \frac{1}{G} \sum_g \mathcal{H}^b(\xi_b^g, \xi_C^g) \Phi\left(z_C^{-1}(\omega_0)\right) \prod_{j=1}^{C-1} \Phi\left(z_j^{-1}\left(z_{j+1}(\xi_{j+1}^g)\right)\right) \tag{F.13}$$