ABSTRACT

Title of dissertation: TWO ESSAYS ON THE ROLE OF INFORMATION TRANSPARENCY IN MARKETPLACE OPERATIONS

Jane Yi Jiang, Doctor of Philosophy, 2024

Dissertation directed by: Professor Wedad J. Elmaghraby
Robert H. Smith School of Business, University of Maryland, College Park

This dissertation encompasses two studies on the crucial role of information within marketplace operations. Collaborating with two platforms, we deliver empirical evidence and offer prescriptive insights into how information is conveyed to and perceived by customers, and the consequent impacts on sellers and the marketplace at large.

The first study analyzes the introduction of the novel blockchain tracing technology into an online grocery marketplace. Our findings indicate that credible supply chain transparency encourages consumers to more readily buy traced products, especially those that are handling-sensitive or offered in less-trusted markets. Consequently, adopting third-party sellers experienced an average monthly revenue increase of up to 23.4%. By utilizing structural estimation to understand how consumers assess product attributes and quality, we highlight that consumer responses (and welfare effects) vary in sophistication and size based on their prior experience with the product category. Additionally, we establish that consumers deem blockchain-based.

The second study analyzes the unintended transparency issue associated with the pricing structure of bundle discounts and its consequences on product purchases and returns.

Our findings reveal that customers tend to overlook complex pricing structures, leading to impulsive buying and increased returns. Enhancing customer attentiveness of pricing can decrease the Retailer's return rates by 20.9%. Moreover, improving customer attentiveness to pricing benefits retailers by enabling them to create more versatile bundle offers, further optimizing their sales strategy.

TWO ESSAYS ON THE ROLE OF INFORMATION TRANSPARENCY IN MARKETPLACE OPERATIONS

by

Jane Yi Jiang

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2024

Advisory Committee:
Professor Wedad J. Elmaghraby, Chair/Advisor
Professor Ken Moon
Professor Ozge Sahin
Professor Ginger Zhe Jin
Professor Ashish Kabra

Acknowledgments

Embarking on this PhD journey, I encountered countless individuals whose support and guidance were pivotal to the successful completion of my dissertation. While it's nearly impossible to express my full gratitude or to name everyone who played a role in this journey, I will attempt to acknowledge their invaluable contributions.

First and foremost, my heartfelt thanks go to my advisor, Wedad J. Elmaghraby. Wedad has been an extraordinary mentor and collaborator, offering not only her academic and professional wisdom but also unwavering support and reliability, qualities I deeply admire. I extend my gratitude to my coauthors, Ken Moon and Ozge Sahin. Ken provided patient guidance during the initial stages of my research, helping me navigate through unfamiliar topics and methodologies. His advice on developing research projects and presentation skills will always hold a special place in my heart. Ozge, likewise, has been incredibly supportive, continuously encouraging me to pursue excellence in research and a career in academia.

My appreciation also extends to Ashish Kabra, who was always available to share his advice and insights on both research and career paths. I am deeply thankful for Ginger Jin, whose research and classes inspired me to develop my own research ideas. She has always been generous in listening and providing feedback. I am grateful to the entire Industrial & Organization Group within the Economics Department, especially Professors Andrew Sweeting and Chenyu Yang, for including me in their group discussions and offering guidance and feedback throughout my research journey.

The support from the DO & IT Department and the PhD Program was crucial; I was fortunate to be surrounded by love and support from all the faculty and my fellow PhD students. I cherish the companionship of all the PhD student friends who were with me through this journey: Wenchang Zhang, Daehoon Noh, Ziwei Cao, Huan Cao, Sahar Hemmati, Shubham Akshat, Jiannan Xu, Eunseong Jang, Efe Serkaya, Fred Bao, Liangzi Zhu, Hugo Mainguy, and Jiaxuan Sun. Special thanks to my cohort, Farzad Fathi, and my Information System PhD cohorts, Gujie Li and SungHyun Kwon. Farzad, in particular, provided significant academic and emotional support.

Lastly, my gratitude extends to all the professors who generously offered me opportunities in research upon my arrival in the U.S. I am grateful to chemistry professors Paris Svoronos and Sujun Wei, physics professors David Lieberman and Vazgen Shekoyan, biology professors Urszula Golebiewska and Patricia Schneider, and math professor Mercedes Franco of Queensborough Community College; physics professor So Takei, and math professors Kenneth Kramer and Stefan Ralescu of Queens College; Professor Arjun Yodh of the University of Pennsylvania and Professor Kane Jennings of Vanderbilt University, who welcomed me into their labs for summer research; and Professor Assaf Zeevi of Columbia Business School, who took me on as a research assistant. Their collective generosity and support opened the door for me to pursue an academic career.

Receiving such abundant love and support has been a profound blessing. To everyone who has been part of my PhD journey, I extend my deepest thanks.

Table of Contents

Chapter 1: Introduction

This dissertation delves into the crucial role of information in marketplace operations, highlighting its transformative impact on transaction dynamics across various platforms. In the era of rapid marketplace evolution, a myriad of tools and mechanisms has reshaped the traditional shopping experience, altering the way information is disseminated and perceived by consumers. My research presents two comprehensive papers that explore the nuances of information delivery within marketplaces, investigating the reception and perception of information by customers, and analyzing the subsequent influence on sellers and the platform at large.

Chapter 1, the first paper, examines the adoption of blockchain tracing within a grocery marketplace and its consequent effects on all marketplace participants. The introduction of blockchain tracing represents an innovative approach to enhancing trust and transparency in the marketplace, with potential benefits for all parties involved – consumers, sellers, and the platform itself. Yet, the impact is not uniform across all stakeholders. There exists a gap between those who can effectively utilize this technology and those who need it the most. From the consumers' perspective, this disparity hinges on their ability to understand and trust the information provided by blockchain tracing. More sophisticated or experienced marketplace users may navigate this new tool more adeptly, but it is the less sophisticated and the less experienced who need it the most. Similarly, from the sellers' perspective, established businesses with high-quality products might find it easier to embrace new technologies, though the benefits could be more significant for newcomers to the market. This paper offers a nuanced analysis of blockchain technology's introduction to the marketplace, focusing on its implications for consumer welfare and market dynamics.

Chapter 2 shifts the focus to the informational aspects of pricing structures, which subtly influence consumer purchasing and return decisions, thereby affecting retailers. Certain discount strategies come with inherently complex pricing structures. For instance, customers wishing to return part of a bundle discount purchase must pay the full price for any items they retain, embedding a 'cost' to returns that may not be immediately apparent due to the initial presentation of a bundle price. This lack of transparency can lead customers to overestimate the value of a bundle, resulting in increased purchases but even higher returns. This paper explores this unintended issue of price transparency of discount strategies and its impact on consumers' understanding of product pricing. Contrary to the traditional notion that retailers prefer less rational customers, our findings suggest that retailers stand to benefit from a customer base that is more attentive to pricing details.

Together, these papers contribute to a deeper understanding of the pivotal role information plays in shaping marketplace dynamics, offering insights into how transparency can influence all stakeholders in marketplaces.

Chapter 2: An Empirical Study of Blockchain-Driven Transparency in a Consumer Marketplace

Marketplaces have been reshaped continually by technological advancements ranging from railroads to the mobile Internet and search. Now, marketplace transparency is set to be transformed by blockchain technology, which offers a uniquely credible platform for tracing the flow of products through firms' supply chains. The popular belief is that blockchain will imminently enable consumers to comprehensively trace food products from farm to table, reveal counterfeits and contamination, and more accurately identify socially responsible products and firms. However, no previous empirical research has examined blockchain technology's impact on consumer marketplaces, and implementing a blockchain-based solution for supply chain transparency can be costly and complex. In collaboration with a leading Chinese technology company that operates the top online consumer grocery marketplace in China, we analyze its recent efforts to introduce blockchain tracing among its sellers. Our findings indicate that credible supply chain transparency encourages consumers to more readily buy traced products, especially those that are handling-sensitive or offered in less-trusted markets. Consequently, adopting third-party sellers experienced an average monthly revenue increase of up to 23.4%. By utilizing structural estimation to understand how consumers assess product attributes and quality, we highlight that consumer responses (and welfare effects) vary in sophistication and size based on their prior experience with the product category. Additionally, we establish that consumers deem blockchain-based information to be more reliable than sellers' claims.

blockchain, consumer learning, consumer trust, consumer welfare, empirical operations management, food safety, information asymmetry, online marketplace, social good, structural estimation, supply chain transparency

2.1 Introduction

Since antiquity, marketplaces have been crucial in matching supply with demand, catalyzing economic growth and helping lift hundreds of millions, if not billions, out of poverty over the last century alone [15]. To achieve these results, marketplaces as institutions embody a duality between transparency and opacity. Buyers enjoy transparent and immediate access to a wide array of competing products from various sellers. Yet, simultaneously, markets encourage anonymity and thrive on the willingness to transact with strangers, seeking the best products at the right prices, irrespective of their origins.

Online marketplaces have since leveraged technological advancements to fuel remarkable growth. Mobile technology has broadened markets, allowing millions of consumers and sellers online access [23]. Algorithm-driven recommendation systems, by personalizing and streamlining the shopping experience, have become key drivers of this expansion [5]. These systems expertly match supply with demand, revolutionizing how consumers interact with online platforms. Technology thus remains vital in maintaining transparency and accessibility as markets continue to evolve.

Blockchain, as an emerging technology, promises to elevate marketplace transparency well beyond simply replicating the transparency of physical marketplaces in expansive online ones. Unlike technologies like RFID sensors, which only enable the physical tracing of products in transit, blockchain provides a unified, decentralized platform for record-keeping and data sharing. It functions as an immutable, cryptographically secured ledger outside centralized control, accessible to all involved parties, and authenticated in a fully distributed manner. Such features position blockchain as a promising tool for secure and detailed recordkeeping across the entire supply chain. For instance, sensors on fishing vessels can create unchangeable records of wild seafood provenance by detailing the exact geoposition and weight of the catch as it is hauled in [67]. For smallholder coffee growers, vending machine-style kiosks equipped with advanced grading machinery can be stationed in villages to evaluate and securely report the growers' individual contributions in terms of quantity and quality [20]. Although the data are only as accurate as the input, granular and immutable records of the production, shipping, and processing steps along a product's journey make the manipulation of individual records extremely difficult. Owing to its potential as a tamper-proof backbone for supply chain transparency, blockchain has stirred considerable excitement for marketplace applications. The popular belief is that blockchain innovations will soon enable consumers to trace agricultural products from farm to table, uncover counterfeits and contamination, and more accurately verify the social responsibility of products and firms.

However, little empirical work to date examines blockchain's genuine impact on consumer marketplaces, and there are valid reasons to question its most optimistic predictions. For instance, agricultural products like seafood often suffer from rampant mislabeling and misidentification of origins [80, 81] and adulteration [85], considering that they typically pass through 10 to 15 intermediaries before reaching consumers [89]. While customers widely mistrust such information, particularly online [91], blockchain-backed data regarding supply chains can be complex. It is uncertain whether providing this information necessarily leads to better customer understanding. Absent significant impacts on customer purchases, firms may be reluctant to invest in costly blockchain tracing. Furthermore, sellers might fear that increased supply chain transparency could inadvertently compromise their competitive edge or expose them to scrutiny and criticism [16].

Secondly, firms also lack clear guidance regarding the effects of blockchain adoption. For example, it is plausible that new sellers, lacking long track records, may benefit from blockchain-verified claims to quickly gain reputation. Conversely, blockchain-backed revelations could widen the advantage of incumbent firms with well-established supply chains. Additionally, gains for individual sellers might hinge on specific product mix details, such as the need for cold chain storage. These are empirical questions, and the answers provide crucial insights into blockchain's impact on sellers and the marketplace.

Lastly, technological progress often carries unintended consequences, functioning as a double-edged sword. For example, automation has led to gains in manufacturing efficiency but also contributed to a widening wage gap, particularly disadvantaging men without higher education [90]. Information technology has spawned a digital divide, creating unequal access to digital literacy, electronic devices, and internet connectivity, resulting in disparities in opportunities and life outcomes across society. The debate over new technology's role in social inequality continues, weighing its transformative potential against its mixed effects on socioeconomic disparities. In the realm of marketplaces, radical transparency could disproportionately favor savvy and experienced customers or hinder emerging sellers still building their market presence and reputation.

Over the past two years, we have collaborated with JD.com (JD), a leading platform technology company that operates one of China's largest online consumer marketplaces. As part of its ongoing quality improvement efforts, JD recently initiated a program introducing blockchain tracing on its platform. Our study focuses on the adoption of JD's blockchain tracing technology among sellers of fresh produce, particularly seafood products, and leverages JD's proprietary data to explore three research questions:

1. **Effect on marketplace transparency:** Does blockchain-backed product information significantly reduce information asymmetry regarding product value? Do con-

sumers regard blockchain-backed claims as more trustworthy than other claims made by sellers?

2. **Effect on sellers:** How does blockchain tracing benefit sellers who adopt it? What impacts does it have on established versus upstart sellers, and how are these effects influenced by the sellers' product mix?

3. **Effect on consumers:** How do consumers respond to complex blockchain-backed supply chain information, and do they interpret it effectively? Does the introduction of blockchain either widen or bridge the information gap between novice and experienced marketplace buyers?

The remainder of the paper proceeds as follows. Section 3.2 summarizes the literature. Section 2.3 describes how blockchain tracing is implemented in practice, and Section 2.4 introduces the platform and setting we study. Section 2.5 presents our empirical analysis, while Sections 2.6 and 2.7 develop and apply a structural estimation model of Bayesian learning to delve into how experienced and inexperienced customers respond to blockchain tracing and the implications for consumer welfare. Finally, we conclude in Section 2.8.

2.2 Literature Review

Our work relates to several streams of research in operations management, marketing, and industrial organization. We build upon studies on information design, trust in information sharing, the economics of technological change, and applications of blockchain technology in supply chains.

A vast body of literature explores strategies to diminish information asymmetry in marketplaces (for reviews, see 14, 47, 113). Classical remedies like third-party certification, buyer protection programs, and consumer feedback systems help build seller or product reputations and guide consumers in identifying quality products. Yet, none of these mechanisms offers a complete solution in practice. Certifications and buyer protection programs can be manipulated or corrupted [30, 55, 58]. Consumer feedback systems, which are the dominant reputational mechanism for many online marketplaces, are found to be often noisy, inflated, and uninformative [24, 44, 57, 92].

Recent empirical studies focus on innovations that seek to improve these reputation-building mechanisms in marketplaces. Examples of such innovations include unilateral feedback systems [77], reward-for-feedback program [87], and new designs of platform quality certifications [63–65]. To this literature, we introduce credible supply chain transparency as an alternative solution, emphasizing the value of blockchain tracing, particularly for products or sellers yet to establish their market reputation.

Another stream of relevant literature studies the role of trust in information sharing. In operations management, the studies on trust focus mainly on the relationship between supply chain partners (18, 56, 66, 68, 100–103, 108). We study a novel extension of such sharing, wherein a seller shares credible supply chain information with consumers in order to earn their trust and drive sales of its agricultural products. In this regard, our work relates to research on the impact of disclosure on consumer trust and purchase intention (e.g., 28, 79, 88, 96). By developing empirical evidence, we further the understanding of supply chain transparency's role in matching supply and demand, quantifying gains in the cultivation of consumer trust.

A growing body of literature scrutinizes the influence of emerging technologies on societal disparities. In the realm of the labor market, both [3] and [25] argue that automation may displace humans from specific occupations. Conversely, [75] and [99] posit that AI could enhance productivity, particularly for lower-skilled workers, and thereby promote worker equality. In marketplace settings, well-intentioned interventions can unintentionally exacerbate social disparity. For instance, consumer protection-oriented data privacy regulations may inadvertently more heavily burden newcomers or those with limited purchasing abilities [112]. Algorithmic recommendation systems have driven online marketplace growth but face criticism for algorithmic bias [37] and algorithmic exclusion [114], risking the inadvertent perpetuation of societal inequalities. Our research emphasizes the potential of blockchain tracing to narrow the welfare gap, especially aiding novice consumers, who are often more susceptible to manipulation [78], and bolstering products that lack established reputations.

Perhaps most narrowly but with relevance today, our work contributes to a rapidly emerging literature that sheds light on the applications of blockchain technology within supply chains [10, 33, 38]. When used to trace individual product units, blockchain resembles RFID technology [48, 82, 83, 116]. Unlike RFID, however, the blockchain's distributed-ledger system renders the information it stores uniquely tamper-resistant and hence highly credible [10]. Theoretical models have found that blockchain tracing should be effective at signaling product quality or authenticity [40, 104] for wide-ranging purposes [36]. [118] offer empirical evidence that applying blockchain tracing in marketplaces can stimulate demand. Yet despite potentially helping retailers grow profits and reduce food waste [76], blockchain tracing can undesirably reduce retailers' incentives to source from responsible suppliers [111] or suppliers' incentive to provide high-quality products [40].Misalignment of stakeholders' interests may hinder the adoption of blockchain tracing, even when beneficial to overall welfare [22, 39, 46, 70]. By employing a structural model, our research unravels and quantifies the mechanisms driving consumer choices, furnishing further evidence of blockchain tracing's substantial impact on purchases and its potential to incentivize broader firm adoption.

2.3 Blockchain Tracing and Its Application to Supply Chains

Blockchain tracing employs blockchain technology to monitor and document the journey and various phases of products throughout the supply chain. At its core is a decentralized, tamper-resistant, time-stamped ledger that ensures the secure and immutable capture of each detail, enabling accurate source verification. Moreover, blockchain's ability to house extensive data establishes it as a valuable information repository.

Blockchain's inherent decentralization ensures that no single entity has overarching control of the ledger, which in turn fosters trust and minimizes the potential for fraud. This decentralized structure necessitates collaboration across the entire supply chain. From raw material suppliers to retailers, every stakeholder must consistently log relevant data on the blockchain, especially when products change hands or undergo significant alterations. This collective effort yields a transparent and easily verifiable product history, further bolstering trust and reducing the chances of monopolistic behavior or fraudulent actions.

Blockchain tracing, with its vast capacity for data storage, serves as a comprehensive information repository. It accommodates diverse data types, from textual descriptions and images to timestamps, geolocation details, and video feeds, exemplified by platforms like IBM Food Trust and Bext360. This abundance of data enriches blockchain tracing as a source of information, but it also requires that adopters thoughtfully determine what specific data to include before initiating a blockchain tracing system. These decisions are guided by a blend of industry standards, stakeholder feedback, regulatory requirements, and the overall objectives of the blockchain initiative. In the context of food product supply chains, commonly recorded information encompasses aspects like product origin, processing and manufacturing details, logistical insights, certifications, and the characteristics of the final product [67].

Blockchain tracing has been applied in two main ways. Early adopters, such as Walmart, use blockchain tracing as a tool to better manage fresh produce within its own supply chains. It is reported that this tool fortifies product quality assurance and aids in settling disputes among supply chain partners [59]. Additionally, there is a growing trend of extending this heightened transparency directly to consumers. Such transparency stands to equip consumers with insights, enabling informed purchasing choices aligned with their ethical and qualitative benchmarks. In blockchain-enabled markets, consumers might be better positioned to verify product authenticity and corporate ethical stances, sidelining inferior goods and dubious market players.

JD's implementation. JD Technology, an affiliate of JD, developed JD's blockchain infrastructure called JD Chain [71]. Distinctively, as JD Chain seamlessly integrates with JD's retail platform, it has greater consumer focus. Unlike many blockchain solutions where the choice of data inputs might be jointly determined to please multiple stakeholders, JD exercises a strong degree of centralized control over traceability standards, so that it can curate and present uniform and valuable information to their consumers.

In 2019, JD opened and promoted JD Chain access to both the 1P and 3P sellers within the fresh produce sector. Participation is voluntary and free of charge: sellers can choose which products to trace or whether to participate at all. Once a product becomes blockchain-traced, an additional green tag is displayed on the product page. On the product webpage at JD.com, consumers can find an introduction to blockchain tracing technology. Live video feeds of both the sea farm and the processing facilities are also available for consumers to watch. In addition, upon receiving delivery of their purchased products, consumers can scan the QR codes attached to their products' packaging to retrieve a complete history of locations and timestamps for their specific batches of products, as well as their inspection reports, as shown in Figure 2.1.

Consumers may exploit such unprecedented supply chain transparency in at least two ways. First, by verifying seller-made claims and observing supplier facilities, consumers

Figure 2.1: Examples of Blockchain-enabled Information Made Available to Consumers.

can better evaluate the trustworthiness of products and claims. Second, the rich supply chain information may educate and inform consumers, leading them to become more sophisticated at assessing products' quality based on how they are produced and handled within their supply chains.

The offer of JD Chain to all sellers in the fresh produce sector at no cost appealed to the many small sellers on JD's platform. JD aimed to bolster consumer trust by providing consumers with verified, in-depth product data, which could lift the reputations of sellers and the credibility of the platform as a whole. Though JD Chain is available at no charge, participating sellers are tasked with coordinating with their supply chain partners to meet JD's standards. To aid this integration, JD Technology provides hands-on assistance, guiding sellers through the infrastructure setup process.

2.4 Data and Background

Our dataset was obtained from JD, a leading e-commerce company that operates one of China's largest online grocery platforms. While blockchain tracing was adopted across several products within the seafood category, sea cucumbers experienced the highest adoption rate; hence we make it our focal product category. We obtain over 58 million consumer purchase records in the seafood product category spanning January 2019 to April 2021, including 1.4 million records on sea cucumber products.

2.4.1 Product Quality and Supply Chain

Like many other food products, a sea cucumber's quality critically depends on its supply chain. Sea cucumbers are a specialty food delicacy that is believed in East Asia to be nutritious and medicinal. Liaoning Province ("Liao") is the source of the most prized sea cucumbers in China that are deemed to have the highest quality. After sea cucumbers are grown and harvested, they are transformed into one of three product types. A sea cucumber is either (i) dried without salt, (ii) packaged as fresh "ready-to-eat," or (iii) dried with salt via one of many processes that are lower-cost and faster than drying without salt. Both dried-without-salt and ready-to-eat are considered higher quality premium products that sell at higher price points, while the remaining dried products (dried with salt) are deemed non-premium products and are often sold at lower price points (Figure 2.2).

Once the product type is decided, the product moves through multiple handling (transportation and storage) stages before it reaches its final end-consumer market. The product type determines the handling processes needed to support product quality: the final quality of fresh ready-to-eat products relies heavily on proper cold-chain transportation and storage, whereas dried products do not have such stringent logistical requirements. Therefore, the final quality of a product depends primarily on two types of characteristics: those that determine the *product type* (the origin of the product, whether it is dried-without-salt, ready-to-eat, or dried with salt) and the *supply chain dependent* (SC-D) characteristics (the product's transportation, handling, and storage) that influence its state of freshness at the time of delivery.

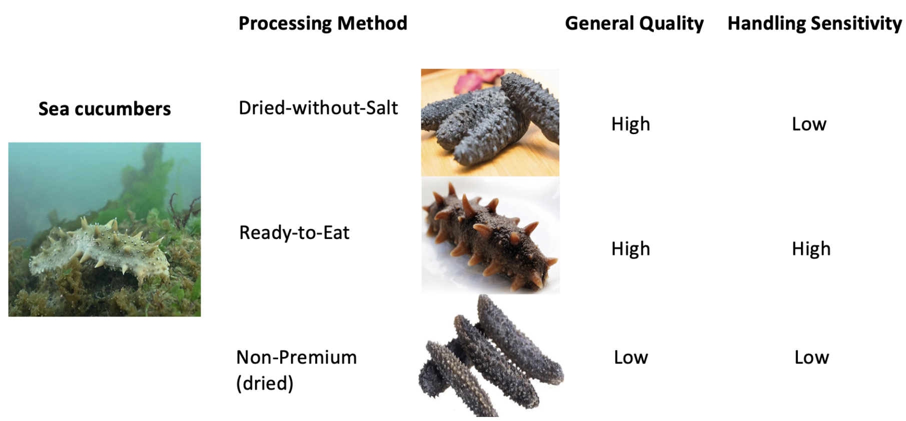

Figure 2.2: Sea Cucumber Types by Their Processing Methods. Photos are from the webpages of the products listed on JD's platform. They are for illustration purposes only and do not represent any specific product.

2.4.2 Marketplace

JD's online business-to-consumer platform consists of two sales channels. In the first-party, or 1P, channel, JD sells directly to consumers an assortment of seafood products that it selects and sources from suppliers. To encourage competition, multiple teams within JD operate the virtual storefronts (1P sellers) that sell these products. In the third-party, or 3P, channel, JD allows third-party sellers (3P sellers) to transact directly with consumers. This dual-channel sales model is prevalent in online retail platforms, capturing researchers' attention [45, 119]. All three types of sea cucumber products are sold on both the 1P and 3P platforms and generate 11.8 million USD in average monthly revenue.

When a consumer searches for sea cucumber products using JD's website or app, the search results draw from both the 1P and 3P sellers alike, thereby maintaining competition between them. Figure 2.3 illustratively depicts a pair of similar products sold by 1P and 3P sellers that were displayed in response to a sea cucumber search query. The red tag in the left-hand image indicates that it shows a 1P product. A similar two-channel structure with common search results is observed on Amazon.com, where 1P corresponds to Sold-by-Amazon and 3P corresponds to Amazon Marketplaces.

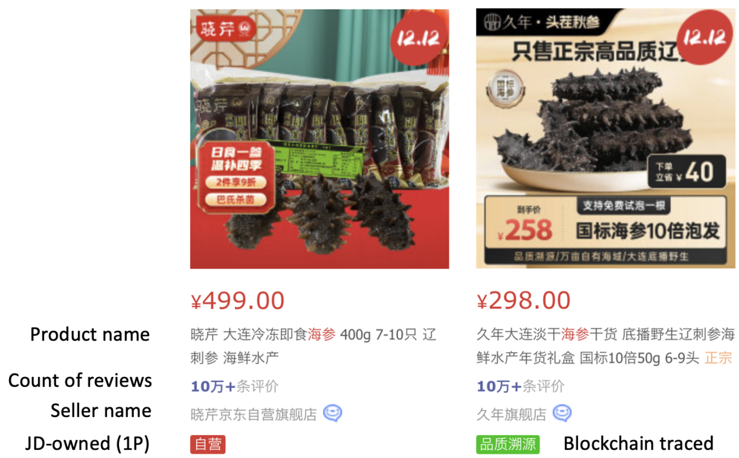

Figure 2.3: Products Displayed on the Search Page. Left. Product name: Xiaoqin (seller name) Dalian (a city in Liao) Frozen Ready-to-eat Sea Cucumbers, 400g 7-10 counts, Liao, sea cucumber, seafood; Count of reviews: 100,000+; Seller name: Xiaoqin JD-operated Flagship Store. Right. Product name: Jiunian (seller name) Dalian (a city in Liao) Dried-without-salt Sea Cucumbers, wild, gift packaging, national standard, 50g 6-9 count, authentic. Count of reviews: 100,000+; Seller name: Jiunian Flagship Store.

A product's quality depends on its product type and how it is handled and maintained throughout the supply chain. However, whether consumers choose to buy the product depends on their ability to acquire and trust information regarding these factors. For instance, a seller may *claim* that its product is of the highest quality type, namely, dried-without-salt from Liao, but consumers may question the seller's credibility and claim. Because 1P products are sold under JD's control, consumers exhibit a preference for 1P products. On average, 1P products sell at higher prices (Table 3.2); additionally, although only 16.5% of all the products sold monthly on the platform are 1P products, they contribute 43.2% of the

Sales Channel	1P		3P	
	mean	std. dev.	mean	std. dev.
Monthly revenue (million USD)	7.74	4.07	10.2	4.86
Number of active sellers per month	23.15	1.52	146.46	20.2
Number of active products per month	69.41	31.32	328.87	181.3
Quantity sold per month (million items)	0.76	0.42	1.68	0.65
Dried-without-salt	49.63%	21.05%	27.39%	18.56%
Ready-to-eat	9.47%	4.16%	9.31%	5.26%
Non-premium	40.90%	15.32%	63.30%	22.97%
Item price (USD)	10.17	13.05	6.07	9.81
Dried-without-salt	14.87	15.65	11.54	13.55
Ready-to-eat	9.39	8.13	9.13	10.68
Non-premium	4.66	6.9	3.25	6.31
High rating rate	98.55%	12.22%	98.58%	11.54%
Dried-without-salt	98.83%	12.43%	98.51%	11.41%
Ready-to-eat	98.01%	13.99%	98.80%	10.89%
Non-premium	98.43%	10.77%	98.68%	12.12%

Table 2.1: As the package size of each product varies, quantity is measured by the count of sea cucumber items. On average, a dried sea cucumber weighs around 5g and a ready-to-eat sea cucumber weighs around 35g. Price is measured on a per-item basis, calculated by dividing the unit product price by the quantity contained in each product. High rating refers to an aggregated rating of four or more stars out of five stars.

total revenue. Owing to consumer trust in JD's own brand and operations, its direct-sale products enjoy a significant advantage in reputation.

JD's platform also maintains a consumer feedback system that helps consumers make informed product purchases. Products' landing pages (see Figure 2.3) display their aggregate ratings and the reviews that previous purchasers voluntarily submit. Consistent with the prior literature, we find that sea cucumber ratings are inflated, which makes the feedback system less informative. More than 98% of the products have high ratings of four or five stars (Table 3.2). For blockchain-traced products, an additional green tag is displayed on the product page (see the right-hand image in Figure 2.3).

2.4.3 Consumers

The ability of consumers to interpret blockchain-backed product quality information heavily determines blockchain tracing's effect on marketplace information asymmetry. Previous literature consistently demonstrates that individuals refine their information-processing capabilities through experiential learning across various contexts [35, 62, 72]. When learning and assessing product (or seller) quality, experience enables consumers to filter out the noise in the information they receive and interpret new information more adeptly than inexperienced consumers. Supply chain information should be no different in this regard, and collectively, this literature suggests that consumers' experience levels should influence how ably they interpret and apply supply chain information about products.

Our dataset contains consumer-order level data that tracks every seafood product order that individual consumers (identified by unique customer codes) place on JD's 1P and 3P marketplaces. Monthly, on average, over 44,000 orders for sea cucumbers are placed by 37,710 consumers. Because sea cucumbers are considered a luxury food item, they are purchased infrequently: on average, 76.8% of the consumers make a single purchase within a year. Broadening the scope to encompass the entire seafood category, 46.1% of consumers purchase only once. This can be attributed to a hesitancy to buy seafood online, given the need for freshness and the elevated logistics requirements. Yet, these first-time consumers represent a notable portion of the demand and significantly influence revenue streams. At the same time, they might have limited capacity to discern quality and trustworthy sellers. Therefore, examining consumers with varying levels of experience becomes crucial to understanding heterogeneous responses to blockchain tracing.

We categorize consumers who placed at least one order in the seafood category in the past three months as being *experienced*; otherwise they are labeled as *inexperienced*. Consumers often base their perceptions of product quality on recent interactions [41]. Purchas-

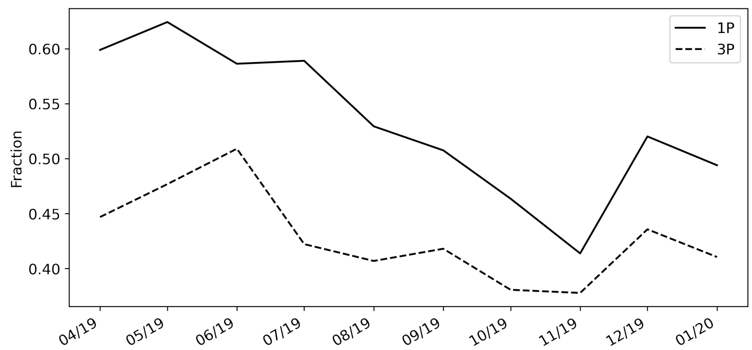

Figure 2.4: Fraction of Orders Placed by Inexperienced Consumers. The fraction of orders by inexperienced consumers is relatively stable over time in both 1P and 3P.

ing a product within the last quarter serves as such an interaction and aligns with current e-commerce segmentation practices. Engaging with the seafood category allows consumers to hone their skills in assessing product quality using factors such as stated origins, product types, and handling processes, all within the context of the same online marketplace. This accrued knowledge informs their future purchases in the seafood category.

The inexperienced consumers form a substantial part of the purchasing population. In an average month, they place 47.2% of all orders (Figure 2.4). Intuitively, inexperienced consumers ought to rely more heavily on JD's own reputation and trustworthiness and gravitate toward purchasing 1P products. Indeed, on average, 1P products derive 54.6% of their orders from inexperienced consumers each month, whereas 3P products derive a relatively smaller 43.6% of their orders from inexperienced consumers.

2.5 Effects of Supply Chain Transparency

We utilize the first wave of en masse adoption of blockchain tracing as a quasi-experiment to examine the effects of supply chain transparency. Two aspects of this adoption event and data uniquely enable our study. JD Technology first developed the technical infrastructure needed to support product blockchain tracing. JD then offered its tracing technology to all fresh produce sellers on its platform: JD Technology worked with participating sellers and their suppliers to incorporate the tracing solution into their supply chains. The tracing solution gives JD's end consumers access to the same supply chain information as the seller. Although the 3P sellers on JD's platform operate independently, the infusion of JD's blockchain tracing technology into their supply chains lets us study an implementation of supply chain transparency that is well developed and highly uniform across participating sellers.

This unique situation lets us empirically examine the introduction of supply chain transparency onto JD's platform as a quasi-experimental treatment effect. Sellers still self-select into adopting transparency via JD's technology or not; however, we crucially avoid the complications of sellers adopting a variety of partial blockchain solutions that vary in strength and credibility based on their motives and the information they might seek to conceal.

In this section, we begin by detailing the adoption event of blockchain tracing in Section 2.5.1. We then outline the methodology employed to assess the treatment effect in Section 2.5.2. Subsequently, in Section 2.5.3, we delve into the impact of blockchain tracing on the revenue of traced products. We pay particular attention to the heterogeneous effects across two dimensions: product reputation and consumer experience. Specifically, we investigate the "cold-start problem", evaluating how products with established reputations and those newly introduced benefit differently from blockchain tracing. We also explore the distinct responses of experienced and inexperienced consumers to transparency.

2.5.1 Quasi-Experiment: First Adoption of Blockchain Tracing

The first mass adoption of JD Chain took place in June 2019. Since JD developed and implemented the technology, all products that adopted JD Chain offered consumers the

	1P	3P
Number of Traced Products	13	70
Dried-without-salt	3	39
Ready-to-eat	5	24
Non-premium	5	7
Number of Adopting Sellers	3	5

Table 2.2: The first mass adoption event of block chain tracing took place in June 2019, which covers all three types of products from both 1P and 3P.

same uniform level of supply chain transparency. We therefore avoid the complications of sellers adopting different blockchain solutions that can partially reveal supply chain information in a manipulative way.

A total of 83 sea cucumber products adopted blockchain tracing. All the traced products were of Liao origin. This is likely attributed to Liao's well-established reputation among consumers as a source of high-quality sea cucumbers. Sellers may want to leverage this recognition by credibly signaling the origin of their products. However, not all Liao sea cucumbers are identical; they vary in both their sales channels and product types (Table 2.2). This diversity offers us a rich assortment of products to study, thereby providing insights into different aspects of the market. In 1P, 13.0% of sellers adopted, and 3.7% of sellers adopted in 3P. The majority of the traced products were premium dried-without-salt and ready-to-eat products, particularly so in 3P. We refer to January to May 2019, the six-month period before the adoption event, as the *pre-treatment period* and to July to December 2020, the six-month period after the adoption event, as the *post-treatment period*.

We examine the heterogeneous treatment effect of tracing adoption on product demand. It is reasonable to posit that a consumer's trust in a seller's product increases her willingness to purchase. Therefore, when supply chain transparency positively affects consumer trust, traced products should experience an empirically measurable increase in their demand. Yet, given the varying reputation levels of different sellers and their products, the impact of blockchain tracing can vary. This variability leads to an intriguing question: does blockchain tracing either bridge or widen the reputation gap between products?

For example, consider illustratively the firm (a ready-to-eat sea cucumber seller) that maintains dual storefronts on JD: a JD-operated 1P seller and a self-operated 3P seller (Figure 2.5). Although it offered a strikingly similar range of products through the 1P and 3P channels, its 3P offerings averaged prices that were about 20% lower during the pre-treatment period, reflecting that consumers perceive a reputational gap between 1P and 3P channels.

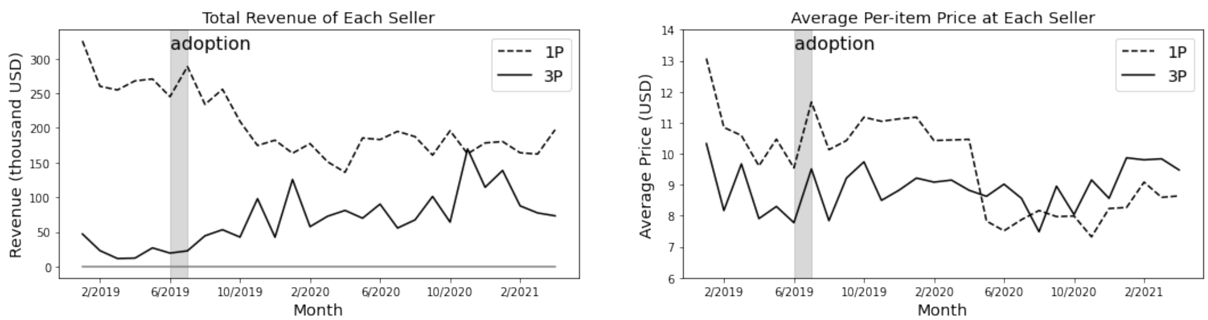

Figure 2.5: Revenue and Price Adjustments for a Firm Offering Through Both 1P and 3P Channels Following Blockchain Tracing Adoption. We tracked the total revenue and average per-item price for the firm's 1P seller and 3P seller from January 2019 to April 2021. The blockchain tracing adoption event occurred in June 2019, which is highlighted in grey. The left panel shows the total revenue for each seller, while the right panel presents the average per-item price for each.

Upon introducing JD Chain, however, this firm's 3P traced products witnessed significant revenue growth that was not mirrored in the 1P channel as illustrated in the left panel of Figure 2.5. Ten months after the adoption event, the 1P prices were adjusted to better align with those of the 3P, depicted in the right panel of Figure 2.5. This anecdotal example of a firm's experience under JD Chain underscores the profound capability of blockchain tracing to potentially close the gaps in perceived reputation that separate actually comparable products.

2.5.2 Robust Synthetic Control Methodology

In order to evaluate the causal effect of blockchain tracing on revenues, we use robust synthetic control (RSC) methods. RSC builds upon the classical synthetic control method (see 1 for a recent review) to better address the challenge of missing and noisy observations [8]. It also relaxes the convex constraints on the weights used under the classical synthetic control method; this relaxation makes the choices of control variables (donor pool) more flexible and generates synthesized units with a better fit. The statistical significance of the effects estimated by RSC are inferred by running placebo tests. The methodological details are provided in Appendix 2.9.

Since sellers and products self-select into JD Chain, care must be taken in identifying the effect of the blockchain tracing program. Specifically, the products that adopt blockchain tracing may already possess higher than average quality (e.g., all of them are from Liao). Therefore, a properly measured effect of blockchain tracing should exclude such pre-existing differences in quality. To address this potential endogeneity bias, we match each of the products in the treatment group (the adopters of blockchain tracing) with a donor pool of controls that best fits its characteristics during the pre-treatment period in terms of product type, sales channel, price point, length of sales on the platform (product tenure), and average rating. Please refer to Appendix 2.9 for details.

For our analysis, we aggregate the data to product-month level [73]. From the matched donor pools, the RSC method constructs for every treated unit a synthesized control unit that mimics the treated unit's outcome variable (e.g., revenues) during the pre-treatment period. Then, the treatment effect is measured using the differences in the outcome variable that emerge between the treated unit and its synthesized control unit during the post-treatment period.

2.5.3 Heterogeneous Effects of Blockchain Transparency

As blockchain tracing enriches the depth and detail of the product information accessible to consumers, it therefore allows them to confidently purchase products that match their quality preferences well. In the post-treatment period, changes to products' aggregated sales and revenues reflect consumers' updated choices in response to the newfound blockchain transparency into products' supply chains.

Figure 2.6 shows, for each of the three product types, the difference in average monthly revenues between the treatment group of traced products and the group of corresponding synthetic control units. During the pre-treatment period, the monthly differences are nearly zero, indicating that the synthetic control units closely match the treatment group well on the whole.

After tracing is introduced, 3P marketplace revenues significantly increase for the traced premium products (either dried-without-salt or ready-to-eat). In contrast, the 3P marketplace revenues of the non-premium products show no improvement over their synthetic controls. In particular, Table 2.3 shows, 3P ready-to-eat products experienced the largest revenue growth from tracing, amounting to 23.4% of their monthly revenue. Given their need for cold-chain logistics (see Section 2.4.1), such ready-to-eat products are the most handling-sensitive among the three product types. Therefore, the newly credible information that tracing provides regarding a product's supply chain dependent (SC-D) quality, such as transportation and handling details, are particularly impactful for the ready-to-eat sea cucumber products. In contrast, consumers appear justifiably indifferent as to whether a non-premium dried product is traced or not (insignificant change in revenues).

Interestingly, in the 1P marketplace, tracing results in no significant change in *overall* revenue for any of the three product types. Note that exactly the same blockchain tracing system was adopted by JD's 1P and 3P sellers, and hence consumers accessed the same

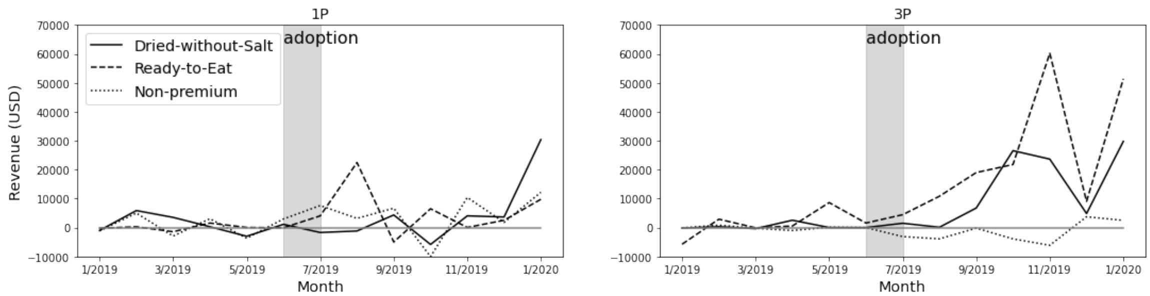

Figure 2.6: Difference in Monthly Average Revenue between Traced and Non-Traced Products.

	1P	3P
Dried-without-salt	4823.27	13312.12**
Ready-to-eat	16599.58	25240.71**
Non-premium	4548.57	-1569.55

Table 2.3: $*p < 0.10$; $**p < 0.05$; $***p < 0.01$. The presented values correspond to the monthly mean revenue differences between the treated group (the traced product) and their corresponding synthetic controls, computed using Eq. 2.4. The statistical significance of these differences is inferred based on Eq. 2.6.

type of new information in both sales channels.

In summary, we find that blockchain tracing positively impacts the 3P marketplace revenues of premium products but produces no significant revenue impact in the 1P marketplace. Our finding suggests that the growth in revenue stems primarily from blockchain adoptions producing new information that generates greater confidence and trust where it was lacking in the 3P marketplace and in the claims made by 3P sellers. In consumers' eyes, blockchain transparency provides credible new information regarding the 3P sellers' products.

Product reputation. We pay special attention to newly introduced products that face the so-called "cold-start problem" highlighted in the literature. On JD's platform, every product's webpage displays both its aggregate rating and a running tally of its units sold to customers. This arrangement inherently advantages products with well-established sales histories. We address the challenge that less established sellers face in promoting newer, but high-quality, products. Blockchain's impact on product revenue ought to vary heterogeneously between newly introduced young products and those with established standing in the marketplace.

We group the products into the two categories of established and young based on their existing reputational exposure on the marketplace. An established product refers to one that has already forged a reputation, either from its own sales history or based on its seller's stature. In other words, if the product's seller possesses reputably high sales volumes across its other products, the new product is more likely to be seen and trusted by customers. Accordingly, we define an established product as one that (i) has been available on the platform for an above-median length of time for its product type or (ii) is sold by a seller with an above-median number of historical orders for that same product type (dried-without-salt, ready-to-eat, or non-premium). The remaining products are categorized as young products. By this definition, established products either possess greater sales records, leading to a clearer track record of reviews and sales, or are sold by more established and visible sellers.

Of the traced products, 35 are classified as young, while 48 are considered established. During the pre-treatment period, young products logged an average monthly revenue of \$38.3 thousand for young products, whereas the established products averaged \$169.5 thousand. An established product's revenue therefore averaged 3.23 times that of a young product, underscoring the cold-start problem.

To fairly compare the magnitude of blockchain's effects across the established products and the young products, we compute the percentage change in monthly revenue between the pre- and post-treatment periods. Specifically, for each product, we calculate the average monthly revenue in the post-treatment period, denoted as $\bar{Y}_{post}$, and divide it by the average monthly revenue in the pre-treatment period, represented by $\bar{Y}_{pre}$. Our resulting outcome variable, $\frac{\bar{Y}_{post}}{\bar{Y}_{pre}} - 1$, equates to the proportional increase or decrease in revenue after the adoption of blockchain tracing. This enables us to determine which group, whether young or established products, garners more benefit in terms of revenue growth.

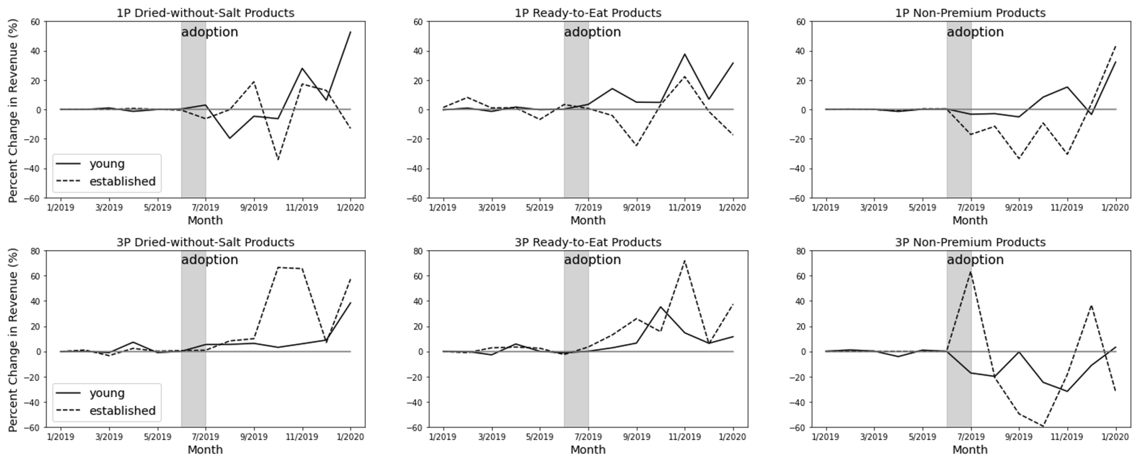

Figure 2.7: Difference in Percentage Change in Average Monthly Revenue by Product Reputation Groups.

	1P		3P	
	Product Reputation Level		Product Reputation Level	
	Young	*Established*	*Young*	*Established*
Dried-without-salt	8.44	-0.66	10.51**	30.80**
Ready-to-eat	14.75**	-3.09	11.04*	24.72**
Non-premium	5.87	-7.85	-14.59	-11.51

Table 2.4: $*p < 0.10$; $**p < 0.05$; $***p < 0.01$. The presented values correspond to the differences in percentage change of monthly mean revenue between the treated group (the traced product) and their corresponding synthetic controls, computed using Eq. 2.4. The statistical significance of these differences is inferred based on Eq. 2.6.

Figure 2.7 illustrates the percentage changes in young and established products' revenue for each product type across JD's two sales channels. We find nuanced differences in blockchain's effects on the two groups across the sales channels.

In the 3P marketplace where blockchain exhibited strongly positive revenue effects, both established and young premium products see notable revenue increases after adopting blockchain tracing. Established products experience a 30.8% surge in monthly revenue for dried-without-salt products and a 24.72% rise for ready-to-eat products. Despite being new to the market, young products also show significantly enhanced revenues, with a 10.51% boost for dried-without-salt products and an 11.04% increase for ready-to-eat products (Table 2.4). These substantial revenue hikes for young products underscore the efficacy of blockchain tracing as a reputation-building tool and a solution to the cold-start problem. However, established 3P products saw substantially larger revenue increases, suggesting that consumers gained even greater confidence and trust in traced products with an established presence on the market.

In contrast, we had found no empirical evidence of overall improvements to revenue in the 1P marketplace. However, we find interesting effects for young products if we consider them independently. Perhaps unexpectedly, young ready-to-eat products sold in the 1P channel logged a statistically significant 14.75% rise in revenue (Table 2.4). However, as previously shown in Table 2.3, we had found no overall average effect, and consistent with that earlier finding, we find no significant change in revenue for the established products in the 1P marketplace. Thus, although it appears that the products sold in the 1P marketplace are generally trusted and hence do not benefit significantly from blockchain tracing, we find that blockchain tracing still ameliorates the cold-start problem for 1P young products.

Thus, we find that young products benefit from blockchain tracing across the board, including to lessen the cold-start problem in trusted marketplaces. In the well trusted 1P marketplace, the net effect of tracing is to help close the perceived reputational gap between young and established products. However, in the 3P marketplace, established products also benefit greatly by using blockchain tracing to credibly reveal their supply chains. While both young and established products benefit, established products can credibly exhibit their advantages to consumers and the gap between them in consumer demand widens.

Consumer experience and sophistication. We further find that that consumers' experience-based knowledge plays a critical role in explaining blockchain's heterogeneous effects on consumer demand. Whereas we also consider other consumer characteristics such as age, gender, education, and platform loyalty, the differential impact is starkest when we segment consumers by their levels of prior purchasing experience with seafood products on the platform. We find that blockchain tracing induces stronger shifts in product demand among relatively more experienced consumers.

We separate the net gains in traced products' revenue into those generated by experienced consumers' orders and inexperienced consumers' orders.[1] At the same time, experienced consumers account disproportionately for the overall consumer response. They account for a larger fraction of the 3P increase in revenue, especially for the ready-to-eat products (Figure 2.8 and Table 2.5). Experienced consumer demand shifted notably towards traced premium products.

In the 1P channel, the unique responsiveness of experienced consumers is even starker. We find a significant increase of 14.3% in average monthly revenue in traced product revenue *solely in experienced consumers' orders placed* for the handling-sensitive, ready-to-eat products. The overall effect is insignificant when inexperienced consumers are included, and we find no significant effects for all instances involving inexperienced consumers (or involving the less handling-sensitive product types) in the 1P marketplace.

	1P		3P	
	Consumer Experience Level		Consumer Experience Level	
	Experienced	*Inexperienced*	*Experienced*	*Inexperienced*
Dried-without-salt	-1546.50	-450.78	4109.36***	1990.66*
Ready-to-eat	12231.88**	140.83	7192.41***	2634.60*
Non-premium	-890.03	4930.10	1092.38	418.08

Table 2.5: $*p < 0.10$; $**p < 0.05$; $***p < 0.01$. The presented values correspond to the monthly mean revenue differences between the treated group (the traced product) and their corresponding synthetic controls, computed using Eq. 2.4. The statistical significance of these differences is inferred based on Eq. 2.6.

[1] We experimented with defining experienced consumers as those who have made a purchase either once in the two months or once in the one month leading up to the current purchase. The pattern of our findings remained robust across these definitions.

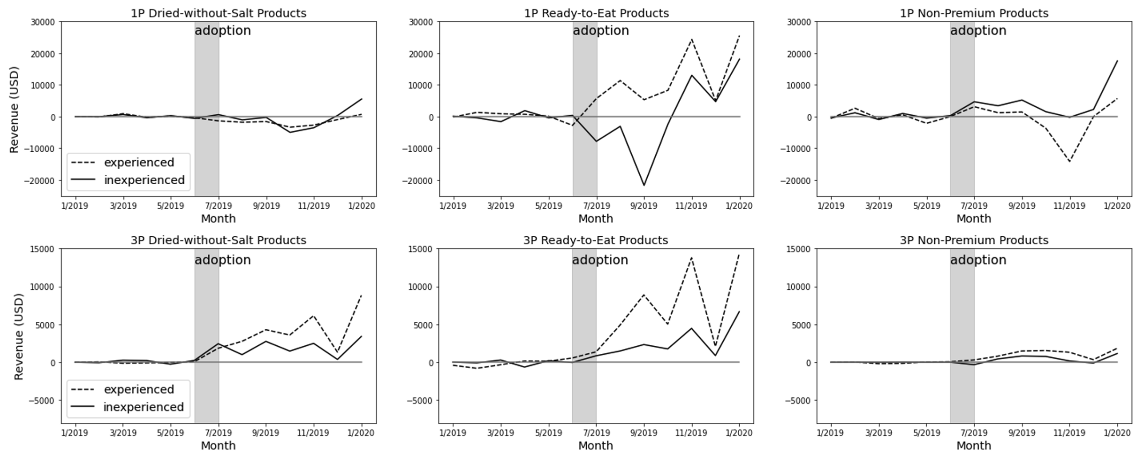

Figure 2.8: Difference in Monthly Average Revenue between Traced and Non-Traced Products by Consumer Experience Groups.

Stepping back, our empirical findings support that blockchain generates credible supply chain information that helps consumers assess both dimensions of product quality. Specifically, the premium products sold in the less trusted 3P channel show across-the-board gains in revenue that reflect basic improvements in the perception of product type quality. Blockchain adoption yields the greatest gains for the 3P market's established products, as tracing resolves consumers' doubts regarding the credibility of sellers' claims. In contrast, trust is already well established in the JD-backed 1P marketplace. Here, blockchain results only in particularized gains in revenue. Exclusively experienced consumers recognize that blockchain tracing reveals important new information regarding the handling of handling-sensitive, ready-to-eat premium products. Thus these savvy consumers exploit new information regarding products' SC-D quality, but we otherwise find no significant changes in 1P consumer demand. In the subsequent section, we introduce a structural model to elucidate the mechanisms behind consumer purchasing decisions.

2.6 Structural Model

To build on our empirical insights, we next develop a structural model of consumer choices based on marketplace information. In particular, the model will tease apart the roles that information and preferences play in consumers' heterogeneous reactions to blockchain tracing. For example, do experienced consumers react more strongly to tracing due to their superior understanding of the newly provided information, or because the products that they particularly prefer are more heavily affected by the new, blockchain-driven disclosures? We then use the estimated model to explore marketplace and welfare implications.

We model consumers' efforts to determine the quality of products prior to purchasing. These assessments should take into account the products' visible attributes, as well as any reputational history and whether the product has been selected for the 1P market channel. We model this assessment as a Bayesian learning process. For each product, a consumer forms beliefs regarding its two-dimensional quality spanning its product type (i.e., its granular quality level based on origin and method of preparation) and its supply chain handling (i.e., how well it is transported, handled, and stored). Ultimately, her preference for the product depends on these quality attributes.

Her beliefs incorporate three sources of available information. First, the consumer considers the product's visible attributes such as its stated origin, ratings, return rates, and tenure. We model these as being jointly distributed with the two quality dimensions that the consumer values. Based on the degrees of correlation between attributes and quality, she forms posterior beliefs regarding quality. Second, consumers account for the fact that JD uses superior information regarding products' quality attributes to select the 1P channel's product assortment. Consequently, the customer's beliefs further incorporate whether the product is being sold through the 1P or 3P channel. Lastly, the customer inspects the product pages to glean information (e.g., photos, written descriptions) that amounts to a noisy signal of the product's quality. We allow the noisiness (equivalently, the accuracy) of her product quality signal to depend on whether the product is blockchain-traced and on her level of prior experience with seafood products on JD.

A customer's final beliefs reflect some residual uncertainty regarding the products' quality. Such uncertainty quantifies, for specific consumers and product types, the granular degree of information asymmetry. Consumers are allowed to value the two quality dimensions heterogeneously.

2.6.1 Product Quality and Consumer Utility

In our model, customers make their purchasing decisions based on product prices and their assessments of product quality. In this section, we model product quality and how customers value it against price. Product quality is defined along two vertical dimensions that we address in turn.

Product type quality. First, consumable food products possess quality characteristics that derive from how they are sourced and prepared. These characteristics determine a product's taste and feel, as well as how it is best enjoyed by the consumer. For example, Norwegian smoked salmon is distinguished by its sourcing in the North Atlantic, potentially from the wild, and from the smoking and packaging processes used to prepare it. A consumer's individual affinity for a smoky flavor and for the quality and typically lower fat content of Norweigian-sourced salmon determines her utility value for this particular salmon product type.

We envision that narrow product categories such as sea cucumbers feature generally vertical gradings of product types that range from the premium quality products down to the cheaper non-premium preparations that compromise on flavor and shelf life. Because customer tastes are subjective, our model allows customers to differ in their sensitivity to such quality and to possess a degree of idiosyncratic tastes that affect their choices.

Between products, labels do not fully determine actual quality. For example, despite Liao's exalted reputation, a poorly prepared sea cucumber product from Liao may rank below a product that is sourced elsewhere but better produced. Therefore, we model every single product as possessing its own true level of product type quality (that is, quality levels are far more granular than the broad product types described in Section 2.4.1). Such quality levels are the latent characteristics of each product, and consumers form beliefs about products' latent product type quality levels based on examining their labels and photos, for example.

Supply chain dependent (SC-D) quality. Beyond sourcing and production, high-quality products must maintain freshness and safe and careful handling throughout their supply chains. Thus, the performance and practices of the sellers' supply chain intermediaries significantly impact product quality [27, 97]. Despite this, however, the differences in supply chain characteristics and quality between products are often practically invisible to end consumers.

Customer utility. Formally, a consumer i's utility from product j depends on the product's two-dimensional quality $Q_j = \left(Q_j^1, Q_j^2\right) \in \mathbb{R}^2$. Specifically, Q_j^1 represents its product type quality, and Q_j^2 represents its SC-D quality. Customers value the two dimensions of quality heterogeneously, so that we capture a customer i's individual valuations in the form of the two-dimensional parameter vector, $\beta_i \in \mathbb{R}_+^2$. From consuming product j at the price P_j, she obtains the net utility payoff

$$u_{ij} = \beta_i{}^\top Q_j - \alpha_i P_j + \varepsilon_{ij},$$

where ε_{ij} is the idiosyncratic utility component that reflects her subjective taste. Subject to the information gathering process described below, the customer seeks to make a purchasing decision that maximizes her expected utility.

2.6.2 Consumer Learning Process

While shopping in the marketplace, a consumer i collects information about the quality of each product j. She forms a *belief* about its quality in the form of a distribution over possible Q_j that we represent using the random variable Q_{ij}^{post}. Her beliefs are based on: viewing the product's visible attributes, including its labels regarding brand, product type, and sourcing; whether the product was selected by the platform for 1P status; and finally her assessment of additional unstructured information, such as that revealed by blockchain tracing or from photos on the product's page. Note that Q_{ij}^{post} is individual-specific, allowing each customer her own belief about quality of product j.

Product attributes.

Whereas a product j possesses a two-dimensional quality Q_j, that quality is never directly revealed to the customer. Instead, the marketplace displays a host of other observable attributes that can inform customers regarding products' underlying quality.

In order to capture this setting, we define each *product* j to be the vector that consists of its quality Q_j and its set of the attributes $A_j \in \mathbb{R}^K$. We denote each product as $X_j = [Q_j^1, Q_j^2, A_j^1, ..., A_j^K]^\top$. In turn, a *product category* is defined by a joint distribution over quality and attributes. The product category thus generates a prior marginal distribution for the quality Q_j of any product drawn from it, absent any further information about the product. It also gives customers the means to update their beliefs about Q_j after observing the product's attributes A_j.

In our model, we let the attributes A_j consist of a product's origin, rating, return rates, complaint rates, how long the product has been sold on the marketplace (product tenure), and how long the seller has been doing business on the marketplace (seller tenure). In keeping with how the customer encounters them, we take these attributes exactly as they are disclosed by the platform and seller. That does not mean, however, that such information is accepted uncritically. When certain types of claims are relatively unreliable, the product category's joint distribution should reflect that these claims, as visible attributes, correlate only weakly with the actual product quality, Q_j. As a result, they exert rationally weaker influence over customers' beliefs about quality, i.e., customers draw rationally weaker conclusions about quality from them. Yet, even highly credible attributes may be informative about only one quality dimension or the other. For example, as described in Section 2.4.1, origin and processing method are highly correlated with product type quality, whereas seller tenure and product tenure may strongly correlate with SC-D quality.

We specify the product category as a multivariate normal distribution $N(\mu,\Psi)$. When there are J products on the market, each product $X_j = [Q_j^1, Q_j^2, A_j^1, ..., A_j^K]^\top,\ j = 1, ..., J$, is deemed to be drawn i.i.d. from this distribution.

Platform selection.

Most customers shop once or a few times per year for sea cucumbers, leading them to rely mainly on the information they see. The platform, however, possesses much richer information regarding the sellers on its marketplace and their products. Moreover, the platform then uses its own superior information to determine which products to sell through the 1P channel. Customers' beliefs regarding products' quality should then rationally incorporate the chosen channels (1P or 3P) through which they are sold.

To model the platform's choice, we assume that it knows the quality and attributes of every product on its platform. On a monthly basis, the platform selects qualifying products for its 1P sellers using a simple classification rule based on a threshold r: product j is sold on the 1P channel if and only if $\Lambda X_j + \nu_j \geq r$, where the vector $\Lambda \in \mathbb{R}^{K+2}$ reflects the weighted influence of the attributes and quality, and ν_j is an error term that is assumed to be normally distributed and independent across products $j = 1, \ldots, J$.[2] This selection rule reflects trade-offs between quality and attributes. For example, if a product is sourced from a supplier that has a long selling history on the market and a low complaint rate (both favorable attributes), then JD may select it for 1P despite it being a non-premium product. Similarly, JD may promote a new supplier through its 1P channel if it assesses that the supplier's products exhibit superior quality.

Unstructured product information including blockchain.

Lastly, a customer incorporates the rich and potentially unstructured information provided by the marketplace about each product. Such information includes written descriptions and pictures and, for traced products, the novel pieces of information (e.g., verified sourcing details, video feeds) produced about their supply chains.

We model that a customer i acquires and interprets such information in the form of a private signal that she receives regarding the quality of product j. The signal is noisy but centered at the product's true quality Q_j. The signal's noisiness depends both on how much information is provided (e.g., whether the product is blockchain-traced) and on the customer's own ability to comprehend it. Therefore, we allow the signal's precision to vary between products and customers. Regarding the latter, because experienced consumers have built some knowledge in the product category [51, 52], we allow that they could derive more precise quality signals than inexperienced customers.

[2] According to interviews with JD managers, in practice they first select suppliers before filtering products. Then JD inspects each batch of products before they become listed in the 1P sellers' storefronts.

Therefore, consumer i receives the private signal q_{ij} regarding product j's quality,

$$q_{ij} = Q_j + \Delta_{ij}, \tag{2.1}$$

in which $\Delta_{ij} = [\Delta_{ij}^1, \Delta_{ij}^2]^\top$ is the signal noise that follows a bivariate normal distribution $N(0, \Sigma_{ij}^\Delta)$. As prefaced above, we let the matrix Σ_{ij}^Δ vary heterogeneously across products j and based on whether consumer i is experienced in the product category.

Bayesian updating and information asymmetry.

Through the learning process just described, a customer i forms her final beliefs regarding product j's quality Q_j based on observing its sales channel (1P or 3P), its attributes A_j, and finally the private signal q_{ij}.

We now specify her beliefs after they are derived through Bayesian updating as follows. Appendix 2.11.1 provides more precise details regarding the Bayesian updating procedure.

First, the customer's beliefs are updated in a single step based on product j's visible attributes A_j *and* channel (1P or 3P). Unlike the private signal that we address subsequently, this first Bayesian update is common to all customers who consider buying product j. If we were to take in the product's visible attributes A_j alone, the Bayesian update would be exceedingly simple: based on the category's jointly normal distribution $N(\mu, \Psi)$ of attributes A_j and quality Q_j, customers would update their priors in the form of Q_j's marginal distribution to the posterior distribution of Q_j conditional on A_j. (Recall for the following that the product covariates X_j include both A_j and Q_j.) Instead, we also need to contend with the fact that the channel-choice classification rule $\Lambda^\top X_j + \nu_j \geq r$ introduces both its own independent stochastic term ν_j and a classification based on truncated multivariate normal distributions over X_j that shift depending on the rule's parameters, Λ and r.

However, we make this extension tractable by expanding the product covariates from

X_j alone to (X_j, ν_j), that is, to cover both a product's quality and attributes and the random filtering component applied when selecting its marketplace channel. We then easily extend the definition of a product category to provide the multivariate normal distribution over the space of (X_j, ν_j), wherein ν_j follows an independent normal distribution. Within this space, the classification rule $\Lambda^\top X_j + \nu_j \geq r$ produces a half-space defined by a hyperplane; in turn, the distribution of products in 1P is given by the truncated multivariate normal distribution defined by the half-space as its support. (The classification rule produces a similarly truncated distribution for the 3P products.) Given a product's publicly visible channel and attributes A_j, we use these distributions to obtain the product's common posterior quality Q_j^{Public}. The exact mechanics of calculating Bayesian updates over these truncated spaces, and integrating out ν_j to do so, are given in the appendix.

Based on her private signal $q_{ij} = Q_j + \Delta_{ij}$, a customer i makes her final Bayesian update regarding product j's quality. This final application of Bayes' rule results in her personal posterior beliefs Q_{ij}^{post}. Both Q_j^{Public} and Q_{ij}^{post} are distributions over the possible values of Q_j that we are representing in the form of random variables.

In estimating the customer's learning process, a parameter of major interest is the matrix Σ_{ij}^{Δ} that governs the level of noisiness in Δ_{ij}; such noisiness represents the level of uncertainty in customer i's private signal of product j's quality. Thus, matrix Σ_{ij}^{Δ} represents the degree of information asymmetry in customer i's learning process. Moreover, its separate influences on the components Δ_{ij}^1 and Δ_{ij}^2 break down such information asymmetry into a customer's uncertainty regarding a product's quality in terms of product type and supply chain handling, respectively.

Thus, we focus foremost on Σ_{ij}^{Δ} and how it varies between products and between experienced and inexperienced customers. We also focus on whether blockchain tracing reduces information asymmetry and, if so, along which of the two dimensions of quality.

2.6.3 Consumer Purchases

Based on her beliefs regarding the quality of all available products $j = 1, \ldots, J$, a consumer i selects the product $j_i^* \in \text{argmax}_{j \in [J]} \mathbb{E}\left[u_{ij}\right]$ that maximizes her expected utility

$$\mathbb{E}\left[u_{ij}\right] = \beta_i{}^{\top} \mathbb{E}\left(Q_{ij}^{post}\right) - \alpha_i P_j + \varepsilon_{ij}, \tag{2.2}$$

where the idiosyncratic utility taste ε_{ij} is distributed i.i.d. extreme value. Through β_i and α_i, this utility function allows consumers to place heterogeneous value on the product quality dimensions and vary in price elasticity. A consumer chooses not to buy at all if the largest utility is less than the outside option's normalized utility value of zero.

Conditional on consumer i's posterior quality beliefs Q_{ij}^{post} and her valuations β_i and α_i regarding quality and price, the probability that her decision y_i is to buy product j is

$$L_i\left(y_i = j | Q_{ij}^{post}, \beta_i, \alpha_i\right) = \frac{exp\left(\beta_i \mathbb{E}(Q_{ij}^{post}) - \alpha_i P_j\right)}{1 + \sum_k exp\left(\beta_i \mathbb{E}(Q_{ik}^{post}) - \alpha_i P_k\right)}. \tag{2.3}$$

In addition to Σ_{ij}^{Δ}, we are interested in estimating the distribution of the valuation parameters β_i and α_i. See Appendix 2.11 for further details regarding estimation and identification.

2.7 Estimation Results

In this section, we present our estimation results. Beyond estimating the variation in how customers value the dual dimensions of product quality, we focus on the parameters that determine the levels of signal noise, or information asymmetry, affecting their perceptions of product quality. We compare such information asymmetry levels across the 1P and 3P sales channels, the different product types, and different consumer experience levels. By comparing the accuracy of consumer quality perceptions for products that are blockchain-traced against those that are not, we quantify supply chain transparency's effect on the measured degree of information asymmetry.

2.7.1 Overall Patterns

Consumer valuation estimates. Generally, the quality valuation parameters that we estimate for consumers appear reasonable (Table 2.6). Overall, consumers derive relatively consistent value from the two dimensions of product quality. Whereas there exists a modest degree of minor variation in the magnitudes of β_1 and β_2 over individual consumers, the levels are quite stable across the two marketplaces and levels of consumer experience. Inexperienced consumers tend to be modestly more price-sensitive and show less appreciation for SC-D quality, i.e., they display less sensitivity to how products are handled. This may stem from inexperienced consumers finding it harder to accurately perceive SC-D quality from the same information.

Information asymmetry as noise. In general, it appears to be substantially more difficult for consumers to accurately ascertain a product's SC-D quality than its product type quality. In Table 2.6, we observe that the uncertainty in consumer perceptions regarding SC-D quality ($\sigma^{\Delta}_{2,ij}$) is always greater than the uncertainty on product type quality ($\sigma^{\Delta}_{1,ij}$). Additionally, across the 1P and 3P marketplaces, information asymmetry is more pronounced on average for 3P products than for 1P products. In both markets, experienced customers receive significantly more accurate private signals regarding product quality than inexperienced customers. In short, experienced customers better perceive products' multi-dimensional quality from the same information; moreover, all consumer perceptions of quality are noisier in the 3P marketplace and for the quality dimension that captures the quality of product handling.

Valuation α_i, β_i				
	Consumer Experience Level			
	Experienced		*Inexperienced*	
	mean	std. dev.	mean	std. dev.
α (price)	1.6496	0.1039	1.7951	0.1134
β_1 (product type quality)	2.3718	0.2248	2.3818	0.2504
β_2 (SC-D quality)	2.4175	0.2183	2.0395	0.1971
Signal noise Δ_{ij}				
	Experienced		*Inexperienced*	
	σ_1^Δ	σ_2^Δ	σ_1^Δ	σ_2^Δ
1P	0.1555	0.2340	0.1643	0.2503
3P	0.1720	0.2461	0.2037	0.2840

Table 2.6: The parameters are estimated using 5% of the orders placed in the post-treatment period. Prices are divided by 100. The signal noise Δ_{ij} is distributed $N(0, \Sigma_{ij}^\Delta)$, where $\Sigma_{ij}^\Delta = [[\sigma_{1,ij}^\Delta, \rho_{ij}^\Delta], [\rho_{ij}^\Delta, \sigma_{2,ij}^\Delta]]$. Here we report $\sigma_{1,ij}^\Delta$ and $\sigma_{2,ij}^\Delta$, the noisiness of the signal that a consumer receives on product type quality and SC-D quality respectively. The signal noise values are weighed by the mean value of quality Q_j. The general patterns we observed on the noise parameters are: (i) $\sigma_{2,ij}^\Delta > \sigma_{1,ij}^\Delta$, (ii) both $\sigma_{1,ij}^\Delta$ and $\sigma_{2,ij}^\Delta$ are greater for 3P than for 1P, and (iii) both $\sigma_{1,ij}^\Delta$ and $\sigma_{2,ij}^\Delta$ are greater for inexperienced consumers than for experienced consumers.

2.7.2 Effects on Information Asymmetry

We assess the effects of supply chain transparency on information asymmetry by comparing the accuracy of consumer quality perceptions of products that are blockchain-traced against that of the non-traced products during the post-treatment period.

Average effect.

Consumers perceive the quality of both traced and non-traced products with noise. As components of the matrix Σ_{ij}^Δ, the parameters $\sigma_{1,ij}^\Delta$ and $\sigma_{2,ij}^\Delta$ control the levels of noise that infiltrate a consumer i's perception of the quality of product j (across the two quality dimensions, respectively). In terms of interpretation, these levels of noise reflect the presence of consumers' unresolved doubts regarding the credibility of the sellers' claims. Greater uncertainty (i.e., larger values of $\sigma_{1,ij}^\Delta$ and $\sigma_{2,ij}^\Delta$) reflects lower trust in the quality signal q_{ij}

that consumers derive from products' claims and cited attributes.

Such noisiness is, on average, smaller in magnitude for blockchain-traced products than it is for their non-traced counterparts. Thus, an average consumer more accurately perceives the quality of a traced product, all else being equal, than that of a comparable non-traced product, and transparency reduces information asymmetry. Table 2.7 shows that in the 3P marketplace, transparency affords products, on average, an estimated 15.6% reduction in the noise affecting consumer perceptions of product type quality ($\sigma^{\Delta}_{1,ij}$) and a 13.4% reduction for SC-D quality ($\sigma^{\Delta}_{2,ij}$). In the 1P marketplace, by contrast, the informational benefits to tracing are relatively muted: it reduces the noise affecting perceptions of product type quality by about 5.4% and of SC-D quality by 6.5%.

The accuracy of perceived product quality improves by a greater margin in the 3P marketplace; this differential indicates that the 3P marketplace suffered from a more significant credibility gap that could be addressed through credible transparency. Supply chain transparency has similar but smaller effects in the 1P marketplace.

		$\sigma^{\Delta}_{1,ij}$	$\sigma^{\Delta}_{2,ij}$
1P	Traced	0.1555	0.2341
	Non-Traced	0.1644	0.2503
	Reduction in signal noise (%)	5.38	6.48
3P	Traced	0.1720	0.2461
	Non-Traced	0.2859	0.3039
	Reduction in signal noise (%)	15.56	13.35

Table 2.7: The percentage of reduction in the noisiness of signal is calculated by taking the difference of the noise parameters of the traced and non-traced products, and dividing by the noise parameter of the non-traced products.

Heterogeneous effects by consumer experience level.

Supply chain transparency impacts information asymmetry differently for experienced and inexperienced consumers. Table 2.8 breaks down the reductions in uncertainty that tracing achieves in product type quality $\sigma^{\Delta}_{1,ij}$ and in SC-D quality $\sigma^{\Delta}_{2,ij}$ by the marketplace, product type, and customer experience level.

Several key take-aways stand out in Table 2.8. First, we see that supply chain transparency consistently reduces consumer uncertainty about product quality virtually across the board. Second, we once again confirm that quality uncertainty diminishes by a substantially larger margin in the 3P marketplace than in 1P; blockchain tracing introduces credible transparency that remedies a general lack of trust in the quality of goods sold by 3P sellers and in their unverified claims.

Third, transparency gives inexperienced customers relatively larger reductions in $\sigma^{\Delta}_{1,ij}$, which represents the level of uncertainty regarding product characteristics such as a product's origin and how it is produced. Such gains in accuracy are especially pronounced for the premium (dried without salt and ready-to-eat) sea cucumber products, for which uncertainty decreases by up to 23-25% in the less trusted 3P marketplace.

Fourth, transparency especially enhances experienced customers' ability to accurately perceive SC-D quality. The reductions in quality uncertainty in the SC-D dimension are starkest for the handling-sensitive ready-to-eat products, and are more substantial in magnitude in the 3P marketplace, with a decrease of up to 18%.

Therefore, inexperienced consumers benefit to a greater extent than others when tracing clarifies the credibility of seller-made claims, especially in the 3P marketplace. On the other hand, across both the 1P and 3P marketplaces, experienced consumers particularly benefit from their ability to more accurately assess products' SC-D quality using the new information that blockchain tracing provides.

		Consumer Experience Level			
		Experienced		*Inexperienced*	
		$\sigma^{\Delta}_{1,ij}$	$\sigma^{\Delta}_{2,ij}$	$\sigma^{\Delta}_{1,ij}$	$\sigma^{\Delta}_{2,ij}$
1P	Dried-without-salt	0.07	5.74	0.87	3.81
	Ready-to-eat	11.48	13.51	14.15	6.99
	Non-premium	10.20	5.08	-2.42	0.21
3P	Dried-without-salt	7.30	11.92	23.03	12.15
	Ready-to-eat	18.74	17.86	24.96	12.53
	Non-premium	10.48	14.36	12.71	6.06

Table 2.8: The percentage of reduction in the noisiness of signal is calculated by taking the difference of the noise parameters of the traced and non-traced products, and dividing by the noise parameter of the non-traced products. The values of all the noise parameter Σ^{Δ}_{ij}'s are reported in Appendix Table 2.17.

Heterogeneous effects by product type.

Consistent with our earlier empirical findings, supply chain transparency maximally alleviates consumer information asymmetry when it comes to the ready-to-eat products. Table 2.8 shows that, across both the 1P and 3P markets and for both experienced and inexperienced customers, the ready-to-eat sea cucumber products consistently realize the largest reductions in both dimensions of quality uncertainty when supply chain transparency is introduced.

Therefore, consumers appear to very rationally incorporate new supply chain information into their decision-making. The quality of ready-to-eat sea cucumbers is indeed heavily reliant on their careful handling with proper cold chain logistics throughout their supply chains. Supply chain transparency lets consumers access significantly better information regarding such handling, and Table 2.8 shows this informational effect.

2.7.3 Effects on Building Trust

The sections above indicate that 3P sellers can use supply chain transparency to enhance their credibility and build consumer trust in their products. Compared to seller-made claims, consumers appear to deem tracing information more credible and thus willingly apply it to discern a product's quality. This next empirical analysis more directly compares the credibility of blockchain-driven supply chain transparency against sellers' untraced claims about their supply chains.

Before JD made blockchain tracing available, some sellers of ready-to-eat sea cucumber products tried to deliver more supply chain information to their consumers by describing their cold-chain logistics. However, although all ready-to-eat products must be transported via cold-chain logistics and are more handling-sensitive than dried products, not all sellers specified such information.

These seller decisions allow us to compare the effects of such supply chain disclosures against those achieved through the full transparency of blockchain. Essentially, the supply chains of the products that had their cold-chain logistics specified ("specified products") were more transparent than those that had not ("non-specified products"). If consumers believe sellers' claims regarding their own logistics, then the information asymmetry regarding the specified products' SC-D quality should be less than that of the non-specified products.

Similar to the comparison between traced and non-traced products, we estimate the noise parameters (Σ_{ij}^{Δ}) in the pre-treatment period and compare the level of noise a consumer perceives when she tries to assess cold-chain specified products against the non-specified ones. Table 2.9 shows that a 1P seller's claim can reduce signal noise up to 11%, but this effect is much smaller if the claim is made by a 3P seller.

We calculate the effect of blockchain tracing in reducing uncertainty of quality perception in percentages for the ready-to-eat products and compare them against sellers' claims. Note that the content of blockchain tracing information is richer than cold-chain logistics, so this is not quite a direct comparison between the sellers' claims and blockchain tracing. However, we can still shed light on how consumers perceive information from the two sources differently.

Table 2.9 shows that, for 3P ready-to-eat products, blockchain tracing has a large effect in reducing signal uncertainty on product type quality for both experienced and inexperienced consumers, illustrating that consumers find such information credible and use it to verify product type characteristics. This effect is particularly pronounced among inexperienced consumers, who are less familiar with the marketplace and thus have a greater need to identify trustworthy sellers. The uncertainty reduction is much smaller in response to seller-made claims, indicating that consumers remain uncertain about product type quality. Without the same credibility, we see little reduction in uncertainty regarding products' SC-D quality even when cold-chain logistics information is disclosed.

In contrast, for more trusted 1P products, the claims made by sellers can reduce the uncertainty in product type quality to a similar degree as blockchain tracing. This indicates that consumers consider the claims made by 1P sellers to be trustworthy and may use them to assess SC-D quality. As a result, we find that the perceived SC-D quality increased by 9.79% for experienced consumers and 6.27% for inexperienced consumers. This suggests that consumers believe and use 1P sellers' claims to better discern SC-D quality. The fact that blockchain tracing still reduces uncertainty by a larger magnitude may reflect that blockchain tracing provides more extensive information than sellers' typical disclosures regarding their cold-chain logistics.

2.7.4 Impact on Consumer Welfare

Using the estimated model, we proceed to analyze the impact of reducing information asymmetry on consumer welfare. As shown in Section 2.7.2, our findings indicate that consumers, regardless of their levels of experience, can benefit from supply chain transparency. By calculating the consumer surplus for both experienced and inexperienced con-

		Consumer Experience Level			
		Experienced		*Inexperienced*	
		$\sigma^{\Delta}_{1,ij}$	$\sigma^{\Delta}_{2,ij}$	$\sigma^{\Delta}_{1,ij}$	$\sigma^{\Delta}_{2,ij}$
1P	Seller's Claim	11.45	9.79	10.88	6.27
	Blockchain Tracing	11.48	13.51	14.15	6.99
3P	Seller's Claim	3.47	7.71	3.40	-0.12
	Blockchain Tracing	18.74	17.86	24.96	12.53

Table 2.9: Seller's claim is on cold-chain logistics only and blockchain tracing contains logistics and other information. The noise parameters for cold-chain logistics specified products and for the non-specified products are presented separately in Appendix Table 2.18. The percentages here are calculated by taking the difference of the values of the two groups and dividing by the non-specific group. The percentages for blockchain tracing is the same as the ones calculated for ready-to-eat products in Table 2.8.

sumers, we demonstrate in this section that supply chain transparency brings benefits to all consumers, with a particularly significant impact on those with less experience in the marketplace.

Table 2.10 summarizes the impact of reducing information asymmetry on consumer welfare.[3] In the first row of the table, we present the average consumer surplus in a counterfactual scenario where no products are traced, serving as a baseline for comparison. The current adoption of blockchain tracing has led to a 0.98% increase in consumer surplus. This indicates that even with a small portion of products being traced (less than 10% of all products listed on the marketplace when the adoption event took place), there is a noticeable enhancement in the average consumer surplus. In the counterfactual case where all 3P products are traced, an average consumer would experience a 6.08% increase in surplus compared to the baseline. If we expand the tracing to cover all 1P *and* 3P products on the marketplace, the average surplus gain would further rise to 8.41%.

By comparing the surplus gain of tracing all 3P products to the surplus gain of tracing all 1P and 3P products on the marketplace, we observe that the reduction of information

[3]To account for the difference between product quality and a consumer's quality belief, we utilize a method developed by [84] to calculate average consumer surplus. Please refer to Appendix 2.13 for details.

	Consumer Experience Level		
	Experienced	*Inexperienced*	Average
(1) Baseline: no tracing	13.36	9.70	11.90
Percent increase in consumer surplus (%)			
(2) Current adoption of tracing	0.52	1.55	1.01
(3) All 3P products get traced	5.22	7.09	6.04
(4) All products get traced	6.82	10.43	8.41

Table 2.10: The rows (1), (3), and (4) present counterfactual scenarios based on the estimated parameters. Row (1) displays the consumer surplus values in the baseline case where no products are traced. Rows (2), (3), and (4) demonstrate the percentage increase in consumer surplus corresponding to different levels of blockchain tracing adoption. Please refer to Appendix 2.13 for calculation details and Appendix Table 2.19 for the values of average consumer surplus.

asymmetry in the 3P marketplace makes a significant contribution to the overall increase in consumer welfare. This suggests that addressing information asymmetry specifically in the marketplace where consumer trust is relatively weak has a substantial impact on improving consumer welfare.

Furthermore, while experienced consumers tend to have a higher surplus compared to inexperienced consumers, we consistently observe that inexperienced consumers enjoy a larger increase in consumer surplus in all scenarios. Specifically, when all products are traced, the increase in consumer surplus for an average inexperienced consumer is 53% more than an average experienced consumer. This finding suggests that supply chain transparency is effective in bridging the gap between experienced and inexperienced consumers on a marketplace. By providing inexperienced consumers with a new tool to identify trustworthy sellers and assess product information, supply chain transparency significantly enhances their ability to ascertain product quality. As a result, inexperienced consumers can make more informed purchasing decisions, leading to an overall improvement in their welfare. This highlights the transformative potential of blockchain-driven supply chain transparency in empowering consumers and creating a more equitable marketplace experience for all.

On the supply side, the introduction of blockchain tracing has shifted the market towards high-quality products and sellers. With heightened transparency, products of subpar quality might be bought initially, but they won't sustain in the market compared to superior quality counterparts. From our structural model, we estimated the quality distribution of products transacted immediately after the blockchain tracing's implementation. As shown in Table 2.11, products priced above the median but in the bottom quartile of quality sustained only 44% of sales into the first half of 2020. In contrast, 65% of top quartile quality products persisted. Similarly, only 39% of sellers with products in the bottom quality quartile continued sales, compared to 67% in the top quartile. This trend doesn't apply to budget-friendly products, suggesting cost sometimes trumps quality for consumers.

Quality level	Products	Sellers
Top quartile	65.4	67.2
Bottom quartile	44.0	35.5

Table 2.11: The table displays the percentages of products priced above the median and their respective sellers that continued sales into the first half of 2020, which is half a year after the adoption event. Quality distribution is calculated from estimated results as shown in Appendix Table 2.15.

To serve as a benchmark, we also use product rating as a measure of quality to do the same comparison. 2.12 illustrates a narrowed disparity between top and bottom quartile products. Notably, more sellers with top quartile ratings sustained sales into the first half of 2020 than those in the bottom quartile, underscoring the market's shift towards quality under the influence of blockchain tracing.

Rating level	Products	Sellers
Top quartile	56.4	57.3
Bottom quartile	56.6	52.1

Table 2.12: The table displays the percentages of products priced above the median and their respective sellers that continued sales into the first half of 2020.

2.8 Concluding Remarks

Our findings affirm that blockchain-driven tracing technology can broadly benefit consumer marketplaces. Despite the challenges associated with comprehending complex supply chain information, blockchain tracing proves to be a valuable tool for all consumers. By effectively reducing information asymmetry on the marketplace, blockchain induces consumers to more confidently buy traced products that lack outside assurances such as an extensive sales record or JD's own reputational backing for 1P products.

We contribute several key findings. First, for sellers, we empirically substantiate that blockchain tracing especially benefits products and sellers seeking to establish their reputations. These products and sellers either operate within a less trusted environment like the 3P channel or have limited sales history. Although the upstart sellers and products that adopt blockchain do generate additional revenues for themselves in all markets, the implications for their marketplace competitiveness are decidedly mixed. On one hand, established sellers do not benefit from additional credibility in the JD-backed 1P marketplace, which lets adopting upstart sellers and products close the gap in relative demand and market share. On the other hand, the established sellers in the 3P marketplace reap the largest revenue gains of all adopters and widen the revenue gap that separates them from their upstart competitors. Thus, platforms should be aware that while third-party marketplaces see the largest and broadest revenue lifts from blockchain adoption, blockchain-backed supply chain transparency cannot replace the necessary investments that platforms undertake to continually attract and maintain new sellers and products.

Second, the adoption of blockchain tracing quite effectively reduces marketplace information asymmetry. Handling-sensitive, ready-to-eat products can see average revenues increases of 23.4%. Blockchain aids consumers' efforts to assess product quality, by sharpening the precision of quality perceptions by up to 15.6% for 3P products and up to 6.5%

for 1P products. We further find that blockchain traced products' information is deemed to be 26% more credible and informative than sellers' unsupported self-made claims.

Lastly, we find that although experienced consumers shift their demand more strongly towards traced products, inexperienced consumers are the ones who reap the largest gains in welfare from blockchain adoption. Experienced consumers react with greater sophistication, for example by better acting on detailed supply chain information that affects handling-sensitive products. However, inexperienced consumers benefit from increased certainty regarding product quality, and overall blockchain narrows the informational divide between experienced and inexperienced consumers.

We extend our sincere gratitude to the fresh produce sector of JD.com for their invaluable support and collaboration. Profound thanks to Professors Ginger Jin, Andrew Sweetings, and Chenyu Yang for their insightful feedback. We are also grateful to the participants of the 2023 Eighth Marketplace Innovation Workshop, the 2023 POMS Annual Meeting, the 2022 INFORMS Annual Meeting, as well as colleagues at London Business School and University College London for their constructive comments.

Online Supplement for "An Empirical Study of Blockchain-Driven Transparency in a Consumer Marketplace"

2.9 Appendix I: Robust Synthetic Control Method

We use robust synthetic control ("RSC") to analyze the effects of supply chain transparency on product revenue (7, 8, 17). RSC builds on the classical synthetic control method by implementing a novel algorithm for synthesizing control groups. It overcomes the challenges of missing and noisy observations. Moreover, RSC relaxes the convex constraints on the weights in classical synthetic control method, which makes the choices of control variables more flexible and thus generates synthesized units with a better fit. However, the interpretation of results from RSC is no different from the classic synthetic control method. Below, we will discuss how we select proper donor pools of controls, estimate the average treatment effect, and infer its significance.

2.9.1 Donor Pools of Controls

As explained in Section 2.5.2, to overcome any potential bias caused by self-selection into blockchain tracing, we pair each traced product with similar non-traced products. This helps us create a synthetic control group using only units that share the same characteristics as the traced products. Specifically, we group all products based on their product type ty, sales channel ch, price level p, product tenure level t, and rating level r, and use the index g to denote such groups. Note that the grouping is based on levels not values. For example, if a product is sold at a price and has a tenure in the top 50% in its sales channel and product type, it is labeled as 1 for high in price and 1 for long in tenure. Similarly, if a product maintains high ratings of four or more stars during every month of its tenure, it receives the label 1 for being highly rated.

The traced products in a given group g are then matched with a pool of non-traced products from the same group g as the donor pool of controls. For example, if a traced 3P dried-without-salt product has a price and a tenure in the top 50% of all the 3P dried-without-salt products and all its monthly average ratings are above 4, it is labeled as a $(p,t,r) = (1,1,1)$, and will be matched with a pool of non-traced products that are 3P, dried-without-salt, and $(p,t,r) = (1,1,1)$. The matching results are summarized in Tables 2.13 and 2.14.

2.9.2 Average Treatment Effect

The metric used to determine the average treatment effect is the mean monthly revenue, which serves as the outcome variable. The calculation of the average treatment effect of supply chain transparency involves taking the difference between the mean monthly revenue of the treatment group and the synthetic control group and averaging this difference across the post-treatment period.

Let the months that span both the pre-treatment period and the post-treatment period be $t = 1, 2, ...T_0, T_0 + 1, ..., T$, where the pre-treatment period starts from $t = 1$ and ends in T_0, and the post-treatment period starts from $t = T_0 + 1$ and ends in T. For each group g and month t, we calculate the mean of the outcome variable of the treatment group (traced products) and denote it as Y_{gt}. The outcome variable of the synthetic control group is denoted as $\widehat{Y}_{gt}$, which is computed according to the RSC algorithm that follows a two-step procedure: first denoising the data through singular value thresholding, and then learning a relationship between the treated units and the donor pool [8]. For any period of time between t_1 and t_2, $0 \le t_1 \le t_2 \le T$, the average treatment effect of group g is measured by

$$\tau_g(t_1, t_2) = \frac{1}{t_2 - t_1 + 1} \sum_{t=t_1}^{t_2} (\widehat{Y}_{gt} - Y_{gt}).$$

The effect of supply chain transparency of treated units in group g is thus $\tau_g(T_0 + 1, T)$, the average treatment effect in the post-treatment period.

In order to demonstrate the varying effects of supply chain transparency across different product types (ty) and sales channels (ch), we group the data by combining these two factors into (ty, ch) combinations. Suppose a (ty, ch) combination contains W groups, and each group contains ω_g products, we calculate a weighted average of $\tau_g(T_0 + 1, T)$ as the average treatment effect of the traced products:

$$\tau_{ty,ch}(T_0+1,T) = \frac{1}{\sum_{g=1}^{W}\omega_g}\sum_{g=1}^{W}\omega_g\tau_g(T_0+1,T). \tag{2.4}$$

2.9.3 Significance Test

To determine the significance of the treatment effect, we conducted placebo tests. The proposed placebo test is similar in concept to classical permutation inference, which involves creating a distribution of a test statistic by randomly permuting the assignment of sample units to intervention and non-intervention groups. This allows us to test the null hypothesis that the treatment had no effect and determine whether the observed treatment effect is statistically significant. [2]. In particular, To conduct the placebo tests, we randomly select non-traced products and designate them as the treated group (“placebo units”). We then apply the same methodology of the synthetic control group analysis on the placebo units, as described in Section 2.9.2 above, treating them as if they were the actual treatment group. This enables us to compare the actual treatment group’s results to those of the placebo units and test the significance of the treatment effect. A quality of fit measure, which is similar to the average treatment effect measure described earlier, is calculated for both the actual treatment group and the placebo units.

It is expected that only the treated units will experience a significant deviation in the post-treatment period, as the treatment (blockchain tracing) is applied to these units. As a result, the quality of fit measure for a treated unit is expected to be lower than that of the placebo units in the post-treatment period. By comparing the quality of fit of the treated unit and the placebo units, we can infer the significance of the treatment effect of the treated unit. Below, we will explain how the quality of fit measure is constructed.

Following [2] and [1], we use the root mean squared prediction error (RMSPE) method to measure the quality of the fit, which works by comparing the post-treatment fit to the pre-treatment fit. Suppose there are L traced products in a group g. We randomly choose L products from the non-traced products, take them as if they were treated, and calculated the mean outcome variable (a placebo). We repeat the process and form 30 placebo units for each group g. Let $j = 1, ..., J+1$ denote all the units within a group g, where $j = 1$ refers to the treated unit (the traced products), and $j = 2, ..., J+1$, $J = 30$, are all the placebo units. We will treat all these $J+1$ units as the treatment groups, and perform RSC on each of them one at a time. If the treatment effect is significant, we should observe a substantial difference between the outcome variable of the actual treated unit and the synthetic control group, while there should be no such difference for the placebo units. In other words, we should observe a statistically significant difference only for $j = 1$, but not for $j = 2, ..., J+1$. This approach allows us to determine whether the observed treatment effect is genuine or simply a result of chance or randomness.

Borrowing the notation from [1], for $0 \leq t_1 \leq t_2 \leq T$ and $j = 1, ..., J+1$, let

$$R_j(t_1, t_2) = \left(\frac{1}{t_2 - t_1 + 1} \sum_{t=t_1}^{t_2} \left(Y_{jt} - \widehat{Y}_{jt} \right)^2 \right)^{1/2}, \tag{2.5}$$

where Y_{jt} is the outcome variable of the treatment group on time t, and $\widehat{Y}_{jt}$ is that of the synthetic control. Equation 2.5 defines RMSPE for a unit j and a time period between t_1 and t_2. The ratio between the post-treatment (starting from month $T_0 + 1$ to month T) RMSPE to the pre-treatment (starting from month 1 to month T_0) RMSPE for unit j is thus

$$r_j = \frac{R_j(T_0 + 1, T)}{R_j(1, T_0)},$$

where r_j measures the quality of the fit in the post-treatment period relative to the pre-treatment period of unit j.

From this procedure, we obtain r_j for the treated unit and each of the placebo units for every group g. Thus, a p-value can be calculated from the distribution of r_j by

$$p = \frac{1}{J+1} \sum_{j=1}^{J+1} I_{+}(r_j - r_1), \tag{2.6}$$

where $I_{+}(\cdot)$ is an indicator function that equals to 1 if the argument inside is non-negative, and 0 otherwise. When aggregating the groups to (ty, ch) combinations, we choose the group that has the largest p value (the lowest significance) to infer the significance level. By selecting the largest p value among the groups, we are essentially taking a conservative approach that assumes that the effect of supply chain transparency is not significant unless there is strong evidence to suggest otherwise.

2.10 Appendix II: Bayesian Updating Process

Consumers value their consumption of products based on their beliefs on the products' quality. Consumers do not observe quality directly; their beliefs on a product's quality are formed based on the product's visible attributes, which are correlated with the hidden quality. Additionally, consumers take into account the sales channel through which a product is sold, which can provide insight into the platform's assessment of the product's quality. Finally, consumers also use online product assortments to obtain private signals about the quality of each product, such as reviews and ratings from other customers.

With the foregoing sources of information, a consumer i forms her final belief on the quality level of a product $j \in \mathscr{J}$ on the marketplace as a two-stage Bayesian updating process. In the first stage, a consumer $i \in \mathscr{N}$ observes the attributes A_j of a product j and its sales channel classification c_j, where $c_j = 1$ if j is a 1P product or $c_j = 0$ if it is a 3P product. From attributes A_j and sales channel c_j, a consumer i forms her first-stage posterior belief Q_j^{public}, which is common among all consumers. In the second stage, each consumer i receives a private signal about quality of a specific product (q_{ij}). The consumer then uses such quality signal to updates her belief and forms the second stage posterior belief on quality (Q_{ij}^{post}). This process generates consumers with heterogeneous posterior beliefs on the same product.

2.10.1 First-Stage Update on Quality Belief

In the first stage, a consumer updates her beliefs on quality based on the product's attributes and sales channel (1P or 3P), both are observable by all consumers arriving at the marketplace. After this first-stage update, all consumers will share a common belief Q_j^{public} about the quality of a product j with certain attributes and sales channel. In the following, we will first describe the distribution of products on each sales channel, and then discuss how a consumer updates her belief about product quality based on the product's channel and attributes.

A product $j \in \mathscr{J}$ is modeled as a multivariate normal random variable $X_j = [Q_j^1, Q_j^2, A_j^1, \ldots, A_j^K]^\top$. To account for the latent nature of quality, we expand the product X_j to the extended space $\tilde{X}_j = (X_j, \nu_j)$, where ν_j is distributed independently from X_j as $N \sim (0, \sigma_\nu)$. This error term reflects the uncertainty or variation that exists in the process of quality assessment and helps to capture the unobservable aspects of quality that cannot be directly measured or observed. We define an extended quality $\tilde{Q}_j = (Q_j, \nu_j)$ to encompass the two unobserved elements in a product X_j. $\tilde{Q}_j$ is distributed trivariate normal $(\tilde{\mu}_Q, \tilde{\Sigma}_Q)$, where $\tilde{\mu}_Q = [\mu_Q, 0]^\top$ and $\tilde{\Sigma}_Q = [[\Sigma_Q, 0], [0, \sigma_\nu]]$. Together, A product is distributed on the extended space $\tilde{X}_j = (\tilde{Q}_j, A_j)$, where $\tilde{Q}_j$ is latent and A_j can be observed. $\tilde{X}_j$ describes the pool of all the products despite their sales channel, which is distributed multivariate normal

$N(\tilde{\mu}, \tilde{\Psi})$ with

$$\tilde{\mu} = \begin{bmatrix} \tilde{\mu}_Q \\ \mu_A \end{bmatrix} \text{ and } \tilde{\Psi} = \begin{bmatrix} \tilde{\Sigma}_Q & \tilde{\Sigma}_{QA} \\ \tilde{\Sigma}_{AQ} & \Sigma_A \end{bmatrix}. \tag{2.7}$$

Given the observed product attributes $A_j = a_j$, the consumer's rationally updated quality beliefs $\tilde{Q}_j \mid A_j = a_j$ are the conditional distribution that is multivariate normal $N(\tilde{\mu}_0^{a_j}, \tilde{\Sigma}_0)$ with

$$\begin{aligned} \tilde{\mu}_0^{a_j} &= \tilde{\mu}_Q + \tilde{\Sigma}_{QA}\Sigma_A^{-1}(a_j - \mu_A), \\ \tilde{\Sigma}_0 &= \tilde{\Sigma}_Q - \tilde{\Sigma}_{QA}\Sigma_A^{-1}\tilde{\Sigma}_{AQ}. \end{aligned} \tag{2.8}$$

We use $f_0(Q_j)$ to denote the density function of this distribution.

In addition to the attributes A_j, the consumers observe the sales channel of each product c_j, which is the result of the platform's classification. The platform selects for the 1P market the products that satisfy the threshold-based classifier $\Lambda X_j + \nu_j \geq r$. That is,

$$c_j = \begin{cases} 1 & \text{if } \Lambda_Q Q_j + \Lambda_A A_j + \nu_j \geq r \\ 0 & \text{if } \Lambda_Q Q_j + \Lambda_A A_j + \nu_j < r, \end{cases} \tag{2.9}$$

where $\Lambda = \{\Lambda_Q, \Lambda_A\}$.

Such a classification (2.9) truncates the multivariate normal distribution of $\tilde{Q}_j \mid A_j \sim (\mu_0^{a_j}, \Sigma_0)$ to two truncated multivariate normal distributions, $\tilde{Q}_j \mid A_j, c_j \sim TN(\tilde{\mu}_0^{a_j}, \tilde{\Sigma}_0, S^{1P})$ or $TN(\tilde{\mu}_0^{a_j}, \tilde{\Sigma}_0, S^{3P})$, where S^{1P} denotes the support of the quality distribution of 1P products and S^{3P} denotes that of 3P products, with

$$\begin{aligned} S^{1P} &= \{\tilde{\Lambda}_Q \tilde{Q}_j \geq r - \Lambda_A a_j\} \\ S^{3P} &= \{\tilde{\Lambda}_Q \tilde{Q}_j < r - \Lambda_A a_j\}, \end{aligned} \tag{2.10}$$

where $\tilde{\Lambda}_Q = [\Lambda_Q, 1]$ and $\tilde{Q}_j = [Q_j, \nu_j]^\top$. We use $f_0^{ch}(Q_j)$, $ch = 1P$ or $3P$, to refer to the density function of $\tilde{Q}_j \mid A_j, c_j \sim TN(\tilde{\mu}_0^{a_j}, \tilde{\Sigma}_0, S^{ch})$. Note that $f_0^{ch}(Q_j)$ is proportional to $f_0(Q_j)$ on its corresponding support, i.e.,

$$f_0^{ch}(Q_j) = \begin{cases} C f_0(Q_j) & \text{if } Q_j \in S^{ch} \\ 0 & \text{otherwise,} \end{cases} \tag{2.11}$$

where C is a normalization constant.

Through the process described above, consumers arrive at their first-stage posterior belief on quality $Q_j^{public} = Q_j \mid A_j, c_j$, which can be obtained by integrating out the idiosyncratic term ν in $\tilde{Q}_j \mid A_j, c_j$:

$$Q_j^{public} = \int f_0^{ch}(Q_j) d\nu. \tag{2.12}$$

2.10.2 Second-Stage Update on Quality Belief

In the second stage, a consumer builds on her belief on quality and proceeds to collect information about the quality of this product j and updates her belief again. Once this second-stage update completes, a consumer i will arrive at her final posterior belief on the quality of product j, denoted as Q_{ij}^{post}, which is her private information about the quality of product j.

In the end of the first stage, all consumers have a common belief on the unobserved extended quality of a product $\tilde{Q}_j \mid A_j, c_j \sim TN(\tilde{\mu}_0^{a_j}, \tilde{\Sigma}_0, S^{ch})$, which depends on the product's observable attributes and sales channel classification. After collecting and inspecting the information regarding product j, a consumer i receives a private signal $\tilde{q}_{ij} = (q_{ij}, q_\nu)$ on the latent components of a product, i.e., the two dimensions of quality and the error term ν that captures the uncertainty in assessing quality. This quality signal $\tilde{q}_{ij}$ is centered at the product's true quality but incorporates a noise term $\tilde{\Delta}$, i.e.,

$$\tilde{q}_{ij} = \tilde{Q}_j + \tilde{\Delta}_{ij},$$

where $\tilde{\Delta}_{ij} \sim N(0, \tilde{\Sigma}_{ij}^{\Delta})$, a trivariate normal distribution. We describe probability density of the signal $\tilde{q}_{ij}$ with the likelihood function $g(q_{ij}|Q_j)$.

Sampling from signals with such likelihood function $g(q_{ij}|Q_j)$, a customer i updates her first-stage posterior density function on quality $f_0^{ch}(Q_j)$ and forms a second-stage posterior belief $\tilde{Q}_j \mid A_j, c_j, \tilde{q}_{ij}$. According to the Bayes rule, the density function of this posterior belief, denoted as $f_1^{ch}(Q_j)$, is

$$f_1^{ch}(Q_j) \propto g(q_{ij}|Q_j) f_0^{ch}(Q_j). \tag{2.13}$$

To better characterize the distribution of this second-stage posterior belief of quality $\tilde{Q}_j \mid A_j, c_j, \tilde{q}_{ij}$ with the density function (2.13), note that $f_0^{ch}(Q_j)$ is proportional to $f_0(Q_j)$ on the support of the sales channel product j belongs to as shown in (2.11). If we omit the sales channel classification in the first-stage update, we obtain $\tilde{Q}_j \mid A_j$ with density function $f_0(Q_j)$ in the end of the first-stage. In the second stage, the same signal with likelihood function $g(q_{ij}|Q_j)$ is used to update a consumer's belief. Then the second-stage posterior belief is $\tilde{Q}_j \mid A_j, q_{ij}$ with density function $f_1(Q_j)$), where

$$f_1(Q_j) \propto g(q_{ij}|Q_j) f_0(Q_j). \tag{2.14}$$

The distribution of $\tilde{Q}_j \mid A_j, q_{ij}$ is distributed multivariate normal $N(\mu_{1,ij}, \Sigma_{1,ij})$, with

$$\begin{aligned} \mu_{1,ij} &= \frac{\tilde{\Sigma}_0}{\tilde{\Sigma}_0 + \tilde{\Sigma}_{ij}^{\Delta}} \tilde{q}_{ij} + \frac{\tilde{\Sigma}_{ij}^{\Delta}}{\tilde{\Sigma}_0 + \tilde{\Sigma}_{ij}^{\Delta}} \tilde{\mu}_0^{a_j}, \\ \Sigma_{1,ij} &= \frac{\tilde{\Sigma}_{ij}^{\Delta}}{\tilde{\Sigma}_0 + \tilde{\Sigma}_{ij}^{\Delta}} \tilde{\Sigma}_0. \end{aligned} \tag{2.15}$$

Combining (2.11), (2.13), and (2.14), we observe the following proportional relationship:

$$f_1^{ch}(Q_j) \propto g(q_{ij}|Q_j) f_0^{ch}(Q_j) \propto g(q_{ij}|Q_j) f_0(Q_j) \propto f_1(Q_j) \quad \forall Q_j \in S^{ch}. \tag{2.16}$$

The relationship (2.16) shows that, the density function of the second-stage posterior belief $\tilde{Q}_j \mid A_j, c_j, q_{ij}$ is proportional to that of $\tilde{Q}_j \mid A_j, q_{ij}$ on the support of the sales channel product j classified to, and we know $\tilde{Q}_j \mid A_j, q_{ij}$ is distributed $N(\mu_{1,ij}, \Sigma_{1,ij})$. In intuitive terms, (2.16) implies that upon receiving the platform's classification, a consumer's belief about the quality of a product is constrained to a distribution specific to the sales channel, and any subsequent adjustments to her quality beliefs are also limited by this same range of values.

Finally, after characterizing $\tilde{Q}_j \mid A_j, c_j, q_{ij}$, a consumer's final posterior belief on product quality $Q_{ij}^{post} = Q_j \mid A_j, c_j, q_{ij}$ can be obtained by integrating out the ν in $\tilde{Q}_j = (Q_j, \nu)$, i.e.,

$$Q_{ij}^{post} = \int f_1^{ch}(Q_j) d\nu. \tag{2.17}$$

2.11 Appendix III: Structural Estimation

By estimating the structural model described in Section 2.6, our objective is to learn the signal noise Δ_{ij}, the joint distribution of products' quality levels Q_j and attributes A_j that enables consumers to learn about quality, and the distribution of consumers' preferences β_i and price elasticities α_i that together drive their purchasing decisions. Because neither Δ_{ij}, Q_j, β_i, nor α_i is directly observable to us as empirical researchers, we must instead infer them from the available data using our structural model. We will begin by introducing the maximum likelihood model, followed by a discussion of the EM algorithm which we use to estimate the parameters of the model.

2.11.1 Maximum Likelihood Estimation

From the dataset, we observe each consumer i's choice of product $x_i = j^*$, as well as its visible attributes A_j and its price P_j. We also observe whether each product j is classified into the 1P ($c_j = 1$) or 3P ($c_j = 0$) sales channel. Using a maximum likelihood model, we aim at estimating the parameters that maximize the probability of consumers making observed product choices and the platform accurately classifying products to observed sales channels.

Both each consumer's choice of product and the platform's classification of sales channel reveal the latent product quality. In particular, a product's quality levels Q_j are correlated with its visible attributes A_j; to consumers, the attributes therefore partially reveal products' hidden quality levels. Moreover, because the products are sorted into sales channels based on the platform's knowledge and assessment of products' quality levels Q_j and attributes A_j, such classifications are informative to consumers about the products' underlying quality. Lastly, by inspecting the product assortment (online), consumers obtain private signals about the quality levels of each product, and the precision of such private signals depends both on consumers' level of experience, e_i, which we observe, and on whether a given product is blockchain-traced. Instead of directly reflecting the true quality levels Q_j, consumers' purchasing decisions hinge on the beliefs about Q_j that they derive from the foregoing sources of information. In order to study the operations of the marketplace under information asymmetry, we model the learning process that both generates consumers' beliefs and the decisions that they make in response.

Let θ represent the model's parameters. They are partitioned into: θ

N, which governs the distribution of the preference parameters β_i, α_i in the consumer population; $\theta_{\mathscr{J}}$, which governs the joint distribution of product quality Q_j and attributes A_j in the product category; θ_Δ, which governs the distribution of noisiness in consumers' private signals regarding a product j's quality, Δ_{ij}, based on whether the shoppers are experienced and on whether the product is blockchain traced; and θ

P, which specifies the platform's classification rule for selectively sorting products into its 1P and 3P marketplaces. We aim to estimate θ.

Given the observed purchasing choices, $\{x_i : i \in \mathscr{N}\}$, and product classifications $\{c_j : j \in \mathscr{J}\}$, we obtain the maximum likelihood estimates

$$
\begin{aligned}
\theta^*_{MLE} :=&_\theta L\Big(\{x_i : i \in \mathscr{N}\}, \{c_j : j \in \mathscr{J}\} \mid \{A_j : j \in \mathscr{J}\}, \{e_i : i \in \mathscr{N}\}; \theta\Big) \\
=&_\theta \prod_{i \in \mathscr{N}} p\Big(x_i \,|\, \{(c_j, A_j) : j \in \mathscr{J}\}, e_i; \theta\Big) \cdot \prod_{j \in \mathscr{J}} p(c_j | A_j; \theta
\end{aligned}
$$

P).

The modeled likelihood explains the purchase decisions $\{x_i : i \in \mathscr{N}\}$ and the product classifications $\{c_j : j \in \mathscr{J}\}$ while taking the products' attributes $\{A_j : j \in \mathscr{J}\}$ and consumers' experience levels with the product category $\{e_i : i \in \mathscr{N}\}$ as given. Here, $p\Big(x_i \,|\, \{(c_j, A_j) : j \in \mathscr{J}\}, e_i; \theta\Big)$ represents the probability that, under θ, a consumer i of the observed experience level e_i makes the choice x_i given the products' observed attributes, $\{A_j : j \in \mathscr{J}\}$ and their marketplace classifications, $\{c_j : j \in \mathscr{J}\}$. Next, $p(c_j \,|\, A_j; \theta$

P) is the probability that, under the platform classification rules θ

P, a product j of the observed attributes A_j receives the marketplace classification c_j.

Under the model, we further expand the above choice-based likelihood into four underlying components as follows:

$$\begin{aligned}\theta^*_{MLE} =_\theta \int \prod_{i \in \mathscr{N}} p(x_i \mid \mathbb{E}Q^{Post}_{ij}, j \in \mathscr{J}, \beta_i, \alpha_i) \cdot \prod_{i \in \mathscr{N}, j \in \mathscr{J}} p(\mathbb{E}Q^{Post}_{ij} \mid Q_j, A_j, c_j, e_i; \theta_{\mathscr{J}}, \theta_{\mathscr{P}}, \theta_\Delta) \\ \cdot \prod_{j \in \mathscr{J}} \Big[p(c_j \mid Q_j, A_j; \theta_{\mathscr{P}}) \cdot p(Q_j \mid A_j; \theta_{\mathscr{J}}) \Big] \cdot \prod_{i \in \mathscr{N}} p(\beta_i, \alpha_i \mid e_i; \theta_{\mathscr{N}}) \\ \cdot d\Delta_{1,1} \cdots d\Delta_{|\mathscr{N}|, |\mathscr{J}|} \cdot dQ_1 \cdots dQ_{|\mathscr{J}|} \cdot d\beta_1 \cdots d\beta_{|\mathscr{N}|} \cdot d\alpha_1 \cdots d\alpha_{|\mathscr{N}|}. \qquad (2.19)\end{aligned}$$

The components involve latent variables, such as product quality, that the data cannot report on directly but that drive consumers' utility-based purchasing decisions. We explain how each component is calculated.

Choice probabilities.

The first component $p(x_i \,|\, \beta_i, \alpha_i, \mathbb{E}Q^{Post}_{ij}, j \in \mathscr{J})$ that the model specifies is consumer i's choice probability based on her preferences and the quality levels of product j expected under her fully updated beliefs:

$$p(x_i = j^* \mid \beta_i, \alpha_i, \mathbb{E}Q_{ij}^{Post}, j \in \mathscr{J}) = \frac{\exp\left(\beta_i \mathbb{E}Q_{ij^*}^{post} - \alpha_i P_{j^*}\right)}{1 + \sum_j \exp\left(\beta_i \mathbb{E}Q_{ij}^{post} - \alpha_i P_j\right)}. \tag{2.20}$$

As a conditional probability, it takes as given the expected quality levels $\mathbb{E}Q_{ij}^{Post}$ that consumers have learned.

Beliefs under consumer learning.

The second component $p(\mathbb{E}Q_{ij}^{Post} \mid Q_j, A_j, c_j, e_i; \theta_{\mathscr{J}}, \theta_{\mathscr{P}}, \theta_{\Delta})$ captures the consumer learning process that produces a consumer's expected quality for product j, $\mathbb{E}Q_{ij}^{Post}$, at the time of her purchase decision. Whereas Q_j denotes product j's true quality, we use the notation Q'_{ij} to denote consumer i's beliefs regarding a product j's quality in the form of a random variable that manifests her beliefs as a probability distribution over the possible quality levels. We use the notation $Q'_{ij} \mid O$ to denote the consumer's beliefs regarding the same when they are updated by conditioning on the set of observed variables O.

To arrive at her final posterior belief Q_{ij}^{Post} regarding product j's quality, a consumer i incorporates three pieces of information: namely, the product j's visible attributes A_j, its observed platform classification c_j, and the private signal q_{ij} that she obtains upon inspecting the product online:

$$Q_{ij}^{Post} := Q'_{ij} \mid A_j, q_{ij}, c_j, e_i. \tag{2.21}$$

The private signal is specified with the individually idiosyncratic additive noise Δ_{ij}, such that $q_{ij} = Q_j + \Delta_{ij}$. The magnitude of the noise depends on whether the product is blockchain-traced and on the consumer's experience level e_i. Given the latter dependence, it is appropriate to condition on e_i when defining the final beliefs Q_{ij}^{Post}: the degree of noise within her private signal q_{ij} influences how strongly consumer i should use it to update her beliefs Q'_{ij}. Aside from the private signal, the beliefs Q_{ij}^{Post} incorporate an update from A_j based on the joint distribution of Q_j and A_j governed by $\theta_{\mathscr{J}}$ and an update from c_j based on the classification rule $\theta_{\mathscr{P}}$.

More explicitly, we calculate the expected quality levels $\mathbb{E}Q_{ij}^{Post}$ in three steps that each incrementally update the beliefs Q'_{ij}. In step one, we obtain the update $Q'_{ij} \mid A_j$ based on the jointly multivariate normal distribution $N(\mu, \Psi)$ of the products' quality Q_j and attributes A_j, where $\mu = [\mu_Q, \mu_A]^\top$ and $\Psi = [[\Sigma_Q, \Sigma_{QA}], [\Sigma_{AQ}, \Sigma_A]]$. (Thus, the vector of distributional parameters for the product characteristics, $\theta_{\mathscr{J}}$, comprises $\{\mu, \Psi\}$.) Given the observed product attributes $A_j = a_j$, the consumer's rationally updated quality beliefs $Q'_{ij} \mid A_j = a_j$ are given by the conditional distribution that is multivariate normal $N(\mu_0^{a_j}, \Sigma_0)$ with

$$\begin{aligned} \mu_0^{a_j} &= \mu_Q + \Sigma_{QA}\Sigma_A^{-1}(a_j - \mu_A), \\ \Sigma_0 &= \Sigma_Q - \Sigma_{QA}\Sigma_A^{-1}\Sigma_{AQ}. \end{aligned} \tag{2.22}$$

Note that the mean of this posterior belief $\mu_0^{a_j}$ is product-specific because it depends on the product's observed attributes a_j. On the other hand, the covariance Σ_0 is common to all products.

In step two, we obtain the further updated beliefs $Q'_{ij} \mid A_j, q_{ij}, \Sigma_{ij}^{\Delta}$ that incorporate the consumer's private quality signal q_{ij}. The signal q_{ij} is centered at the product's true quality Q_j with the addition of mean-zero, idiosyncratic noise Δ_{ij} that is distributed $N(0, \Sigma_{ij}^{\Delta})$. The magnitude of Σ_{ij}^{Δ} describes the noisiness of the signal the consumer receives, which depends both on her ability to understand the signal and on the quality of the information that produced it. Thus, we allow Σ_{ij}^{Δ} to differ based on consumer i's experience level e_i, product j's channel c_j and type ty_j (a dried-without-salt, a ready-to-eat, or a non-premium product), and whether this product is traced or not tr_j. Together, there are 24 Σ_{ij}^{Δ}'s for all the combinations of (e_i, c_j, ty_j, tr_j).

The distribution of the conditional beliefs $Q'_{ij} \mid A_j, q_{ij}, \Sigma^{\Delta}_{ij}$ is distributed multivariate normal $N(\mu_{1,ij}, \Sigma_{1,ij})$, with the following mean and variance:

$$\begin{aligned}
\mu_{1,ij} &= \frac{\Sigma_0}{\Sigma_0 + \Sigma^{\Delta}_{ij}} q_{ij} + \frac{\Sigma^{\Delta}_{ij}}{\Sigma_0 + \Sigma^{\Delta}_{ij}} \mu_0^{a_j}, \\
\Sigma_{1,ij} &= \frac{\Sigma^{\Delta}_{ij}}{\Sigma_0 + \Sigma^{\Delta}_{ij}} \Sigma_0,
\end{aligned} \tag{2.23}$$

where $\mu_0^{a_j}$ and Σ_0 are given in (2.22), and Σ^{Δ}_{ij} is the covariance of signal noise Δ_{ij}.

Finally, in step three, we obtain the last update

$$Q'_{ij} \mid A_j, q_{ij}, \Sigma^{\Delta}_{ij}, c_j = Q'_{ij} \mid A_j, q_{ij}, \Sigma^{\Delta}_{ij}, c_j, e_i \tag{2.24}$$

$$= Q'_{ij} \mid A_j, q_{ij}, c_j, e_i \tag{2.25}$$

$$= Q^{Post}_{ij}, \tag{2.26}$$

which incorporates the platform's choice c_j of the sales channel of product j, such that $c_j = 1$ for 1P and $c_j = 0$ for 3P. Under the selection parameters $\theta_{\mathscr{P}} = \{\Lambda_Q, \Lambda_A, r\}$, the platform selects for the 1P market the products that satisfy the threshold-based classifier $\Lambda_Q Q_j + \Lambda_A A_j + \nu_j \geq r$. That is,

$$c_j = \begin{cases} 1 & \text{if } \Lambda_Q Q_j + \Lambda_A A_j + \nu_j \geq r \\ 0 & \text{if } \Lambda_Q Q_j + \Lambda_A A_j + \nu_j < r. \end{cases} \tag{2.27}$$

The idiosyncratic component ν_j is normalized to be drawn from the standard normal distribution.

For a product j sold in the 1P marketplace, we calculate the expected values of its quality levels under consumer i's final posterior beliefs, $\mathbb{E}Q_{ij}^{post}$, as follows:

$$\mathbb{E}Q_{ij}^{post} = \mathbb{E}(Q'_{ij} \mid A_j, q_{ij}, e_i, c_j = 1) = \frac{\mathbb{E}(Q'_{ij} 1_{\{c_j=1\}} \mid A_j, q_{ij}, e_i)}{\mathbf{Pr}(c_j = 1 \mid A_j, q_{ij}, e_i)} \tag{2.28}$$

$$= \frac{\mathbb{E}(Q'_{ij} 1_{\{\Lambda_Q Q'_{ij} + \Lambda_A A_j + \nu \geq r\}} \mid A_j, q_{ij}, e_i)}{\mathbf{Pr}(\Lambda_Q Q'_{ij} + \Lambda_A A_j + \nu \geq r \mid A_j, q_{ij}, e_i)} \tag{2.29}$$

$$= \frac{\int_\nu \int_x x f_{Q'|A,q,e}(x) \phi(\nu) 1_{\{\nu \geq r - \Lambda_Q x - \Lambda_A A_j\}} dx d\nu}{\int_\nu \int_x f_{Q'|A,q,e}(x) \phi(\nu) 1_{\{\nu \geq r - \Lambda_Q x - \Lambda_A A_j\}} dx d\nu} \tag{2.30}$$

$$= \frac{\int_x x f_{Q'|A,q,e}(x) \left[1 - \Phi\left(r - \Lambda_Q x - \Lambda_A A_j\right)\right] dx}{\int_x f_{Q'|A,q,e}(x) \left[1 - \Phi\left(r - \Lambda_Q x - \Lambda_A A_j\right)\right] dx}, \tag{2.31}$$

where $f_{Q'|A,q,e}(\cdot)$ denotes the probability density function of $Q'_{ij} \mid A_j, q_{ij}, \Sigma_{ij}^{\Delta}$, which is multivariate normal with mean and variance given in 2.23, and $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal distribution's probability density and cumulative density functions, respectively.

Similarly, for a product j sold in the 3P marketplace, the expected values of its quality levels under consumer i's final posterior beliefs, $\mathbb{E}Q_{ij}^{post}$, are given by

$$\mathbb{E}Q_{ij}^{post} = \mathbb{E}(Q'_{ij} \mid A_j, q_{ij}, e_i, c_j = 0) = \frac{\mathbb{E}(Q'_{ij} 1_{\{c_j=0\}} \mid A_j, q_{ij}, e_i)}{\mathbf{Pr}(c_j = 0 \mid A_j, q_{ij}, e_i)} \tag{2.32}$$

$$= \frac{\int_x x f_{Q'|A,q,e}(x) \Phi\left(r - \Lambda_Q x - \Lambda_A A_j\right) dx}{\int_x f_{Q'|A,q,e}(x) \Phi\left(r - \Lambda_Q x - \Lambda_A A_j\right) dx}. \tag{2.33}$$

Both (2.31) and (2.33) must be numerically evaluated, and we relegate the details to Section 2.11.3.

To close, we return to the calculation of the second component, $p(\mathbb{E}Q_{ij}^{Post} \mid Q_j, c_j, A_j, e_i; \theta_{\mathscr{J}}, \theta_{\mathscr{P}}, \theta_\Delta)$. By inspecting (2.23), (2.31), and (2.33), it is easily verified that for any fixed values of A_j, e_i, and c_j, the expected quality levels $\mathbb{E}Q_{ij}^{post}$ under the consumer's final beliefs are increasing (element-wise) in the (corresponding elements of) the private signal, q_{ij}. Then, for any fixed values of A_j, Q_j, e_i, and c_j, the expected quality levels $\mathbb{E}Q_{ij}^{post}$ under the consumer's final beliefs are similarly monotonically increasing (element-wise) in the (corresponding elements of) the private signal's noise, Δ_{ij}. Thus, for any fixed values of A_j, Q_j, e_i, and c_j, we can define the inverse function $\Delta_{A,Q,c,e}$ that maps given values of $\mathbb{E}Q_{ij}^{post}$ back to the unique values of signal noise Δ_{ij} that can produce them.

Thus, treating $\Delta_{A,Q,c,e}(x)$ as a fixed quantity, the calculation is given by

$$p(\mathbb{E}Q_{ij}^{Post} = x \mid Q_j, A_j, c_j, e_i; \theta_{\mathscr{J}}, \theta_{\mathscr{P}}, \theta_\Delta) = p(\Delta_{A,Q,c,e}(x) \mid Q_j, A_j, c_j, e_i; \theta_{\mathscr{J}}, \theta_{\mathscr{P}}, \theta_\Delta) \tag{2.34}$$

$$= \frac{\exp\left(-\frac{1}{2}\Delta_{A,Q,c,e}(x)^\top (\Sigma_{ij}^\Delta)^{-1} \Delta_{A,Q,c,e}(x)\right)}{\sqrt{(2\pi)^2 |\Sigma_{ij}^\Delta|}}. \tag{2.35}$$

Product characteristics.

The third component $p(c_j \mid Q_j, A_j; \theta_{\mathscr{P}}) \cdot p(Q_j \mid A_j; \theta_{\mathscr{J}})$ addresses platform classification and product quality. The first subcomponent, $p(c_j \mid Q_j, A_j; \theta_{\mathscr{P}})$, is the modeled probability that product j receives its actually observed marketplace classification c_j conditional on its quality and attributes. The platform classification parameters $\theta_{\mathscr{P}} = \{\Lambda_Q, \Lambda_A, r\}$ accounts for such quality and attributes and provides for a normalized idiosyncratic component ν_j. The resulting marketplace classification probabilities are

$$p(c_j \mid Q_j, A_j; \theta_{\mathscr{P}}) = \begin{cases} 1 - \Phi(r - \Lambda_Q Q_j - \Lambda_A A_j) & \text{if } c_j = 1 \\ \Phi(r - \Lambda_Q Q_j - \Lambda_A A_j) & \text{if } c_j = 0. \end{cases} \tag{2.36}$$

Straightforwardly in comparison, the second subcomponent is the simple conditional probability of the product quality Q_j based on its joint distribution with the given product attributes A_j:

$$p(Q_j = x \mid A_j; \theta_{\mathscr{J}}) = \frac{\exp\left(-\frac{1}{2}(x - \mu_0^{a_j})^\top \Sigma_0{}^{-1}(x - \mu_0^{a_j})\right)}{\sqrt{(2\pi)^2 |\Sigma_0|}}, \tag{2.37}$$

where the mean $\mu_0^{a_j}$ and covariance Σ_0 are given in (2.22).

Consumer preferences.

The final component $p(\beta_i, \alpha_i \mid e_i, \theta_{\mathscr{N}})$ completes the likelihood calculation by including the probability of the individual preference parameters, β_i, α_i, where $\beta_i \in \mathbb{R}^2$ is the utility coefficient for the two-dimensional expected posterior belief of quality $\mathbb{E}Q_{ij}^{post}$, and $\alpha_i \in \mathbb{R}$ is the price elasticity. The preference parameters are distributed multivariate normal $N(\mu_i^N, \Sigma_i^N)$ and are allowed to differ based on consumer i's experience level e_i. Let $\theta_{\mathscr{N}} = \{\mu_i^N, \Sigma_i^N\}$, this probability is

$$p([\beta_i, \alpha_i]^\top = x \mid \theta_{\mathscr{N}}) = \frac{\exp\left(-\frac{1}{2}(x - \mu_i^N)^\top (\Sigma_i^N)^{-1}(x - \mu_i^N)\right)}{\sqrt{(2\pi)^3 |\Sigma_i^N|}}. \tag{2.38}$$

2.11.2 EM Algorithm

In order to compute the maximum likelihood estimator θ_{MLE}^* in 2.19, we utilize the EM algorithm. This algorithm is an iterative procedure that is a computational alternative to directly maximizing the likelihood function. It is especially useful when the likelihood function integrates over a large set of latent variables that are used by the structural model to explain observed choices.

The integrand of expression (2.19) gives the likelihood jointly in terms of the observed outcomes $\{x_i : i \in \mathscr{N}\}, \{c_j : j \in \mathscr{J}\}$ and the model's underlying latent variables, wherein the latent variables are subsequently integrated out to obtain the likelihood function without latent variables. The EM algorithm makes direct use of exactly this integrand, which is simply the joint likelihood of the data $\{x_i : i \in \mathscr{N}\}, \{c_j : j \in \mathscr{J}\}$ as observed outcomes and the latent variables $\{\mathbb{E}Q_{ij}^{Post}, Q_j, \beta_i, \alpha_i : i \in \mathscr{N}, j \in \mathscr{J}\}$ as unseen outcomes.

In overviewing the algorithm, let us denote the data $\{x_i : i \in \mathscr{N}\}, \{c_j : j \in \mathscr{J}\}$ collectively as O and the latent variables $\{\mathbb{E}Q_{ij}^{Post}, Q_j, \beta_i, \alpha_i : i \in \mathscr{N}, j \in \mathscr{J}\}$ as Y. Further let us use $L(O, Y \mid \theta)$ to mean their joint likelihood function (i.e., the integrand of (2.19)) evaluated at the parameter values θ; for simplicity, we suppress that the joint likelihood is actually defined conditional on the product attributes and consumer experience as exogenously reported by the data—i.e., that it is more formally $L(O, Y \mid \{A_j : j \in \mathscr{J}\}, \{e_i : i \in \mathscr{N}\}; \theta)$.

Each iteration t of the EM algorithm updates a current parameter value that we label as θ_t. Specifically,

$$\theta_{t+1} =_\theta \mathbb{E}_{Y|O,\theta_t}[\log L(O, Y \mid \theta)]. \tag{2.39}$$

Thus, each round of the algorithm chooses the next parameter value θ_{t+1} to maximize a specially constructed expectation. Typically, in practice, the expectation is evaluated through simulation. In this case,

$$\theta_{t+1} =_\theta \sum_{b=1}^{B} \log L(O, Y_b \mid \theta), \tag{2.40}$$

where the values of Y_b, $b = 1, \ldots, B$, are drawn, by simulation, from the posterior distribution of the latent variables Y conditional on O and θ_t. Then, in order to compute θ_{t+1}, the EM algorithm first performs an E-step that samples such Y_b, $b = 1, \ldots, B$ and constructs the objective function. The M-step then performs the maximization. In early iterations of the EM algorithm, θ_{t+1} is often chosen to improve on the objective value over θ_t (instead of fully maximizing the objective), as this is sufficient to guarantee an improvement in the associated likelihood value and accelerates the algorithm's overall progress by shortening its earlier iterations.

We next describe each step.

E-step.

The purpose of the E-step is to sample from the posterior distribution of Y conditional on O and θ_t. Various sampling algorithms can be used to achieve this objective.

Several of these algorithms involve Monte Carlo techniques that, given a current Y_k, propose candidate values Y' that may be accepted or rejected based on the ratio of $L(Y' \mid O, \theta_t)$ over $L(Y_k \mid O, \theta_t)$ (or equivalently and more computationally conveniently, the difference between the two log likelihoods). For such algorithms, it is often immensely useful that

$$\frac{L(Y' \mid O, \theta_t)}{L(Y_k \mid O, \theta_t)} = \frac{L(O, Y' \mid \theta_t)/L(O \mid \theta_t)}{L(O, Y_k \mid \theta_t)/L(O \mid \theta_t)} = \frac{L(O, Y' \mid \theta_t)}{L(O, Y_k \mid \theta_t)}. \tag{2.41}$$

In other words, the algorithm need only be able to compute the joint (log) likelihood whose computation is explicitly laid out in Section 2.11.1. Then, by invoking this computation repeatedly at various values of Y', such Monte Carlo-based algorithms produce sequences of the form $\{Y_k\}$ such that, by applying burn in and sampling at appropriate intervals, they yield subsequences Y_b, $b = 1, \ldots, B$ that are effectively drawn from the desired posterior of Y conditional on O and θ_t.

M-step.

The EM algorithm iteration is then completed by obtaining θ_{t+1}:

$$\theta_{t+1} =_{\theta} \sum_{b=1}^{B} \log L(O, Y_b \mid \theta), \tag{2.42}$$

in which the values $Y_b, b = 1, \ldots, B$ were obtained and fixed by the prior E-step so that the parameter values of θ are the only remaining variables of the objective function.

As (2.19) lays out in detail, the joint likelihood function $L(O, Y_b \mid \theta)$ decomposes into the product of subcomponents of the form $L(O \mid Y_b) \cdot L(Y_b \mid \theta)$. Then, because $\log L(O, Y_b \mid \theta) = \log L(O \mid Y_b) + \log L(Y_b \mid \theta)$, we can greatly simplify the objective (2.42) that is used in the M-step. Indeed, the $\log L(O \mid Y_b)$ terms drop out completely, leaving only $\sum_{b=1}^{B} \log L(Y_b \mid \theta)$ as the M-step's objective function.

2.11.3 Implementation Details

E-step.

In the E-step of iteration t of the EM algorithm, we use Metropolis-Hastings algorithm to sample a sequence of latent variable Y_b, $b = 1, ..., B$ from $L(Y \mid O, \theta_t)$, the posterior distribution of latent variable Y.

The parameter θ_t is taken as given in the E-step. An initial latent variable Y_0 is drawn from θ_t. Then with current latent variable Y_k, we propose a candidate value of the new latent variable Y' by a symmetric jump function $J(Y' \mid Y_k)$, which follows a normal distribution $N(Y_k, \Phi)$. Taking log of (2.41), the new latent variable Y' will be accepted if

$$\log L(O \mid Y', \theta_t) + \log L(Y' \mid \theta_t) - \log L(O \mid Y_k, \theta_t) - \log L(Y_k \mid \theta_t) > 0. \tag{2.43}$$

For either the current latent variable Y_k or the proposed new Y', the log probability of the observed outcomes conditional on Y and θ_t has two parts:

$$\log L(O \mid Y, \theta_t) = \sum_i \log L(x_i \mid \beta_i, \alpha_i, \mathbb{E}Q_{ij}^{post}, j \in \mathscr{J}) + \sum_j \log L(c_j \mid Q_j, \theta_{\mathscr{P}}).$$

The first term gives the likelihood of observing consumer choice outcome x_i conditional on the latent variables (calculation given in (2.20)), and the second term gives the likelihood of observing product classification outcome c_j conditional on latent variables and platform classification parameters (calculation given in (2.36)).

The calculations of $\log L(Y' \mid \theta_t)$ and $\log L(Y_k \mid \theta_t)$ can be arranged into three blocks: Q_j governed by $\theta_{\mathscr{J}}$, α_i and β_i governed by θ

N, and $\mathbb{E}Q_{ij}^{post}$ governed by θ_Δ, each block is independent. $\log L(Q_j \mid \theta_{\mathscr{J}})$ and $\log L(\alpha_i, \beta_i \mid$
θ

N) are straightforwardly the log-likelihood of observing Q_j, α_i and β_i from their corresponding multivariate normal distributions. The calculation of quality belief level $\mathbb{E}Q_{ij}^{post}$ is given in (2.31) and (2.33), which requires numerical evaluation.

The numerical evaluation of equations (2.31) and (2.33) are similar, and we will use (2.31) as an example below.

First, we sample K points $\{x_1, \ldots, x_K\}$ from the distribution of $Q'_{ij} \mid A_j, q_{ij}, \Sigma_{ij}^{\Delta}$, which is distributed multivariate normal $N(\mu_{1,ij}, \Sigma_{1,ij})$ given in (2.23). Note that when Q_j is given, $N(\mu_{1,ij}, \Sigma_{1,ij})$ only depends on Δ_{ij}. For each of these points x_k, we calculate its probability density function $f_{Q'|A,q,e}(x_k)$, and $1 - \Phi\left(r - \Lambda_Q x_k - \Lambda_A A_j\right)$, which is the probability that this product is correctly classified to $c_j = 1$ given x_k. Then, (2.31) can be obtained as a weighted sum of these K points:

$$\mathbb{E}Q_{ij}^{post} = \frac{\sum_{k=1}^{K} x_k f_{Q'|A,q,e}(x_k)\left[1 - \Phi\left(r - \Lambda_Q x_k - \Lambda_A A_j\right)\right]}{\sum_{k=1}^{K} f_{Q'|A,q,e}(x_k)\left[1 - \Phi\left(r - \Lambda_Q x_k - \Lambda_A A_j\right)\right]}. \tag{2.44}$$

Similarly, (2.33) can be calculated by sampling K points $\{x_1,...,x_K\}$ from the same distribution $N(\mu_{1,ij},\Sigma_{1,ij})$, but weighted by the probability of having them classified to $c_j = 0$:

$$\mathbb{E}Q_{ij}^{post} = \frac{\sum_{k=1}^{K} x_k f_{Q'|A,q,e}(x_k)\left[\Phi\left(r-\Lambda_Q x_k - \Lambda_A A_j\right)\right]}{\sum_{k=1}^{K} f_{Q'|A,q,e}(x_k)\left[\Phi\left(r-\Lambda_Q x_k - \Lambda_A A_j\right)\right]}. \tag{2.45}$$

Noticed that $\mathbb{E}Q_{ij}^{post}$ is a function of Δ_{ij} when Q_j, A_j, c_j, e_i are fixed, we can compute the interpolation functions between $\mathbb{E}Q_{ij}^{post}$ and Δ_{ij}. In particular, for each set of parameters θ_t, we compute two interpolation functions for each product j, one for the experienced consumer case and the other for the inexperienced consumer case in the beginning of the E-step. These functions will be used throughout the E-step every time when a new sample Y' is proposed and we need to decide whether to accept it according to (2.43). This way we avoid the time-consuming calculation of $\mathbb{E}Q_{ij}^{post}$ for every proposed sample.

M-step.

The purpose of M-step is to update the values of the parameters based on the samples of latent variables obtained from E-step. At an iteration t, after the E-step, we obtain a sequence of latent variables Y_b, $b = 1,...,B$. The parameters for the next step is set to be the ones that maximize the M-step objective function:

$$\begin{aligned}
\theta_{t+1} &=_{\theta} \sum_{b=1}^{B} \log L(O, Y_b \mid \theta_t) \\
&=_{\theta} \sum_{b=1}^{B} [\log L(O \mid Y_b; \theta_t) + \log L(Y_b \mid \theta_t)] \\
&=_{\theta} \sum_{b=1}^{B} [\log L(x_i : i \in \mathscr{N} \mid Y_b) + \log L(c_j : j \in \mathscr{J} \mid Y_b; \theta_{\mathscr{P},t}) + \log L(Y_b \mid \theta_t)] \\
&=_{\theta} \sum_{b=1}^{B} [\log L(c_j : j \in \mathscr{J} \mid Y_b; \theta_{\mathscr{P},t}) + \log L(Y_b \mid \theta_t)],
\end{aligned} \tag{2.46}$$

because $L(x_i : i \in \mathscr{N} \mid Y_b)$ is not a function of the parameters θ_t. Except for determining the platform classification parameters $\theta_{\mathscr{P}}$ for the $t+1$ iteration, we only need $\sum_{b=1}^{B} \log L(Y_b \mid \theta)$ as the M-step's objective function.

To determine the parameters other than $\theta_{\mathscr{P}}$, let $Y_b^N = \{Q_j, \Delta_{ij}, \beta_i, \alpha_i\}$. Note that in M-step we update Δ_{ij} instead of $\mathbb{E}Q_{ij}^{post}$ because when Q_j, A_j, c_j, e_i are fixed, the randomness in $\mathbb{E}Q_{ij}^{post}$ comes solely from Δ_{ij}, and the value of $\mathbb{E}Q_{ij}^{post}$ can be calculated from (2.44) or (2.45). The elements in Y_b^N are all from normal distributions, where Q_j is from $N(\mu_0^a, \Sigma_0^a)$, Δ_{ij} is from $N(\mu^\Delta, \Sigma^\Delta)$, and β_i and α_i are from $N(\mu^{\mathscr{N}}, \Sigma^{\mathscr{N}})$.

Let the means be denoted as $\mu^\theta = \{\mu_0^{a_j}, \mu^\Delta, \mu^{\mathscr{N}}\}$, and the variances denoted as $\Sigma^\theta = \{\Sigma_0, \Sigma^\Delta, \Sigma^{\mathscr{N}}\}$. The parameters that maximize the M-step objective function are

$$\mu_{t+1}^\theta = \frac{1}{B} \sum_{b=1}^{B} Y_b^N,$$

and

$$\Sigma_{t+1}^\theta = \frac{1}{B} \sum_{b=1}^{B} (Y_b^N - \mu_t^\theta)(Y_b^N - \mu_t^\theta)^\top.$$

To determine $\theta_{\mathscr{P}}$, the objective function is:

$$\theta_{\mathscr{P},t+1} =_\theta \sum_{b=1}^{B} [\log L(c_j : j \in \mathscr{J} \mid Q_j; \theta_{\mathscr{P},t}) + \log L(\mathbb{E}Q_{ij}^{post} \mid \theta_{\mathscr{P},t})]. \tag{2.47}$$

It is evaluated numerically with the optimizer “minimize” from the Python SciPy package, using BFGS method.

2.12 Appendix IV: Estimation Results

We present the results on parameters estimated according to the EM algorithm presented in Section 2.11. In Section 2.12.1, we discuss the estimated joint distribution of products' quality levels Q_j and attributes A_j on the two sales channels 3P and 1P. In Section 2.12.2, we report the estimated parameters related to signal noise Δ_{ij}.

2.12.1 Results on Quality Distribution and Platform Classification

We estimated the joint distribution of $(Q_j, A_j) \sim N(\mu, \Psi)$, where $\mu = [\mu_Q, \mu_A]^\top$, and $\Psi = [[\Sigma_Q, \Sigma_{QA}], [\Sigma_{AQ}, \Sigma_A]]$. In estimating such joint distribution of quality and attributes, note that the two dimensions of quality are both constructed, so their means μ_Q are only meaningful when comparing with each other. In addition, the attribute means μ_A are observed from data. We are the most interested in Σ_{QA}, namely how the quality level of each product is correlated with its attributes, because this is what enables consumers to learn quality by observing attributes. We report the estimated values of Σ_{QA} weighted by Σ_Q in Table 2.15.

In our model, we let the attributes A_j consist of a product's rating, return rates, experienced-related complaint rate, product-related complaint rate,[4] how long the product has been sold on the marketplace (product tenure), how long the seller has been doing business on the marketplace (seller tenure), and the product's origin (Liao). We estimated how A_j is correlated with Q_1 (product type quality) and Q_2 (SC-D quality). As shown in Table 2.15, both dimensions of quality are positively correlated with the attributes rating, seller tenure, product tenure, and Liao. Comparing the two dimensions of quality with each other, SC-D quality is more correlated with seller tenure and product tenure than product type quality, which makes sense because we would expect that sellers that are more experienced in the

[4] We are provided with data on whether an order is associated with a complaint filed after receiving the product. Though the content of each complaint is not available, we are provided with the complaint types. These complaints fall into two categories: those related to a product's quality and those related to user experience. The former includes intrinsic quality issues, such as product preservation and chemical contamination, packaging problem, such as improper packaging material, and misleading description such as size or weight of the product not matching seller's claim. The latter is more subjective. Examples include unsatisfactory taste or texture and overprice. If a complaint is associated with both groups, it is counted once in each group.

business have products with higher SC-D quality.

Both dimensions of quality are negatively correlated with the two consumer complaint rates, where product type quality is more correlated with product-related complaints and SC-D quality is more correlated with experience-related complaints as expected. Whereas product type quality is negatively correlated with return rate, counter-intuitively, we find that SC-D quality is positively correlated with return rate. This result may suggest that sellers with higher SC-D quality products are more willing to facilitate return requests from consumers, which leads to a higher return rate.

The platform employs a classification rule $\Lambda X_j + \nu_j \geq r$ to select products sell through its 1P channel and the remaining products are sold on 3P, where Λ consists of Λ_Q and Λ_A, the coefficient of Q_j and A_j respectively. In table 2.16, we report the estimated results on Λ_Q and Λ_A weighed by r.

With the estimated joint distribution of quality and attributes and the platform's classification rule, we can recover the quality distributions on 1P and 3P. In Figure 2.9, we plot the marginal distribution of quality on the two sales channels. When comparing the quality distribution between 1P and 3P, we can see that the majority of the quality distribution on 3P and 1P overlaps, but the 1P distribution is higher in both quality dimensions.

2.12.2 Estimation Results of the Noise Parameters

The signal noise Δ_{ij} is distributed $N(0, \Sigma_{ij}^{\Delta})$, where the noise parameter $\Sigma_{ij}^{\Delta} = [[\sigma_{1,ij}^{\Delta}, \rho_{ij}^{\Delta}], [\rho_{ij}^{\Delta}, \sigma_{2,ij}^{\Delta}]]$. We report the estimated noise parameters Σ_{ij}^{Δ} in this section. Noise parameters represent the level of information asymmetry a consumer experiences when she tries to determine a product's quality. Smaller values in noise parameters indicate a lower level of information asymmetry. In Appendix Table 2.17, we compare the level of information asymmetry between blockchain traced products and non-traced products.

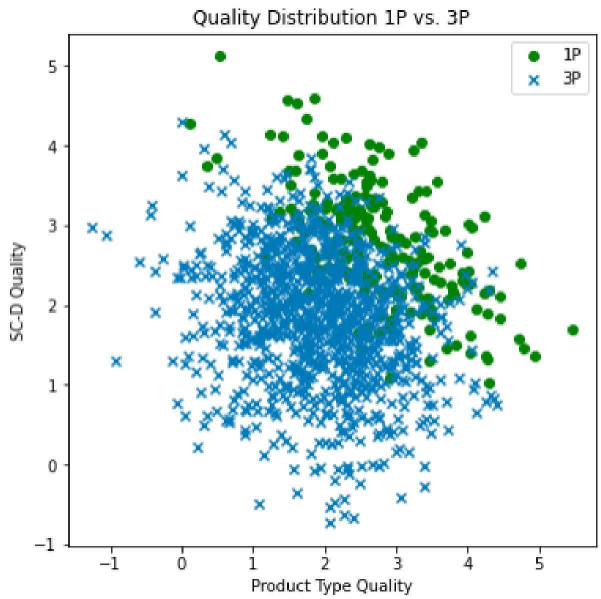

Figure 2.9: Quality Distribution Generated from Estimated Parameters.

We estimated 24 Σ_{ij}^{Δ}'s for all the combinations of (e_i, c_j, ty_j, tr_j) (please refer to Section 2.11.1 for details about the modeling of the noise parameters) using 5% of the orders placed in the post-treatment period. The sample contains a monthly average of 571 consumers placing orders on 616 products. In Table 2.17, we report $\sigma_{1,ij}^{\Delta}$ and $\sigma_{2,ij}^{\Delta}$, the noisiness of product type quality and SC-D quality respectively. The values of he noise parameters are weighed by the mean value of quality Q_j.

Similarly, 5% of the orders placed in the pre-treatment period were used to estimate the signal noise of the ready-to-eat products. The sample contains orders placed by 513 consumers on 120 products. We estimated 8 different Σ_{ij}^{Δ} for all the combinations of consumer experience level, sales channel, and whether the cold-chain logistics is specified by the seller. In Table 2.18, we report $\sigma_{1,ij}^{\Delta}$ and $\sigma_{2,ij}^{\Delta}$, the noisiness of product type quality and SC-D quality respectively. The values of he noise parameters are weighed by the mean value of quality Q_j.

2.13 Appendix V: Consumer Surplus Calculations

In our model, a consumer's belief of quality $\mathbb{E}Q_{ij}^{post}$ differs from the true product quality Q_j. While $\mathbb{E}Q_{ij}^{post}$ determines her purchasing decision, the true quality Q_j determines the consumer's *ex post* utility. To account for this difference between quality belief and quality, we use an adjusted formula to calculate consumer welfare [11, 12, 84].

Let V_{ij} denote a consumer i's mean utility, excluding the random shock. When a consumer makes her decision based on her posterior belief of quality, her mean utility is

$$V_{ij}(\mathbb{E}Q_{ij}^{post}) = \beta_i \mathbb{E}Q_{ij}^{post} - \alpha_i p_j.$$

Had she observed the true quality, her mean utility is

$$V_{ij}(Q_j) = \beta_i Q_j - \alpha_i p_j.$$

Then, a consumer i's expected surplus weighed by price elasticity α_i is

$$\mathbb{E}(CS)_i = \frac{1}{\alpha_i}\left[\log\left(\sum_{j=1}^{J}\exp(V_{ij}(\mathbb{E}Q_{ij}^{post}))\right) + \sum_{j=1}^{J}\pi_{ij}\left(V_{ij}(Q_j) - V_{ij}(\mathbb{E}Q_{ij}^{post})\right)\right], \quad (2.48)$$

where

$$\pi_{ij} = \frac{\exp\left(V_{ij}(\mathbb{E}Q_{ij}^{post})\right)}{1+\sum_{j=1}^{J}\exp\left(V_{ij}(\mathbb{E}Q_{ij}^{post})\right)}.$$

The first term of (2.48) is the standard log sum formula to calculate consumer surplus in a discrete choice model with a type-I extreme value random term. The second term captures the fact that a consumer's purchasing decision is made under her belief $\mathbb{E}Q_{ij}^{post}$ rather than the true quality Q_j.

Each consumer's surplus is calculated according to (2.48). Table 2.19 below reports the average values of the two consumer experience groups and the average values across all consumers.

	Type	Group (p,t,r)	Treated	Count	Price	Tenure	Rating
1P	Dried-without-salt	(1,1,1)	Treated	2	18.87	54.00	1.00
					(1.18)	(1.41)	(0.00)
			Pool	12	18.36	52.35	1.00
					(3.19)	(8.18)	(0.00)
		(0,1,1)	Treated	1	12.25	42.00	1.00
					-	-	-
			Pool	26	12.14	44.32	1.00
					(4.22)	(12.47)	(0.00)
1P	Ready-to-eat	(1,1,1)	Treated	4	16.26	40.75	1.00
					(1.64)	(10.11)	(0.00)
			Pool	15	16.11	38.65	1.00
					(2.87)	(13.62)	(0.00)
		(0,1,1)	Treated	1	7.52	42.00	1.00
					-	-	-
			Pool	22	7.03	38.05	1.00
					(3.81)	(10.02)	(0.00)
1P	Non-premium	(1,1,1)	Treated	1	5.45	30.00	1.00
					-	-	-
			Pool	23	5.78	25.75	1.00
					(2.11)	(12.31)	(0.00)
		(0,1,1)	Treated	4	3.17	45.25	1.00
					(1.68)	(11.06)	(0.00)
			Pool	33	3.25	32.71	1.00
					(2.41)	(13.18)	(0.00)

Table 2.13: This table summarizes the key characteristics (prices, tenure, and rating) of the treated unit and the donor pool within each group for 1P products. For each of the characteristics, we present the group mean with standard deviation in parenthesis. Price is measured on a per-item basis, calculated by dividing the unit product price by the quantity contained in each product. Tenure is measured by the number of months since the product was first listed on the platform.

3P	Dried-without-salt	(1,1,1)	Treated	21	16.63	46.71	1.00
					(2.32)	(8.74)	(0.00)
			Pool	133	15.55	40.10	1.00
					(4.23)	(15.29)	(0.00)
		(1,1,0)	Treated	1	16.76	41.00	0.67
					-	-	-
			Pool	18	16.12	40.20	0.76
					(3.22)	(7.22)	(0.29)
		(0,1,1)	Treated	17	9.91	45.07	1.00
					(2.22)	(8.74)	(0.00)
			Pool	313	8.31	35.50	1.00
					(3.07)	(18.77)	(0.00)
3P	Ready-to-eat	(1,1,1)	Treated	8	12.02	46.33	1.00
					(1.78)	(11.57)	(0.00)
			Pool	34	12.15	46.87	1.00
					(2.32)	(14.31)	(0.00)
		(0,1,1)	Treated	4	7.21	50.50	1.00
					(0.58)	(7.78)	(0.00)
			Pool	102	7.35	46.50	1.00
					(2.86)	(20.61)	(0.00)
		(1,0,1)	Treated	12	10.54	46.33	1.00
					(1.18)	(11.57)	(0.00)
			Pool	34	10.05	36.75	1.00
					(2.55)	(13.01)	(0.00)
3P	Non-premium	(1,1,1)	Treated	3	5.88	44.50	1.00
					(0.54)	(14.85)	(0.00)
			Pool	59	4.36	40.50	1.00
					(1.37)	(17.55)	(0.00)
		(0,1,1)	Treated	1	2.07	56.00	1.00
					-	-	-
			Pool	204	2.43	42.66	1.00
					(1.07)	(22.85)	(0.00)
		(1,0,1)	Treated	1	5.01	25.00	1.00
					-	-	-
			Pool	38	4.95	22.50	1.00
					(0.89)	(10.24)	(0.00)
		(0,0,1)	Treated	2	3.36	25.00	1.00
					(0.78)	(0.00)	(0.00)
			Pool	104	3.23	20.33	1.00
					(1.57)	(15.19)	(0.00)

Table 2.14: This table summarizes the key characteristics (prices, tenure, and rating) of the treated unit and the donor pool within each group for 3P products.

A	rating	seller tenure	product tenure	complaints (exp)
Q_1 (product type)	0.0277	0.0233	0.0251	-0.0110
Q_2 (SC-D)	0.0333	0.0395	0.0345	-0.0571
A	complaints (prod)	return rate	Liao	
Q_1 (product type)	-0.0274	-0.0101	0.0210	
Q_2 (SC-D)	-0.0043	0.0022	0.0166	

Table 2.15: Σ_{QA}: Covariance between quality Q and attributes A

λ_{Q_1} (product type)	λ_{Q_2} (SC-D)	λ_{A_1} (rating)
0.1027	0.1465	0.0960
λ_{A_2} (seller tenure)	λ_{A_3} (product tenure)	λ_{A_4} (complaints (exp))
0.1065	0.1910	-0.1135
λ_{A_5} (complaints (prod))	λ_{A_6} (return rate)	λ_{A_7} (Liao)
-0.1169	-0.1104	0.1240

Table 2.16: Λ: The platform's classification parameters (weighted by r)

				Experienced		Inexperienced	
			Count	σ_1^Δ	σ_2^Δ	σ_1^Δ	σ_2^Δ
1P	Dried-without-salt	traced	10	0.1374	0.2104	0.1784	0.2572
		non-traced	18	0.1375	0.2232	0.1800	0.2674
	Ready-to-eat	traced	5	0.1444	0.2070	0.1788	0.2608
		non-traced	10	0.1631	0.2394	0.2083	0.2804
	Non-premium	traced	5	0.1424	0.2348	0.1828	0.2797
		non-traced	14	0.1586	0.2473	0.1785	0.2803
3P	Dried-without-salt	traced	28	0.1676	0.2293	0.1715	0.2621
		non-traced	289	0.1808	0.2603	0.2228	0.2984
	Ready-to-eat	traced	25	0.1651	0.2305	0.1811	0.2674
		non-traced	80	0.2032	0.2806	0.2413	0.3057
	Non-premium	traced	7	0.1761	0.2437	0.2027	0.3077
		non-traced	125	0.1968	0.2846	0.2322	0.3276

Table 2.17: Noisiness in Quality Signal: Traced vs. Non-traced Products.

			Experienced		Inexperienced	
		Count	σ_1^Δ	σ_2^Δ	σ_1^Δ	σ_2^Δ
1P	specified	7	0.1671	0.2324	0.2093	0.2909
	non-specified	9	0.1862	0.2552	0.2321	0.3091
3P	specified	31	0.2017	0.2770	0.2788	0.3108
	non-specified	73	0.2087	0.2984	0.2883	0.3104

Table 2.18: Noisiness in Quality Signal: Cold-Chain Logistics Specified vs. Non-specified for Ready-to-eat Products.

	Consumer Experience Level		
	Experienced	*Inexperienced*	Average
(1) Baseline: no tracing	13.36	9.70	11.90
(2) Current adoption of tracing	13.43	9.85	12.02
(3) All 3P products get traced	14.06	10.39	12.62
(4) All products get traced	14.35	10.88	13.03

Table 2.19: Average Consumer Surplus.

Chapter 3: Impact of Price Strategy on Returns: a Price Transparency and Valuation Uncertainty Story

We study how discount strategies affect *net* sales – sales after accounting for returns. We evaluate two prevalent discount strategies: bundle discounts, which apply discounted prices exclusively to product bundles, and single-item discounts, where the discount extends to all products. Two primary factors are considered in analyzing the discount strategies' performance in enhancing net sales: customers' post-purchase valuation updates, which influence their return behavior, and their attentiveness to the intricacies of the discount pricing structure. The latter is particularly relevant under bundle discounts, where customers retaining any items must pay full price, a detail that could be overlooked by those inattentive to pricing, affecting both purchase and return decisions. We adopt a forward-looking approach to customer behavior by modeling the interconnected purchase and return decisions, incorporating post-purchase valuation changes and potential pricing inattentiveness. Partnering with one of Turkey's largest fashion retailers, we apply structural estimation to ascertain how customers reassess product valuations post-purchase and their attentiveness to pricing details. Furthermore, we simulate generic retail scenarios to discern the optimal discount strategy for retailers. Our findings reveal that factoring in returns drastically shifts the retailer's preference in discount strategies. By tailoring bundles that account for customers' post-purchase valuation, the Retailer's net sales can be enhanced by 15.6%. We also find that customers tend to overlook complex pricing structures, leading to impulsive buying and increased returns. Enhancing customer attentiveness of pricing can decrease the Retailer's return rates by 20.9%. Moreover, improving customer attentiveness to pricing benefits retailers by enabling them to create more versatile bundle offers, further optimizing their sales strategy.

discount strategy, bundling, product return, price transparency, bounded rationality, structural estimation

3.1 Introduction

The fashion industry is under growing scrutiny owing to its significant environmental impact, a considerable proportion of which stems from waste. Astonishingly, 92 million tons of the 100 billion garments produced globally each year are discarded, contributing to escalating landfills. Returns constitute a notable portion of this waste, accounting for 17 million tons of discarded garments in the United States alone [105]. In addition to their environmental repercussions, product returns present a multifaceted logistical and financial challenge for retailers. A substantial 5% to 15% of in-store purchases and a staggering 15% to 40% of online purchases are returned, accumulating to a whopping $816 billion in returned merchandise in 2022 alone [69].

The problem of product returns is further exacerbated during sales seasons, where retailers deploy significant discounts and additional promotions to liquidate inventory. While these price reductions are effective in moving stock, they also tend to result in a significant increase in returns. Consequently, the *net* sales – the volume of sales after accounting for returns – may not experience a significant increase. Compounding this issue is the seasonal character of these products, which often renders them non-saleable upon return, thereby amplifying the logistical and financial burdens faced by retailers. There is an option to mitigate these issues by tightening return policies or increasing the cost of returns for customers. However, retailers may be hesitant to enforce more stringent return policies, concerned that such measures could alienate customers and subsequently reduce demand [26].

Intriguingly, the very strategies retailers employ to discount products also influence returns. Take the bundle discount, for instance, a prevalent strategy in retail where customers are incentivized with a price discount only upon purchasing multiple designated items together. This price structure implies that customers who opt to return a part of a bundle must pay the full original price of the items they retain. However, such intricacies might not be immediately apparent during the purchasing process, as customers are presented with a collective bundle price, leading to a potential hidden "cost" of return that may not be initially recognized. This complexity in pricing structure is not unique to bundle discounts but also exists in other prevalent discount strategies, such as Buy One Get One (BOGO) offers. In contrast, the price structure of a straightforward single-item discount, where each item is sold at a reduced price, is more transparent by its nature.

Bundling, the strategy of selling multiple items for a collective price, is a common practice in the retail industry. It has been shown to effectively enhance firm profits by categorizing customers according to their diverse willingness to pay, thus facilitating the extraction of consumer surplus. Despite extensive research on the efficacy of different bundling strategies in optimizing firm sales, there is a notable gap in understanding how the price structure of bundles affect customers' grasp of pricing intricacies, especially when customers might not be fully attentive to the pricing details. Specifically, customers who pay full attention to price details would perceive the bundled products as less valuable when considering the additional "cost" associated with future partial returns. Conversely, initial demand may not be affected among inattentive customers. The realization of the hidden cost associated with partial returns only surfacing at the return stage for inattentive customers, which can either discourage them from making returns or, in some cases, lead them to return the entire bundle. Therefore, the price transparency issue inherent to the pricing structure of bundle discounts can profoundly affect customers' purchase and return decisions, consequently impacting net sales.

In collaboration with one of Turkey's largest fashion retailers (the "Retailer"), our research delves into the impact of return considerations on a retailer's choice of discount strategies. We juxtapose bundle discounts against single-item discounts, evaluating their effectiveness in terms of both sales and net sales quantities. Recognizing the intricate pricing structure inherent in bundle discounts, our study examines how customer attentiveness to pricing nuances impacts the retailer's superior discount strategy. Our aim is to identify and recommend discount strategies that can enhance net sales for retailers, while accommodating the varying degrees of customer attentiveness to pricing.

Our study makes several contributions. Firstly, we demonstrate that the inclusion of return considerations significantly shapes a retailer's choice of the most effective discount strategy. After a purchase, customers often reassess their initial valuation of the product, leading to post-purchase updates. For retailers focusing on net sales, ensuring that the purchased products deliver a positive surplus to customers after the purchase, rather than solely relying on ex-ante valuation, is crucial. Our empirical findings reveal a discernible pattern: products initially highly valued by customers at the point of purchase are likely to receive positive valuation updates, whereas products deemed less valuable initially are prone to further devaluation. This insight is particularly valuable for retailers in predicting the ex-post value of different product categories, aiding in more informed decision in bundle design.

Secondly, we identify that customer inattentiveness to pricing intricacies influences their purchasing and returning behavior, thereby affecting a retailer's net sales. Our empirical findings suggest that customers generally do not pay close attention to detailed pricing when faced with complex structures like bundle pricing. Despite being aware of the option to return, most customers struggle to accurately factor in the pricing details associated with returns for bundle purchases. Consequently, they often overvalue the bundle, leading to impulsive buying and subsequent returns. Challenging the traditional view that retailers favor non-strategic customers, our findings indicate that rational customers, those who pay close attention to price structures and nuances, can actually benefit retailers.

Thirdly, our analysis underscores the necessity for retailers to consider both post-purchase valuation updates and customer inattentiveness to complex pricing structures when devising discount strategies. If a retailer prioritizes sales volume only, targeting a base of inattentive customers with bundle discounts can prove more profitable. However, to enhance net sales, attentive customers are found to be more beneficial due to their higher final purchase quantities after returns. Furthermore, engaging with a more attentive customer segment allows retailers the flexibility to incorporate a wider array of product categories into their bundle offers. This strategy not only caters to a diverse range of customer preferences but also enables the clearance of items that are less favored or those experiencing negative post-purchase surpluses. Effectively, this approach helps retailers strategically manage inventory, moving products that are otherwise challenging to sell.

Lastly, our study delves into the bounded rationality of customers through structural estimation. Acknowledging customers' forward-looking behavior – that is, the anticipation of potential returns at the time of purchase – our model incorporates both the purchase and return decisions of customers, including the valuation updates that occur in between. We also account for the likelihood that customers might miscalculate the cost implications of partial returns when faced with complex pricing structures. Given the complexity of the model, we employ the Monte-Carlo Expectation Maximization (MCEM) algorithm for its computational efficiency in handling high-dimensional data. Through this method, we gain valuable insights into the subtle intricacies of consumer behavior, offering a deep understanding of how retailers can tailor their discount strategies to optimize their net sales outcomes.

The structure of this paper is laid out as follows: We begin with a review of the related literature in Section 2. In Section 3, we delve into the data, offering initial empirical evidence into the impact of bundles on product net sales. In Section 4 we present our model of customers' purchase and return decisions. We describe our structural estimation method in Section 5. The results obtained from our model estimation are thoroughly detailed in Section 6. This is followed by Section 7, which presents a series of counterfactual analyses to explore the potential enhancements in bundle discount strategies and the pivotal role of customer attentiveness. Section 8 presents managerial insights illustrated by simulations using generic product categories, demonstrating how accounting for returns and customer attentiveness significantly affect a retailer's optimal discount strategy. Finally, we encapsulate our findings and present our conclusions in Section 9.

3.2 Literature Review

The extensive body of literature in operations and marketing explores various firm strategies for managing product returns through the design of return policies. To mitigate product returns, many studies suggest charging customers a restocking fee or a shipping fee for returned products as an economically optimal solution, as evidenced in studies like [42, 110], and [6] and [98]. However, [26] offer a contrasting viewpoint, suggesting that allowing customers to return products for free can lead to an increase in their future purchases. On the other hand, imposing return charges might lead to a decrease in subsequent purchases from that retailer. Such finding elucidates why, despite the benefits of imposing return costs, many retailers opt for more lenient return policies. Our research contributes a fresh angle to this literature by examining the "hidden costs" inherent to discount strategies with complex pricing structures. Our study finds that pricing strategies can subtly influence return behaviors, offering actionable insights for retailers to strike a balance between managing returns and nurturing customer relationships.

Another stream of literature on product returns focuses on understanding customers' return behavior. [9] point out that the cost of return is the key to understand customers' return behavior. They model individual customer's purchase and return behavior and measure the cost of return. In the retail industry setting, studies have shown that return behavior can be influenced by various factors that affecting customers' decision-making, such as product reviews [106], online image presentation [49], the form of discounts [54], web technology [43], and return period policy [53]. Our research introduces a distinctive factor influencing customer returns, which has not been extensively explored in existing literature – customers' inattentiveness to price details. Employing a model similar to [9], we quantify the impact of this attentiveness to price on product sales and returns, thereby contributing a novel perspective to the understanding of customer return behavior.

Our study, which focuses on the impact of bundle discounts on product sales and returns, is built upon the extensive literature on product bundling. Since [4], a primary advantage of bundling has been identified as its ability to enhancing a firm's profit by reducing consumers' valuation dispersion and extracting their surplus. For a review of the early literature in this domain, please refer to [109]. Subsequent theoretical research has delved into the performance of bundling, particularly examining its two forms: pure bundling, where only the bundle is sold and not the individual products, and mixed bundling, where both the bundle and the separate products within it are sold individually [e.g., 13, 21, 61, 93, 115].

More recent literature has explored various operational aspects of product bundling, such as decision-making under demand uncertainty [94] and supply constraints [31]. Beyond product sales, [117] investigate how bundling can be utilized to manage customer search behaviors. Our research contributes to this field by expanding the perspective of retailers to include the product return stage. Specifically, we explore how the bundling of products can influence customers' understanding of pricing and, consequently, their purchase and return decisions. This addition offers a new dimension to the existing body of knowledge on product bundling, particularly in the context of retail returns management.

Moreover, our research adds to the existing literature on the impact of information transparency on customers, with a specific focus on price transparency and its influence on consumer behavior. Previous studies emphasize how the disclosure and detailed explanation of the components that make up a price can affect customer preferences, such as breaking down a price into gross retail proceeds, royalties, and taxes [32], or dividing a product's price into its base cost plus shipping and handling fees [19]. In a related vein, [96] explores cost transparency, demonstrating that providing customers with sensitive information about product costs can build trust and consequently boost their willingness to buy. We delve into the effects of price structure on consumer behaviors, considering both monetary and operational perspectives, thereby intersecting with the operational transparency literature. This body of work illustrates how making a firm's operational processes known to consumers can amplify perceptions of product quality [29], elevate the likelihood of making a purchase [28], and enhance the success of crowdfunding campaigns [95]. Our findings indicate that rendering the price structure more transparent to customers can significantly boost sales and mitigate returns for retailers, thereby offering a new dimension to the understanding of transparency in consumer-retailer dynamics.

Finally, our research contributes to the empirical literature examining consumers' limited rationality in marketplaces. Previous empirical studies demonstrate that general withholding of pricing information can make it challenging for consumers to navigate and understand complex information environments. In experimental studies, [50] demonstrate the difficulties individuals face in processing and comprehending correlated information within intricate informational contexts. [74] show that when faced with intentionally complicated information, recipients often make systematic errors in their understanding and decision-making.

Looking at the marketplace from a broader perspective, it is evident that consumer populations consist of both strategic and non-strategic individuals. In the context of durable goods, [60] identify a segment of strategic consumers who hoard goods in anticipation of future price drops, a behavior not observed in the rest of the population. This phenomenon is similarly observed by [86] in the airline ticket sales market. [34] find a comparable trend among students purchasing and reselling textbooks, although not all take into account the potential resale value of their textbooks. Building upon these insights, our study employs a structural model to delve into consumer purchasing and return behaviors, with a particular focus on their capacity to grasp complex price structures. Our analysis reveals that customers tend to overlook price details, a finding that holds significant implications for retailers in determining their optimal discount strategies.

3.3 Background and Data

3.3.1 Industry Overview

Our study is set in the context of the fashion retail industry. We collaborate with the Retailer, a leading Turkish fashion retailer that Specializes in denim pants but also offers a variety of other apparel including shirts, T-shirts, polos, and fashion pants. The Retailer operates a network of 433 stores throughout Turkey, in addition to an online shopping platform, with inventory, pricing, and promotional decisions centrally managed.

The Retailer frequently offer single-item discounts throughout the year, with an annual end-of-season sale typically occurring from mid-June to July (the "Markdown Season"). During the Markdown Season, additional bundle discounts are introduced. Under single-item discounts, any purchase of a discounted product is charged at the discounted price and fully refundable upon return. In contrast, bundle discounts require the purchase of any two products from designated categories as a combined basket to receive the discount. If a single item from such a two-item basket is returned, the basket no longer qualifies for the discount, and the retained item is charged at full price.

The Retailer upholds a generous return policy. Customers dissatisfied with their purchases can return the items for a full refund within 30 days of purchase, with no shipping or restocking fees incurred. Product returns are prompted by various factors, including defective items, unknown fit attributes, and valuation uncertainty. While the Retailer offers fitting rooms for customers to try garments on, consumers often only fully discern the suitability of an item, especially regarding its fit or other experiential attributes, after they have purchased it and experienced it in their own environment [107]. This delayed realization of product value is a significant factor in the decision to return products.

3.3.2 Overview of the Data

The Retailer provided us with detailed customer-order level data a rich data set with $2,113,930$ data points, encompassing in-store transactions from March 1st to July 31st, 2019, a total of 21 weeks. The Markdown Season spanned the last 6 weeks of this period, from June 15th to July 31st, 2019. During this Markdown Season, 32% of the products qualified for bundle discounts. These bundle discounts were limited to certain categories, as outlined in Table 3.1.

As indicated in Table 3.1, the nature of these bundle discounts varies. Some are within-category, allowing customers to select any two products from the same category for the discount. Others are cross-category, permitting the selection of two products from any of the specified categories. For instance, the “2FOR249” offer for pants is a within-category bundle, where purchasing two pairs of pants would cost a customer 249 Turkish Liras (TL). In contrast, the “2FOR149” offer for shirts and polos exemplifies a cross-category bundle

	2FOR59	2FOR79	2FOR99	2FOR149	2FOR179	2FOR199	2FOR249
Pants							x
Shirts				x	x		
T-Shirts	x	x	x				
Polos				x	x		
Fashion Pants						x	

Table 3.1: The table above outlines 7 types of bundle discounts offered during the Markdown Season. Each bundle type's name reflects the qualifying criteria and the bundled price point. For instance, '2For59' in the T-shirts category implies a bundle offer of 2 T-shirts for 59 Turkish Liras (TL), equating to an average additional discount of 10.11% beyond the single-item discount. The within-category bundles include '2For59', '2For79', and '2For99' for T-Shirts, '2For199' for fashion pants, and '2For249' for pants. The cross-category bundles, applicable to shirts and polos, are '2For149' and '2For179'.

discount. Under this offer, a bundle consisting of a shirt and a polo, a shirt and a shirt, or a polo and a polo would be priced at 149 TL.

Table 3.2 presents summary statistics relevant to our study, offering initial insights into how bundle and single-item discounts might impact customer purchasing and return behaviors. The data suggest that bundle discounts, which only exist in the Markdown Season, might lead to customers purchasing more items per transaction basket and that items sold under bundle discounts generally have a lower return rate compared to those with single-item discounts. However, a higher rate of whole basket returns with bundle discounts indicates that customers may be more likely to return entire orders under these conditions, likely to avoid paying the full price for items they wish to keep. To delve deeper into these behaviors, we employ a regression model to analyze the impact of discount strategy on sales and product returns.

3.3.3 Initial Insights into Customers' Purchasing and Return Decisions

Utilizing a regression model, we examine how discount strategies affect sales and returns across various product categories. Our analysis reveals that both bundle discounts and single-item discounts yield higher sales volumes. However, the effect on return quantities

	Pre-Markdown Season		Markdown Season	
Purchases	mean	std. dev.	mean	std. dev.
Count of items per basket	1.23	0.54	1.44	0.74
Weekly sales quantity per product	199.87	158.72	235.14	278.82
Prices				
Pants	172.34	11.54	154.37	31.91
Fashion pants	133.75	26.74	110.26	16.97
Polos	85.89	12.44	79.57	10.88
Shirts	119.78	23.65	101.44	18.58
T-shirts	44.68	11.19	42.09	9.65
	Single-item Discount		Bundle Discount	
Returns	mean	std. dev.	mean	std. dev.
2-5 Product return rate	0.0903	0.0408	0.0734	0.0284
Whole basket return rate	0.0132	0.0046	0.0202	0.0039

Table 3.2: The table presents summary statistics for both the Pre-Markdown Season and the Markdown Season, the former referring to the period preceding the Markdown Season. During the Pre-Markdown Season, only single-item discounts were offered. In contrast, the Markdown Season featured additional bundle discounts, as detailed in Table 3.1. "Whole basket return" refers to instances where the entire order is returned. This rate is calculated by dividing the total number of whole-basket returns by the total number of orders placed.

is not uniform across categories: while the return rate for pants shows a significant decrease under both bundle discount and single-item discount, it notably increases for T-shirts and polos under bundle discount. These findings not only highlight the diverse effects of bundle discounts but also indicate opportunities for more effective discount strategy design.

To assess the impact of discount strategies on sales quantities and return rates, we conduct an analysis comparing the effects of bundle discounts and single-item discounts. The data is aggregated at the week-product level, and we utilized the following regression model:

$$y_{jt} = \beta_1 + \beta_2 \text{SingleDiscount}_{jt} + \beta_3 \text{BundleDiscount}_{jt} \times \text{BundleType}_j + \text{Item}_j + \text{Week}_t + \zeta_{jt},$$

where j represents a product and t denotes a week. For each of the five product categories, the dependent variable y_{jt} represents either the logarithm of weekly sales quantity or the logarithm of return rate (in percentage) for product j. The variable $\text{SingleDiscount}_{jt}$ is defined as the percentage discount for product j under a single-item discount in week t, while $\text{BundleDiscount}_{jt}$ represents the percentage discount for product j under a bundle discount in the same week. BundleType_j refers to the bundle types described in Table 3.1. Item_j represents item-specific fixed effects, and Week_t accounts for time-specific fixed effects. The error term is represented by ζ_{jt}.

The primary coefficients of interest are β_2 and β_3, which capture the average effect of single-item discount and bundle discount, respectively, on sales quantities and returns rates. Specifically, as SingleDiscount and BundleDiscount are constructed as the percentage discounts, estimating β_2 and β_3 facilitates a direct comparison of discount depth offered in terms of the two different discount strategies. By focusing on these coefficients, our analysis aims to provide a clearer understanding of how the form of discount (i.e., single-item vs. bundle), rather than discount depth, influence consumer behavior in terms of both purchasing and returning products. We estimate the regression model using the full 21 weeks data, with the Markdown Season spanning the final 6 weeks. During the weeks leading up to the Markdown Season, no bundle discounts were available, allowing us to isolate and identify β_2, the coefficient for SingleDiscount. The coefficient for BundleDiscount β_3 can thus be determined using data from the Markdown Season.

The results, as shown in Table 3.3, indicate that deeper discounts, whether they are single-item discount or bundle discount, are generally associated with an increase in sales quantity. However, the effect of these discounts on return rates is less straightforward. While steeper single-item discounts are linked to a decrease in the return rate of pants, they do not significantly influence the return rates of other product categories. In the case of bundle discounts, a similar pattern is observed for pants, with return rates diminishing alongside more substantial discounts. Yet, for T-shirts and Polos, the opposite effect is noted: increased bundle discounts are associated with higher return rates.

Category	Parameters	log(Sales)	log(Return Percentage)
Pants	SingleDiscount	0.0292***	-0.0069***
		(0.002)	(0.002)
	BundleDiscount×2FOR249	0.0550***	-0.0145***
		(0.005)	(0.005)
Shirts	SingleDiscount	0.0091***	0.0115
		(0.002)	(0.060)
	BundleDiscount×2FOR149	0.0100*	-0.0153
		(0.005)	(0.022)
	BundleDiscount×2FOR179	0.0060*	0.0026
		(0.003)	(0.011)
T-shirts	SingleDiscount	0.0130*	-0.0200
		(0.006)	(0.011)
	BundleDiscount×2FOR59	0.0113	0.0419
		(0.009)	(0.020)
	BundleDiscount×2FOR79	0.0178*	0.052*
		(0.009)	(0.025)
	BundleDiscount×2FOR99	0.0049	0.0795***
		(0.007)	(0.020)
Polos	SingleDiscount	0.0672***	0.032
		(0.019)	(0.028)
	BundleDiscount×2FOR149	0.0731***	0.1071***
		(0.011)	(0.024)
	BundleDiscount×2FOR179	0.0642***	0.1273**
		(0.016)	(0.043)
Fashion Pants	SingleDiscount	0.0533***	0.021
		0.018	(0.022)
	BundleDiscount×2FOR199	0.0769***	0.028
		-0.019	(0.029)

Table 3.3: *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

These diverse outcomes pose intriguing questions about the complex dynamics of consumer behavior in response to different discount forms. In particular, they lead us to explore why bundle discounts, while boosting sales, have mixed effects on returns across product categories but single-item discounts do not. One key difference between bundle discount and single-item discount is its pricing structure and the implication of refund at the return stage. To address such a question, we need to delve deeper into understanding customers' purchase and return decisions, especially how customers update their valuation updates post-purchase and whether they understand the pricing structure of bundle discounts. We develop a customer utility model that encompasses both purchase and return behaviors. This model, which we detail in the subsequent section, is designed to shed light on customer behavior patterns and pave the way for more intelligent and effective discount strategy designs.

3.4 Model

In this section, we develop a model that capture the nuances of a customer's purchase and return decisions, which are related but occur at different times. We describe a customer's interaction with the Retailer as a two-stage journey: the purchase stage, where the customer decides what to buy, followed by the return stage, where the decision is made on whether to keep or return the purchased product(s). Our model highlights customers' valuation updates after their purchases, and the possibility that customers might be inattentive to pricing details when given complex pricing structure.

Our model is built upon three key assumptions:

1. Valuation uncertainty: Customers, upon purchasing clothing items, do not have a perfect understanding of their valuation of the product in the purchase stage. This valuation becomes clearer once the customer has the opportunity to use or experience

the clothing in the return stage. It is this delayed realization of the product's true value to the customer that primarily drives their decision to return it.

2. Forward-looking with bounded rationality: We assume that customers are forward-looking, meaning they consider the future possibility of returning a product at the time of purchase. However, they may exhibit inattentiveness to price details. Specifically, a customer's capability to fully understand the cost implications of making a partial return varies. This aspect acknowledges the cognitive limitations and varying levels of information processing ability among customers.

3. Limited basket size: We limit the model to consider that a customer's purchase at any given time consists of at most two items. This assumption is grounded in the observation that 88% of customers at the Retailer purchase no more than two items at once. Restricting the basket size to two items not only aligns with empirical purchasing patterns but also simplifies the choice set, making it more manageable within the confines of discrete choice models. Allowing for larger basket sizes would exponentially increase the complexity and potentially overburden the model's capacity to provide clear insights.

We begin by presenting a utility model specifically designed for understanding the factors influencing a customer's decision to return a product. Then, we introduce a utility model for a customer's purchase decision, which integrates the customer's anticipation of potential future returns into their purchasing choices. We end this section by providing simulations that illustrate the implications of our model.

3.4.1 Model Setup

In our model, the market of the Retailer has J distinct products, denoted by $j = 1, ..., J$. Each customer is represented by $i = 1, ..., I$. Each customer has the option to select a basket,

$b = 1, \ldots, B$, from the entire range of products available in the market. We constrain each basket to contain a maximum of two items, meaning the size of the basket can either be 1 or 2. Consequently, the total number of possible baskets available in the market is $B = \binom{J}{2} + J$, accounting for all combinations of two-item baskets as well as individual item baskets. For baskets that include two items, we label the items within the basket as $1st$ and $2nd$, such that a two-item basket is denoted as $b = \{1st,\ 2nd\}$. In cases where a basket contains only one item, it is simply represented as $b = \{1st\}$. This setup provides a structured framework to analyze customer choices among a variety of combinations, while maintaining a manageable level of complexity in modeling the potential basket compositions in the Retailer's market.

Each customer's decision to purchase or return is shaped by her individual valuation of a product and, in cases where the basket contains more than one item, by the additional valuation stemming from the interaction between the items in the basket. This interacting valuation term is designed to capture the potential complementarity or substitutability effects that may arise when multiple products are combined within the same basket. We first introduce the return stage model, as a customer's decision is myopic at this stage. Following this, we detail the purchase stage model, wherein a customer's decision integrates her considerations of the potential return stage outcomes.

3.4.2 Return Stage Model

In the return stage, customers assess their valuation of the products based on their actual use and experience. The refund amount associated with each return option is also clear to the customers at this stage.

One-item baskets.

We initiate our analysis with the simpler scenario where a purchased basket contains only a single item. In this case, a customer i already purchased a basket b^* that contains one item $1st$. Her option r at the return stage is binary: to keep or return the item. To be consistent with the two-item basket case, we denote the keep option as "keep 1st" and the return option as "keep none". Thus, the value for each option is represented as follows:

$$\begin{aligned} V(r = keep\ 1st)_i &= v_{i,1st} + z_{i,1st} - \alpha_i P_{b^*,1st} \\ V(r = keep\ none)_i &= 0. \end{aligned} \tag{3.1}$$

The value of return, or keep none, is set to be zero, as the Retailer maintains a generous return policy where no shipping or restocking fee is charged. The value of keeping the item contains three terms. The first term, $v_{i,1st}$, represents the purchase stage valuation of the product, known to the customer at the time of purchase. This term captures the initial perceived value derived from the product based on information available before the purchase. The second term, $z_{i,1st}$, is an update to the customer's valuation post-purchase, reflecting any changes in perception after actually using or experiencing the product. This term measures the discrepancy between a customer's pre-purchase valuation and her re-evaluated post-experience valuation. We assume that $z_{i,1st}$ is not observable by the researcher and only becomes apparent to the customer upon post-purchase experience. The third term, $\alpha_i P_{b^*,1st}$, captures the product's price $P_{b^*,1st}$ (which is the $1st$ and only item in the purchased basket b^*) and the customer's sensitivity to this price, denoted by α_i.

Two-item baskets.

For scenarios where the purchased basket b^* includes two items, the customer faces a broader range of options at the return stage, including the possibility of a partial return. Let r represent the customer's return options, with the value for each option $V(r)_i$ being:

$$\begin{aligned}
&V(r = keep)_i = v_{i,1st} + v_{i,2nd} + v_{ib^*} + z_{i,1st} + z_{i,2nd} + z_{ib^*} - \alpha_i P_{b^*}, \\
&V(r = keep\ 1st)_i = v_{i,1st} + z_{i,1st} - \alpha_i P_{b^*,1st}, \\
&V(r = keep\ 2nd)_i = v_{i,2nd} + z_{i,2nd} - \alpha_i P_{b^*,2nd}, \\
&V(r = keep\ none)_i = 0.
\end{aligned} \tag{3.2}$$

In this model, the option of returning both items (referred to as "*keep none*") is assigned a value of zero. The value of keeping both items (denoted as "*keep*") includes the purchase stage valuations of both items ($v_{i,1st}$ and $v_{i,2nd}$), their valuation updates at the return stage ($z_{i,1st}$ and $z_{i,2nd}$), plus an additional interacting term. This interacting term, v_{ib^*}, represents the combined value of having both items in the same basket at the time of purchase, complemented by a post-purchase valuation update, z_{ib^*}. The value of retaining just one item from the two-item basket (termed as "*keep* 1*st*" or "*keep* 2*nd*") is analogous to the option of keeping an item in a one-item basket scenario, which consists of the purchase stage valuation and the post-purchase valuation update specifically for the retained item.

Once products are brought home, a customer's experience with them is not influenced by the discount type under which they were purchased – be it a bundle or a single-item discount. However, the type of discount becomes crucial when considering returns, especially in cases of partial returns, due to different refund amounts. In particular, the values associated with keeping just one item ($V(r = keep\ 1st)_i$ and $V(r = keep\ 2nd)_i$) differ between bundle and single-item discounts. Under a bundle discount, $P_{b^*,1st} = P_{j,1st}$, which is the original price of the first item in the basket, and similarly, $P_{b^*,2nd} = P_{j,2nd}$, the original price of the second item. Conversely, in the case of a single-item discount, $P_{b^*,1st}$ is priced at $(1-d_j)P_{j,1st}$, the discounted price, where d_j represents the discount fraction for product j. Consequently, the value derived from partially returning an item under a bundle discount is always lower than that under a single-item discount. To gain a deeper understanding of how customers decide whether to return products, we analyze their return options in Figure 3.1.

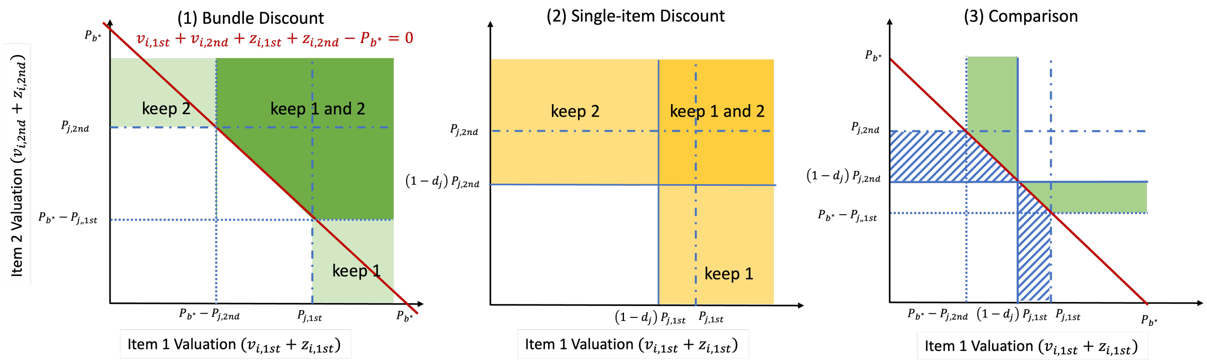

Figure 3.1: Return Decision of a Two-item Basket. This figure visualizes a customer's decision-making process at the return stage when she has purchased a two-item basket. For simplicity, we assume that the customer's price sensitivity $\alpha_i = 1$, and there is no complementarity or substitutability between the two items. The red diagonal line represents the threshold at which consumer surplus is zero; areas to the upper right of this line indicate positive surplus. "Keep 1" and "keep 2" denote the customer's potential decisions to retain either the first or second item in the basket, based on their respective valuations. Plot (3) offers a comparative view by overlaying Plots (1) and (2), highlighting the differences in customer decision outcomes under varying valuations.

In Figure 3.1, the return stage valuations of item 1 and item 2 of a two-item basket are plotted along the x-axis and y-axis, respectively. The total valuation for item 1 is $v_{i,1st} + z_{i,1st}$, and similarly for item 2, it is $v_{i,2nd} + z_{i,2nd}$. We assume price sensitivity $\alpha_i = 1$ and no complementarity or substitutability between these two items for simplicity. Plot (1) in Figure 3.1 illustrates the return options under a bundle discount. The diagonal line marks the borderline where the consumer surplus of such a two-item basket equals zero under bundle discount. The surplus of a basket containing two items is the difference between their combined valuation and the basket's price, i.e., $v_{i,1st} + v_{i,2nd} + z_{i,1st} + z_{i,2nd} - P_{b^*}$. The pentagonal region in the top right quadrant indicates where customers are keeping both items. For comparison, Plot (2) of Figure 3.1 represents the scenario of a single-item discount. In this case, customers keep both items only if $v_{i,1st} + z_{i,1st} \geq 0$ and $v_{i,2nd} + z_{i,2nd} \geq 0$. Notably, the regions indicating the keeping of only one item are larger than in the bundle discount scenario.

Overlaying these two plots, as shown in Plot (3) of Figure 3.1, highlights the differences. The red diagonal line represents the threshold at which consumer surplus is zero under bundle discount. The solid green area marks where customers keep more items under a bundle discount, while the lined region indicates where they would retain more items under a single-item discount. This visual comparison underscores that bundle discounts excel in scenarios where the basket's surplus is positive (to the upper right of the red borderline), resonating with the established wisdom in the literature that bundle discounts efficiently extract consumer surplus. The implication is clear: for bundles to be more effective than single-item discount in terms of reducing returns, they need to offer some room for consumer surplus.

Return decision.

At the return stage, a customer's decision is to select from the set of options as defined in Eq. (3.1) for a one-item basket, or from the options in Eq. (3.2) for a two-item basket. Combining with an idiosyncratic preference term ξ_{ir}, the utility of a customer i choosing return option r is:

$$u_{ir} = V(r)_i + \xi_{ir},$$

where ξ_{ir} is distributed i.i.d. extreme value.

Given that a customer has already purchased basket b^*, the probability that her return decision y_i being option r^* is then

$$p(y_i = r^* | b^*, v_{ij}, z_{ij}, j \in b^*, v_{ib^*}, z_{ib^*}, \alpha_i) = \frac{\exp\left(V(r^*)_i\right)}{1 + \sum_r \exp\left(V(r)_i\right)}. \tag{3.3}$$

3.4.3 Purchase Stage Model

At the purchase stage, a customer's decision-making process involves choosing whether to place an order for a basket containing either one or two products, or to refrain from making a purchase altogether. This decision is based on the utility the customer associates with placing an order for basket b, which is expressed as:

$$u_{ib} = \mathbb{E}[V(Basket)]_{ib} + \varepsilon_{ib}, \tag{3.4}$$

where $\mathbb{E}[V(Basket)]_{ib}$ refers to a customer's expected value of purchasing basket b, and ε_{ib} is included to account for her idiosyncratic preference, which follows an i.i.d. extreme value distribution.

The focus of the model is $\mathbb{E}[V(Basket)]_{ib}$, the expected value of purchasing this basket. At the purchase stage, a customer evaluates the ultimate utility she might derive from using the items in the basket, balancing this against her options to keep or return the items. This consideration involves weighing the probabilities of choosing between the various return options detailed in Section 3.4.2, and integrating the expected utilities from these potential outcomes. The customer's ability in comprehending the pricing details is inherently woven into her decision-making process, influencing her purchase decision.

One-item baskets.

For baskets containing only one item, bundle discounts do not apply to the purchases in such baskets. The only possible discount is single-item discount. To be consistent with the two-item basket case, we denote the item as "1*st*", which is also the only item in basket *b*. The expected value of placing such an order is calculated as follows:

$$\begin{aligned}\mathbb{E}[V(Basket)]_{ib} =& \mathbb{E}[V(keep\ 1st)_{ib}|keep\ 1st]\mathbb{P}(keep\ 1st)_{ib} + \\ & \mathbb{E}[V(keep\ none)_{ib}|keep\ none]\mathbb{P}(keep\ none)_{ib},\end{aligned} \tag{3.5}$$

where

$$\begin{aligned}&V(keep\ 1st)_{ib} = v_{ib,1st} + z_{ib,1st} - \alpha_i P_{b,1st},\\ &V(keep\ none)_{ib} = 0,\end{aligned}$$

and

$$\begin{aligned}&\mathbb{P}(keep\ 1st)_{ib} = \mathbb{P}(V(keep\ 1st)_{ib} \geq V(keep\ none)_{ib})) = \mathbb{P}(v_{ib,1st} + z_{ib,1st} - \alpha_i P_{b,1st} \geq 0),\\ &\mathbb{P}(keep\ none)_{ib} = 1 - \mathbb{P}(keep\ 1st)_{ib}.\end{aligned}$$

In this scenario, the customer knows exactly the price $P_{b,1st}$, which is the amount she pays at the time of purchase for this one item in the basket. Should the customer decide to return the item, she is entitled to a full refund.

Two-item baskets.

When a customer purchases items within a two-item basket b, the discount type – whether it is a single-item discount or a bundle discount – plays a crucial role in determining the pricing implications for returns. With single-item discounts, the process is straightforward: any returned items, whether one or both, are eligible for a full refund. However, the dynamics change under bundle discounts. Here, returning just one item from the bundle results in the loss of the bundle discount, meaning the customer must pay the full price for the item retained. This creates a hidden cost of return for products bought under bundle discounts, a factor that customers might not be fully aware of at the time of purchase. In the subsequent discussion, we construct a model for the customer's expected utility of placing an order, taking into account of potential customer inattentiveness to price details associated with bundle discounts.

The expected value of placing an order of a two-item basket can be articulated as follows:

$$\begin{aligned}\mathbb{E}[V(Basket)]_{ib} = & \mathbb{E}[V'(keep)_{ib} \mid keep]\mathbb{P}(keep)_{ib} + \\ & \mathbb{E}[V'(keep\ 1st)_{ib} \mid keep\ 1st]\mathbb{P}(keep\ 1st)_{ib} + \\ & \mathbb{E}[V'(keep\ 2nd)_{ib} \mid keep\ 2nd]\mathbb{P}(keep\ 2nd)_{ib} + \\ & \mathbb{E}[V'(keep\ none)_{ib} \mid keep\ none]\mathbb{P}(keep\ none)_{ib},\end{aligned} \tag{3.6}$$

where "$keep$" refers to keep both items in the basket, "$keep\ 1st$" refers to keep the $1st$ item only and return the $2nd$, "$keep\ 2nd$ "refers to keep the $2nd$ item only and return the $1st$, and "$keep\ none$" denotes the option to return both items.

The detailed expression of (3.6) is laid out in Appendix 3.10. Below, we focus on $V'(\cdot)$, the value of a return option when a customer anticipates it in the purchase stage. For the return options *keep* $1st$ or *keep* $2nd$, V' does not necessarily equal to V, the return stage options (as expressed in Eq. (3.1) and Eq. (3.2)) because of customers' inattentiveness to price details under bundle discount. For instance, consider a customer who purchases a two-item basket $b = \{1st, 2nd\}$ at a bundle discount price P_b, and later thinks about to keep the $1st$ item while returning the $2nd$ item. At the time of purchase, she may not be fully aware that a partial return will necessitate paying the full price for the $1st$ item. The full price of the $1st$ item might also be unclear at the point of sale as it is sold as part of a bundle. However, upon reaching the return stage, the customer becomes fully aware of the actual price she needs to pay for the item she retains. To reflect this, we define $V'(keep\ 1st)$, the anticipated value of keeping only the $1st$ item, as:

$$V'(keep\ 1st)_{ib} = v_{ib,1st} + z_{ib,1st} - \alpha_i P'_{b,1st} = v_{ib,1st} + z_{ib,1st} - \alpha_i \Delta_i P_{j,1st},$$

where $P'_{b,1st}$ denotes the price that the customer anticipated she would pay should she choose to return item $2nd$, and $P_{j,1st}$ represents the full price of item $1st$ without a bundle discount, which is what she needs to pay at the return stage, where $P'_{b,1st} = \Delta_i P_{j,1st}$.

Intuitively, an inattentive customer may expect that under a bundle discount, she would receive a full refund for any returned item, similar to the refund amount she would receive with a single-item discount. This belief is represented in the model as $P'_{b,1st} = (1 - d_j) P_{j,1st}$, where d_j denotes the depth of the discount for product j (for example, $d_j = 0.2$ for a 20% discount). On the other end of the spectrum, the most attentive customer correctly understands that $P'_{b,1st} = P_{j,1st}$, recognizing that a partial return under a bundle discount would necessitate paying the full price for the retained item. However, most customers likely fall somewhere in between these two extremes. They might be aware that the refund scenario under a bundle discount differs from that of a single-item discount, but they may not precisely calculate the exact price. This intermediate understanding is modeled as

$$\Delta_i = 1 - (1 - \delta_i) d_j,$$

where δ_i ranges from 0 to 1. The variable δ_i captures the level of attentiveness to price of customer i, with 0 indicating complete inattentiveness and 1 representing full attentiveness. Single-item discounts involve no complexity in pricing structure that would challenge a customer's understanding, so we fix δ_i to 0 when an order is placed under single-item discount. Thus, the anticipated value of each return option as perceived by a customer at the purchase stage can be represented as follows:

$$\begin{aligned}
&V'(keep)_{ib} = v_{ib,1st} + v_{ib,2nd} + v_{ib} + z_{ib,1st} + z_{ib,2nd} + z_{ib} - \alpha_i P_b, \\
&V'(keep\ 1st)_{ib} = v_{ib,1st} + z_{ib,1st} - \alpha_i P'_{b,1st} = v_{ib,1st} + z_{ib,1st} - \alpha_i \Delta_i P_{j,1st}, \\
&V'(keep\ 2nd)_{ib} = v_{ib,2nd} + z_{ib,2nd} - \alpha_i P'_{b,2nd} = v_{ib,2nd} + z_{ib,2nd} - \alpha_i \Delta_i P_{j,2nd}, \\
&V'(keep\ none)_{ib} = 0.
\end{aligned}$$

From these expressions, it is evident that $V'(keep)_{ib} = V(keep)_{ib}$ and $V'(keep\ none)_{ib} = V(keep\ none)_{ib}$. This means the anticipated values for the options of keeping both items or returning both items align with their actual values at the return stage, as there is no price ambiguity in these scenarios. However, $V'(keep\ 1st)_{ib} \neq V(keep\ 1st)_{ib}$ and $V'(keep\ 2nd)_{ib} \neq V(keep\ 2nd)_{ib}$, indicating that the anticipated values for keeping just one item out of two differ from their actual values due to customer inattentiveness under bundle discounts.

Purchase decision.

Each customer in our model makes her purchase decision by choosing from all the products or product combinations available in the marketplace, or opting not to make a purchase at all. As setup in Section 3.4.1, a customer chooses from all the available basket b's, which contains one or two items. With the utility model expressed in Eq. (3.4), the conditional probability that a customer i's purchase decision x_i is to buy basket b^* is:

$$p(x_i = b^* | v_{ij}, z_{ij}, v_{ib}, z_{ib}, j \in J, b \in B, \alpha_i, \delta_i) = \frac{\exp\left(\mathbb{E}[V(Basket)]_{ib^*}\right)}{1 + \sum_b \exp\left(\mathbb{E}[V(Basket)]_{ib}\right)}. \tag{3.7}$$

The expected value of a basket, denoted as $\mathbb{E}[V(\text{Basket})]_{ib}$, hinges on several key variables: the valuation at the purchase stage (v), the valuation update at the return stage (z), and the attentiveness to pricing (δ). Intuitively, this expected value increases with both the purchase stage valuation and the return stage valuation update. For baskets purchased under bundle discounts, the expected value diminishes with increased pricing attentiveness; inattentive customers are prone to overvalue the basket as they may not recognize that partial returns could affect the discounted price they initially received. In contrast, the expected value of baskets acquired via a single-item discount is not affected by attentiveness to pricing, as there is no complexity involved in its price structure. Additional simulations are performed to demonstrate the variations in the expected value of a basket based on customers' valuation and attentiveness, as well as to understand the resulting purchase and return behaviors. The findings from these simulations are detailed in the Appendix 3.11.

3.5 Estimation

By estimating the structural model described in Section 3.4, our objective is to learn the valuation parameters and the attentiveness parameter. Below, we first outline our maximum likelihood estimation model, followed by identification and the Monte-Carlo Expectation Maximization (MCEM) method employed to estimate the model's parameters.

3.5.1 Maximum Likelihood Estimation

In our model, customers are differentiated by their attentiveness to pricing details (δ) and price sensitivity (α). We assume that customers' valuations of products vary across different categories, so we aim to estimate valuation distributions that are specific to each product category.

In the market, there are K distinct product categories, denoted as $k = 1, \ldots, K$. For any given product j that belongs to category k, we posit that its purchase stage valuation by a customer follows a normal distribution: $v_{ik} \sim N(\mu_{v_k}, \sigma^2_{v_k})$. Similarly, the valuation update for product j belongs to category k at the return stage is also assumed to follow a normal distribution: $z_{ik} \sim N(\mu_{z_k}, \sigma^2_{z_k})$. In scenarios where a basket contains two items, it is possible for a customer to derive additional value from the interaction between these items. This valuation is hypothesized to vary depending on whether the items in the basket are complements or substitutes. Specifically, we consider tops paired with tops or bottoms paired with bottoms as substitutes, whereas a combination of a top and a bottom is viewed as complements. Accordingly, there are two distinct distributions for interaction valuation. Each interaction valuation distribution is thus expressed as $v_{ic} \sim N(\mu_{v_c}, \sigma^2_{v_c})$, where $c = C$ or S, with C representing complements and S representing substitutes. Similarly, the interaction valuation update is expressed as $z_{ic} \sim N(\mu_{z_c}, \sigma^2_{z_c})$, where $c = C$ or S. Furthermore, we assume that the customer's price sensitivity parameter, α_i, follows a normal distribution: $\alpha_i \sim N(\mu_\alpha, \sigma^2_\alpha)$. For the customer attentiveness parameter, δ_i, we allow for a range between 0 (fully inattentive) and 1 (fully attentive). This parameter is modeled using a beta distribution for purchases made under a bundle discount, $\delta_i \sim Beta(\alpha_\delta, \beta_\delta)$. If a purchase is made under a single-item discount, we simplify the model by setting δ_i to 0. To align the reporting of this parameter with others in our model, we convert the beta distribution's shape parameters, α_δ and β_δ, into its mean μ_δ and standard deviation σ_δ.

Let θ represent the model's parameters. They are partitioned into: θ_{v_k}, which governs the distribution of valuation of product j at the purchase stage that belongs to category k; θ_{z_k}, which governs the distribution of valuation update at the return stage of product j that belongs to category k; θ_{v_c}, which governs the distribution of the valuation of a basket belongs to category combination c at the purchase stage; θ_{z_c}, which governs the distribution of the valuation update of a basket a basket belongs to category combination c at the return stage; θ_α, which governs the distribution of customers' price sensitivity; θ_δ, which governs the distribution of customer attentiveness level to price details. Together, $\theta = \{\theta_{v_k}, \theta_{z_k}, \forall k \in K, \theta_{v_c}, \theta_{z_c}, c = C \text{ or } S, \theta_\alpha, \theta_\delta\}$.

Given the observed purchase choices, $\{x_i : i \in N\}$, and return decisions, $\{y_i : i \in N\}$, we obtain the maximum likelihood estimates:

$$
\begin{aligned}
\theta^*_{MLE} =_\theta & L(x_i, y_i : i \in N \mid \theta) \\
=_\theta & \int \prod_{i \in N} p(x_i | v_{ik}, z_{ik}, v_{ic}, z_{ic}, k \in K, c = C, S, \alpha_i, \delta_i; \theta) \cdot p(y_i | v_{ik}, z_{ik}, v_{ic}, z_{ic}, \alpha_i, b^*; \theta) \\
& \cdot p(v_{ik} \mid \theta_{v_k}, k \in K) \cdot p(z_{ik} \mid \theta_{z_k}, k \in K) \\
& \cdot p(v_{ic} \mid \theta_{v_c}, c = C, S) \cdot p(z_{ic} \mid \theta_{z_c}, c = C, S) \\
& \cdot p(\delta_i \mid \theta_\delta) \cdot p(\alpha_i \mid \theta_\alpha) \\
& \cdot dv_1 \cdots dv_{|K|} \cdot dz_1 \cdots dz_{|K|} \cdot dv_C \cdot dv_S \cdot dz_C \cdot dz_S \cdot d\delta_1 \cdots d\delta_{|N|} \cdot d\alpha_1 \cdots d\alpha_{|N|},
\end{aligned}
\tag{3.8}
$$

where the conditional likelihood of the purchase decision is provided in Eq. (3.7) and the conditional likelihood of the return decision is expressed in Eq. (3.3).

3.5.2 Identification

Our model expands on the typical identifiable discrete choice model by introducing valuation update parameters (θ_{z_k}) that require distinct identification from the initial purchase valuation parameters (θ_{v_k}). The identification of the two groups of parameters relies on leveraging the customers' return decisions as an additional source of information beyond the initial purchase decisions.

At the purchase stage, our model calculates the expected value of a basket. Here, a customer knows her purchase stage valuation (v_{ij}) but faces uncertainty regarding her valuation update (z_{ij}). Thus, the expectation calculation is centered around z_{ij}. The formulation of this expectation is detailed further in the Appendix 3.10. Conversely, in the return stage model, the customer knows both v_{ij} and z_{ij}, which enter into her utility linearly as shown in Eq. 3.1 and Eq. 3.2. By analyzing the combined purchase and return decisions observed for each customer, we can identify θ_{z_k}. Moreover, the parameters for complementarity or substitutability (θ_{v_c} and θ_{z_c}) are identifiable because of the heterogeneity in basket size across customers, i.e., we have customers who purchase a single item j_1 and those who buy both j_1 and j_2. The parameters θ_{v_c} and θ_{z_c} are relevant only in the latter scenario. The parameters thus identified describe how an average customer would behave in terms of purchase and return within such a market.

3.5.3 MCEM Algorithm

To calculate the maximum likelihood estimator θ^*_{MLE} as detailed in Eq. 3.8, we employ the Monte-Carlo Expectation Maximization (MCEM) algorithm. This algorithm serves as an iterative computational solution for situations where direct maximization of the likelihood function is infeasible. It involves a set of latent variables that bridge the parameters we want to estimate the observed outcomes.

The likelihood function described in Eq. 3.8 is defined jointly by the observed outcomes $\{x_i, y_i : i \in N\}$, alongside the model's latent variables $\{v_{ik}, z_{ik}, v_{ic}, z_{ic}, \delta_i, \alpha_i, : i \in N, k \in K\}$. The MCEM algorithm leverages this joint likelihood, considering the latent variables as unobserved outcomes. For the sake of clarity within the algorithm's context, we refer to the observed data $\{x_i, y_i : i \in N\}$ collectively as O and the latent variables $\{v_{ik}, z_{ik}, v_{ic}, z_{ic}, \delta_i, \alpha_i : i \in N, k \in K\}$ as Y. We then define $L(O, Y \mid \theta)$ as the joint likelihood function of both O and Y, evaluated at the parameter values θ, effectively capturing the essence of the likelihood function.

Each iteration t of the algorithm updates a current parameter value that we label as θ_t. Specifically,

$$\theta_{t+1} =_\theta \mathbb{E}_{Y|O,\theta_t}[\log L(O, Y \mid \theta)]. \tag{3.9}$$

Thus, each round of the algorithm chooses the next parameter value θ_{t+1} to maximize a specially constructed expectation. Typically, in practice, the expectation is evaluated through simulation. In this case,

$$\theta_{t+1} =_\theta \sum_{b=1}^{B} \log L(O, Y_b \mid \theta), \tag{3.10}$$

where the values of Y_b, $b = 1, \ldots, B$, are drawn, by simulation, from the posterior distribution of the latent variables Y conditional on O and θ_t.

In order to compute θ_{t+1}, the algorithm first performs an E-step that samples such Y_b, $b = 1, \ldots, B$ and constructs the objective function. Then, the M-step performs the maximization of Eq.3.10. The E-step and the M-step are iteratively performed until no significant changes are observed between iterations. This iterative refinement process ensures that with each cycle, the parameter values θ_{t+1} inch closer to the maximum likelihood es-

	mean	std. dev.		
Price sensitivity (α)	0.07	0.02		
Attentiveness (δ)	0.39	0.04		
	Purchase stage (v)		Return stage (z)	
Valuations	mean	std. dev.	mean	std. dev.
Pants	22.08	4.62	4.05	2.37
Fashion pants	19.28	5.63	1.51	1.21
Polos	4.78	1.68	-0.74	0.31
Shirts	5.54	1.56	-0.23	0.25
T-shirts	2.13	1.41	-0.25	0.14
Substitutes	0.28	0.34	-1.41	0.30
Complements	1.53	0.30	3.24	0.20

Table 3.4: The reported distribution of the attentiveness parameter δ pertains specifically to customers who made purchases under bundle discounts. The price sensitivity parameter α has been estimated across the entire customer base in the market. Valuation parameters have been calculated for each individual product category, as well as for the interactions occurring within a basket. In this context, "substitutes" refer to baskets comprising two items that are either both tops or both bottoms, whereas "complements" refer to baskets that include one top and one bottom.

timates, effectively optimizing the likelihood function based on both the observed data O and the estimated distribution of the latent variables Y. The convergence of the algorithm signifies that the parameter estimates have reached a level of stability, offering a reliable set of parameters θ^*_{MLE} that best describe the underlying data generation process within the specified model framework.

3.6 Estimation Results

In this section, we present our estimation results derived from the structural model detailed in Section 3.4. Our analysis reveals an interesting pattern in how customers update their valuation of products at the return stage. For product categories that customers initially value higher than their sales prices at the purchase stage, there is a tendency for them to positively update their valuation at the return stage. This positive evaluation update renders the purchase even more appealing. Conversely, in categories where customers initially have a lower valuation, they tend to adjust their valuation downwards at the return stage, further diminishing the perceived value of the purchase. Additionally, a significant insight from our findings is the generally inattentiveness to price details exhibited by customers. This inattentiveness often results in customers inadequately accounting for the hidden costs associated with partial returns, leading them to overvalue a bundle at the purchase stage. Consequently, customer inattentiveness contributes to irrational purchasing decisions, with implications for both sales and returns of a retailer.

The estimation results are presented in Table 3.4. The attentiveness parameter reveals customers' inattentiveness to price details, with δ averaging 0.39 (on a scale from 0 to 1). While customers do consider potential future returns, our results suggest they generally struggle with accurately calculating the exact price required for a partial return. Consequently, they tend to overestimate the value of a bundle, which could lead to impulsive purchases and, potentially, an increase in returns.

For product valuations, we estimate the distributions of both the purchase stage valuation and the return stage valuation update for each product category k. We notice that the means of the valuation updates post-purchase (z) are considerably smaller than those of the initial valuations at the purchase stage (v). This observation suggests a relatively minor shift in customer valuations after purchase, implying that changes in customer perceptions post-purchase are not substantial. This phenomenon could be attributed to the nature of the data collected from brick-and-mortar stores, where customers have the opportunity to physically interact with and try on the products.

To further contextualize these results, we define a customer's mean myopic surplus at the purchase stage, $\mu_{v_k} - \mu_\alpha \bar{p}_k$, as a baseline. It represents the difference between the mean purchase stage valuation of a customer and the mean price of a product category $\bar{p}_k$, adjusted by the mean price sensitivity μ_α. As shown in Table 3.5, categories like pants and fashion pants exhibit a positive mean myopic surplus, categorizing them as high

	Purchase stage valuation mean	Return stage valuation mean	Mean myopic surplus ($\mu_{v_k} - \mu_\alpha \bar{p}_k$)	Surplus category
Pants	22.08	4.05	11.27	High
Fashion pants	19.28	1.51	11.56	High
Polos	4.78	-0.74	-0.79	Low
Shirts	5.54	-0.23	-1.56	Low
T-shirts	2.13	-0.25	-0.82	Low

Table 3.5: The values of purchase stage valuation mean (μ_{v_k}), return stage valuation mean (μ_{z_k}), and price sensitivity mean (μ_α) are taken from Table 3.4. The average price of each category ($\bar{p}_k$) is the mean price during the Markdown Season as documented in Table 3.2.

surplus categories. Conversely, categories such as polos, shirts, and T-shirts demonstrate a negative mean myopic surplus, placing them in the low surplus categories. Intriguingly, these classifications align with the mean return stage valuation updates: positive surplus categories exhibit positive mean updates, while negative surplus categories show negative mean updates. This pattern suggests a form of confirmation bias, where initial perceptions at purchase may influence post-purchase feelings about the products.

In addition, we observed that the interaction valuation terms for baskets containing either substitutes or complements are positive, albeit with substitutes showing a smaller magnitude. This somewhat counterintuitive finding might stem from the increased likelihood of customers noticing and purchasing items when they are from categories that are bundled together, thereby accruing added value that supersedes their substitutive nature. However, at the return stage, the substitutive characteristics emerge, negatively impacting a basket's net value. The relatively minor magnitudes of these interaction valuations suggest that, while present, their impact on the overall value of a basket is limited.

3.7 Counterfactual Analysis

The estimation results in Section 3.6 uncovers significant potential for optimizing the Retailer's bundle offerings. Based on the estimated customer valuation parameters of each product category, many current bundles fail to deliver a net surplus post-purchase, which may induce customer returns. By redesigning bundles to simply ensure a positive post-purchase net surplus, we project an increase in net sales quantity by 15.6%. Additionally, there exists an opportunity for the Retailer to enhance customer attentiveness about pricing, which our counterfactual analysis suggests could lower the return rate by 20.88% with little impact on net sales.

For the Retailer, we propose a more comprehensive and flexible bundle discount strategy that allows for within-category and cross-category bundling, provided that the mean net surplus post-purchase is positive. It allows customer to freely choose a product from the high surplus categories as shown in Table 3.5, and bundle it with any other products. This approach creates bundles involving all categories, but deliberately avoiding the creation of bundles within or across low surplus categories like polos, shirts, or T-shirts. For example, a bundle combining a pair of pants and a shirt would be acceptable, as would a bundle of two pairs of pants, but a bundle comprising two shirts would not be considered. A comparison between our proposed bundle strategy and the Retailer's current approach is illustrated in Figure 3.2, highlighting the discrepancy between current bundles and those meeting our positive net surplus criteria.

	Pants	Fashion Pants	Polos	Shirts	T-shirts
Pants	✓✓	✓	✓	✓	✓
Fashion Pants	✓	✓✓	✓	✓	✓
Polos	✓	✓	✓		
Shirts	✓	✓	✓	✓	
T-shirts	✓	✓			✓

Figure 3.2: Proposed bundles vs. Current bundles. The proposed bundle combinations are denoted with black check marks, whereas the current bundle combinations are denoted with red check marks. The current bundles are the same as those listed in Table 3.1.

Through our counterfactual analyses, we delve into comparing the sales and revenue of our proposed bundle strategy against the Retailer's current strategy, utilizing the single-

	Sales	Net Sales	Net Revenue	Return Rates
Current bundles (change in %)	-1.19	-0.69	-0.07	5.33
Proposed bundles (change in %)	15.50	15.61	12.58	-1.60

Table 3.6: The sales, net sales, and net revenue reported for the baseline case represent averages derived from 10 trials, each involving 10,000 customers making purchase and return decisions. The numbers provided are percent changes relative to those under single-item discount, where all products are sold under discounted prices.

item discount as a benchmark for this comparison. The single-item discount serves as a straightforward baseline, offering every product at a reduced price without necessitating bundling as a prerequisite for the discount. The outcomes, including sales quantities, net sales, net revenue post-returns, and return rates, are presented in Table 3.6. Our findings indicate that the Retailer's current bundle strategy slightly underperforms compared to the single-item discount. This is unsurprising, given that the majority of current bundles have low net surpluses. In contrast, our proposed bundle strategy leads to a 15.6% increase in net sales quantity and a 12.6% increase in net revenue, accompanied by a 1.6% reduction in the return rate. These results suggest that at the current level of customer attentiveness ($\delta \sim N(0.39, 0.04)$), customers are discerning enough to understand the real value of bundles. As a result, bundles that offer a post-purchase positive net surplus outperform the single-item discount strategy.

Next, we examine the influence of attentiveness levels on purchase and return decisions. To do this, we simulated customer behavior under two distinct scenarios: one where customers are completely inattentive ($\delta = 0$), meaning they assume products bought under a bundle discount will still qualify for the discount even if they make a partial return, and another where customers are completely attentive ($\delta = 1$), fully aware of the refund amount for partial returns. We use the completely inattentive scenario ($\delta = 0$) as our baseline and compare it with the estimated attentiveness level and the completely attentive scenario. The results are detailed in Table 3.7.

	Strategy	Estimated Attentiveness (change in %)	Full Attentiveness (change in %)
Sales	Current bundle	-0.43	-0.81
	Proposed bundle	-1.38	-2.63
Net Sales	Current bundle	-0.40	-0.56
	Proposed bundle	-0.04	-0.05
Net Revenue	Current bundle	-0.28	-0.43
	Proposed bundle	0.34	0.37
Return Rates	Current bundle	-0.25	-2.08
	Proposed bundle	-10.75	-20.88

Table 3.7: The sales, net sales, and net revenue reported for the baseline case represent averages derived from 10 trials, each involving 10,000 customers making their purchase and return decisions. The numbers shown represent the percent changes relative to those under single-item discount as a baseline. The baseline remains constant across varying levels of δ because purchases and returns under single-item discount are independent of the attentiveness level.

We observe that the most attentive customer base tends to purchase 2.6% less than their inattentive counterparts. However, this decrease is marginal (only 0.05%) when considering net sales. Moreover, the net revenue increases slightly by 0.37%. These positive outcomes are attributed to a significant reduction in return rates, which drops by 20.88% among attentive customers compared to inattentive customers. In the more realistic scenario, where customers are neither completely inattentive nor entirely attentive, the return rate can be reduced by 10.75%, with a 0.04% decrease in net sales but a 0.35% increase in net revenue. This suggests that more attentive customers tend to keep higher-priced products, resulting in a more profitable customer base for retailers. Therefore, fostering a more attentive customer base can not only significantly diminish return rates but also slightly enhance net revenue, with minimal impact on net sales.

As discussed in Section 3.6, the primary influence on customers' valuations of the Retailer's products stems from their initial valuations at the purchase stage, which is likely attributed to the dataset originating from brick-and-mortar stores. In such settings, the tangible experience of physically examining products likely informs these valuations. However, the net surplus of a basket is influenced by various factors, including a customer's initial valuations at the purchase stage, their valuation updates at the return stage, and the complementarity between items in the basket. In the context of online retail, the inherent uncertainty a customer faces regarding the products she acquires is expected to be significantly higher compared to in-store purchases, thereby amplifying the impact of her post-purchase valuation updates. This increased uncertainty accentuates the importance of incorporating considerations of returns and customer attentiveness to pricing details in devising a retailer's discount strategy. The subsequent section will delve deeper into the implications of these factors in retail environment where high level of valuation uncertainty exists, highlighting their critical role in shaping effective discount strategies.

3.8 Simulations and Managerial Insights

The estimation results from Section 3.6 suggest that the Retailer's customers generally exhibit smaller magnitude of post-purchase valuation updates compared to their initial purchase-stage valuations. This could be attributed to the data originating from the Retailer's brick-and-mortar stores, where the tangible experience allows customers to better assess and try products, potentially reducing the inclination to revise valuations after the fact. Yet, the online shopping environment presents a different scenario, characterized by greater uncertainty around the product's attributes and how well they match personal preferences, which can significantly affect product valuation post-purchase and lead to higher returns.

In this section, we explore the impact of a retailer's discount strategies on sales and net sales quantities within retail environments characterized by a high degree of uncertainty. Through the use of simulated generic product categories, we aim to ascertain the most effective discount strategy, taking into consideration varying levels of customer attentiveness to pricing. Our analysis reveals that accounting for returns significantly changes a retailer's choice of discount strategy. Moreover, customizing discount strategies to align with customer attentiveness to pricing can notably improve overall sales performance. A customer base inattentive to pricing details tends to generate higher sales volumes, while a more pricing-attentive customer base is associated with increased net sales quantities. Engaging with attentive customers also grants retailers the flexibility to devise bundle offerings across a wider range of product categories, expanding to include products of lower value that might otherwise be difficult to sell. Cultivating customers' understanding of pricing details and thereby increasing transparency of bundle pricing structure proves to be beneficial for retailers.

Consider a hypothetical scenario where a retailer, drawing on historical sales data, has calibrated its pricing to closely match customer valuations at the time of purchase for four product categories: A, B, C, and D. However, the retailer faces a challenge in predicting how customers' valuations of these categories might shift after the purchase. To mirror this situation, we assume that the retailer sets discounted prices to align with customers' initial valuations, yet remains uncertain about the post-purchase appeal of the product categories. In our simulations, we explore three scenarios (Table 3.8) where customers experience varying degrees of mean valuation update post-purchase, up to a 50% change from the initial product valuation. These scenarios are designed to minimize mean consumer surplus at the purchase stage to nearly zero, allowing for a focused examination of how post-purchase valuation updates influence consumer behavior. According to the model described in Section 3.4, each customer independently determines her valuation (v) and valuation update (z) for each potential one or two-item basket from the four categories. For clarity, we exclude potential complementarity or substitutivity effects within baskets.

Our analysis of bundle discounts versus single-item discounts reveals the comparative sales quantities under each strategy. With single-item discounts, all products are offered at a reduced price. In contrast, bundle discounts apply only when customers purchase two

Product Categories	Scenario I Balanced Updates	Scenario II Positive Updates	Scenario III Negative Updates
A	+50%	+50%	0%
B	+25%	+25%	-12.5%
C	-25%	+12.5%	-25%
D	-50%	0%	-50%

Table 3.8: Each percentage value represents the average change in customer valuation at the return stage. For example, if the mean of the purchase stage valuation (v) is 80, then a +50% update implies that the mean of valuation update (z) is 40. Conversely, a -25% update indicates that the mean of valuation update is -20.

items from the specified categories within the bundle, without extending the discount to individual item purchases. Items from categories not included in the designated bundle still benefit from single-item discounts. For example, in the scenario involving an AB bundle, a customer who buys any combination of two items from categories A or B is eligible for the bundled discount. However, should they opt to purchase only a single item from either category, the price reverts to its original rate. Meanwhile, items from categories C and D, regardless of being purchased singly or in pairs, receive a single-item discount since they fall outside the bundle's scope.

We consider four different bundle designs: A, AB, ABC, and ABCD, and compare their performances in terms of sales and net sales with respect to single-item discount. These particular designs were selected based on the sequence of mean valuation updates, starting with category A, which exhibits the highest update, and subsequently including additional categories in order of decreasing mean valuation. These bundle discounts render the included products less attractive compared to an overarching single-item discount strategy, as single purchases of the bundled items are default to the original price.

Figure 3.3 demonstrates the most advantageous discount strategies for maximizing sales and net sales quantities. Our analysis distinguishes between strategies suited for inattentive and attentive customer bases, with varying responses to pricing intricacies. The shaded areas indicate which customer base – attentive or inattentive – exhibits greater sales volumes

	Scenario I (Balanced Updates)		Scenario II (Positive Updates)		Scenario III (Negative Updates)	
	Sales	Net Sales	Sales	Net Sales	Sales	Net Sales
Attentive	ABC	ABCD	ABC	ABCD	A	Single-item
Inattentive	ABCD	AB	ABCD	ABC	ABCD	A

Figure 3.3: The three scenarios are detailed in Table 3.8. The labels represent discount strategies, with A, AB, ABC, and ABCD referring to various bundle discount configurations and "Single-item" referring to single-item discount. The discount strategies shown above represents the strategy that provides the retailer with the highest sales or net sales quantities. Attentive customers are fully attentive with $\delta = 1$, whereas inattentive customers have $\delta = 0$. The shaded areas in the figure indicate which customer group – attentive or inattentive – generates higher sales or net sales quantities in each scenario.

under each strategy. Further details on the underlying simulations are provided in Appendix 3.12. Generally, bundle discounts surpass single-item discounts in effectiveness, except in scenarios with negative post-purchase updates. This finding confirms our earlier assertion that bundle discounts are preferable when the post-purchase net surplus is non-negative.

Focusing on scenarios that bundle discount is the preferred over single-item discount (balanced or positive updates), the choice of an optimal bundle strategy is influenced by customer attentiveness. A comparison of sales versus net sales reveals that, a retailer makes more sales with inattentive customers but makes more net sales with attentive customers. This suggests that targeting inattentive customers may initially seem more profitable based on sales alone; however, when returns are factored in, catering to a base of attentive customers becomes more profitable.

In examining net sales more closely, the versatility of the optimal bundle appears to hinge on customer attentiveness too. In scenario I (balanced updates), the comprehensive ABCD bundle achieves the highest net sales increase among attentive customers, whereas the AB bundle stands out for inattentive customers. Similarly, scenario II (positive updates) identifies the ABCD bundle as the optimal choice for attentive customers, while the ABC bundle fares best for inattentive ones. These results demonstrate that attentive customers can navigate complex pricing structures and discern true value. For retailers, this suggests a tactical approach: with an attentive customer base, broader and more inclusive bundle discounts can be offered, spanning a range of product categories with varying valuation levels. Conversely, for an inattentive customer base, it is more beneficial to provide bundles that focus on categories with higher inherent valuations, mitigating the potential for these customers to overestimate the bundle's value due to the complexity of the pricing structure.

The implication is twofold: attentive customers not only lead to higher net sales and fewer returns, facilitating effective inventory management during sales seasons, but they also enable retailers to design bundles that encompass lower-value categories, assisting in the sale of products that might otherwise be difficult to sell. These findings suggest that cultivating customer rationality in the retail setting, especially by improving price transparency, can actually benefit the retailer.

3.9 Concluding Remarks

We examine how retailers should choose their discount strategy taking into consideration of returns and customers' inattentiveness to pricing. We find that returns significantly influence the discount strategy a retailer should adopt. To effectively move inventory, a retailer must account for both the ex-ante valuation that drives initial sales and the ex-post valuation that influences returns. Central to managing this balance is understanding how customers update their valuation post-purchase. We provide empirical evidence of a confirmation bias in customer behavior, suggesting that customers are likely to positively update valuations for products they initially favored and negatively for those they did not.

The issue of customer inattentiveness to complex pricing structures emerged as a critical factor in determining a retailer's discount strategy. Our empirical analysis indicates that customers typically overlook intricate pricing details. Data from brick-and-mortar stores show that enhancing customers' understanding of pricing can significantly reduce return rates by 20.9% with only a minimal impact on net sales. Simulations further illustrate that under high post-purchase uncertainty retail settings, attentive customers contribute more to net sales than their inattentive counterparts. Furthermore, an attentive customer base permits the retailer to offer a broader array of bundles, which can effectively include products that are otherwise less appealing.

For retailers, this means that when the primary objective is to clear inventory, particular attention should be given to ex-ante customer valuations. In addition, when discount strategies involve complex pricing structures, it becomes necessary to factor in the general inattentiveness of customers. Our findings suggest a novel perspective on price transparency: How the structure of pricing could inherently obfuscate full comprehension of the customers and thus influence sales and returns. For retailers, cultivating customer attentiveness to pricing is a strategic move that can lead to better sales performance and inventory management.

Online Supplement for "Smarter Bundles Need Smarter Customers"

3.10 Appendix I: Expression of Expected Value of a Two-item Basket

We provide the details of the expected value of a two-item basket in this appendix. As stated in Eq. 3.6, a customer *i*'s expectation of the value of a two-item basket *b* is the following:

At the purchase stage, a customer anticipates her valuation of the products and the basket in the return stage and makes her purchasing decision. Her expectation of valuation of a order of basket b can be expressed as

$$\begin{aligned}\mathbb{E}[V(Basket)]_{ib} =&\mathbb{E}[V'(keep)_{ib} \mid keep]\mathbb{P}(keep)_{ib}+\\&\mathbb{E}[V'(keep\ 1st)_{ib} \mid keep\ 1st]\mathbb{P}(keep\ 1st)_{ib}+\\&\mathbb{E}[V'(keep\ 2nd)_{ib} \mid keep\ 2nd]\mathbb{P}(keep\ 2nd)_{ib}+\\&\mathbb{E}[V'(keep\ none)_{ib} \mid keep\ none]\mathbb{P}(keep\ none)_{ib},\end{aligned} \tag{3.11}$$

where "$keep$" refers to keep both items in the basket, "$keep$ $1st$" refers to keep the $1st$ item only and return the $2nd$, "$keep$ $2nd$ "refers to keep the $2nd$ item only and return the $1st$, and "$keep\ none$" denotes the option to return both items.

In Eq. 3.11, we use $V'(\cdot)$ to denote the value of the return options when a customer anticipates it in the purchase stage. In particular, for each of the return options, we have:

$$\begin{aligned}&V'(keep)_{ib} = v_{ib,1st} + v_{ib,2nd} + v_{ib} + z_{ib,1st} + z_{ib,2nd} + z_{ib} - \alpha_i P_b,\\&V'(keep\ 1st)_{ib} = v_{ib,1st} + z_{ib,1st} - \alpha_i P'_{b,1st} = v_{ib,1st} + z_{ib,1st} - \alpha_i \Delta_i P_{j,1st},\\&V'(keep\ 2nd)_{ib} = v_{ib,2nd} + z_{ib,2nd} - \alpha_i P_{b,2nd} = v_{ib,2nd} + z_{ib,2nd} - \alpha_i \Delta_i P_{j,2nd},\\&V'(keep\ none)_{ib} = 0,\end{aligned}$$

where $\Delta_i = 1-(1-\delta_i)d_j$. δ_i captures the level of attentiveness to price of customer i, ranges from 0 to 1.

When making a purchase decision, a customer knows her valuation of the product v_{ij}, but she is uncertain about the valuation update z_{ij}. Thus, the expectation of $\mathbb{E}[V(Basket)]_{ib}$

is about z_{ij}. The probability of keeping both items is:

$$
\begin{aligned}
\mathbb{P}(keep)_{ib} =& \mathbb{P}(V'(keep)_{ib} \geq \max(V'(keep\ 1st)_{ib}, V'(keep\ 2nd)_{ib}, V'(keep\ none)_{ib})) \\
=& \mathbb{P}(v_{ib,1st} + v_{ib,2nd} + v_{ib} + z_{ib,1st} + z_{ib,2nd} + z_{ib} - \alpha_i P_b \\
& \geq \max(v_{ib,1st} + z_{ib,1st} - \alpha_i \Delta_i P_{j,1st}, v_{ib,2nd} + z_{ib,2nd} - \alpha_i \Delta_i P_{j,2nd}, 0)) \\
=& \mathbb{P}(z_{ib} \geq \alpha_i (P_b - \Delta_i P_{j,1st}) - v_{ib,2nd} - z_{ib,2nd} - v_{ib}, \\
& z_{ib,1st} + z_{ib} \geq \alpha_i (P_b - \Delta_i P_{j,2nd}) - v_{ib,1st} - v_{ib}, \\
& z_{ib,1st} + z_{ib,2nd} + z_{ib} \geq \alpha_i P_b - v_{ib,1st} - v_{ib,2nd} - v_{ib}),
\end{aligned}
$$

and the expectation valuation of keep both is

$$
\begin{aligned}
& \mathbb{E}[V'(keep)_{ib} \mid keep] \\
=& \mathbb{E}[v_{ib,1st} + v_{ib,2nd} + v_{ib} + z_{ib,1st} + z_{ib,2nd} + z_{ib} - \alpha_i P_b \mid v_{ib,2nd} + z_{ib,2nd} + v_{ib} + z_{ib} \geq \alpha_i (P_b - \Delta_i P_{j,1st}), \\
& v_{ib,1st} + z_{ib,1st} + v_{ib} + z_{ib} \geq \alpha_i (P_b - \Delta_i P_{j,2nd}), v_{ib,1st} + v_{ib,2nd} + v_{ib} + z_{ib,1st} + z_{ib,2nd} + z_{ib} \geq \alpha_i P_b] \\
=& v_{ib,1st} + v_{ib,2nd} + v_{ib} - \alpha_i P_b + \mathbb{E}[z_{ib,1st} + z_{ib,2nd} + z_{ib} \mid z_{ib} + z_{ib,2nd} \geq \alpha_i (P_b - \Delta_i P_{j,1st}) - v_{ib,2nd} - v_{ib}, \\
& z_{ib,1st} + z_{ib} \geq \alpha_i (P_b - \Delta_i P_{j,2nd}) - v_{ib,1st} - v_{ib}, z_{ib,1st} + z_{ib,2nd} + z_{ib} \geq \alpha_i P_b - v_{ib,1st} - v_{ib,2nd} - v_{ib}].
\end{aligned}
$$

The probability of keeping the first item only is:

$$
\begin{aligned}
&\mathbb{P}(keep\ 1st)_{ib} = \mathbb{P}(V'(keep\ 1st)_{ib} \geq \max(V'(keep)_{ib}, V'(keep\ 2nd)_{ib}, V'(keep\ none)_{ib}))\\
=&\mathbb{P}(v_{ib,1st} + z_{ib,1st} - \alpha_i \Delta_i P_{j,1st}\\
\geq&\max(v_{ib,1st} + v_{ib,2nd} + v_{ib} + z_{ib,1st} + z_{ib,2nd} + z_{ib} - \alpha_i P_b, v_{ib,2nd} + z_{ib,2nd} - \alpha_i \Delta_i P_{j,2nd}, 0))\\
=&\mathbb{P}(z_{ib,2nd} + z_{ib} \leq \alpha_i (P_b - \Delta_i P_{j,1st}) - v_{ib,2nd} - v_{ib},\\
&z_{ib,1st} - z_{ib,2nd} \geq \alpha_i \Delta_i (P_{j,1st} - P_{j,2nd}) - v_{ib,1st} + v_{ib,2nd},\\
&z_{ib,1st} \geq \alpha_i \Delta_i P_{j,1st} - v_{ib,1st}),
\end{aligned}
$$

and the expected value of keep first:

$$
\begin{aligned}
&\mathbb{E}[V'(keep\ 1st)_{ib} \mid keep\ 1st]\\
=&v_{ib,1st} - \alpha_i \Delta_i P_{j,1st} + \mathbb{E}[(z_{ib,1st} \mid z_{ib,2nd} + z_{ib} \leq \alpha_i (P_b - \Delta_i P_{j,1st}) - v_{ib,2nd} - v_{ib},\\
&z_{ib,1st} - z_{ib,2nd} \geq \alpha_i \Delta_i (P_{j,1st} - P_{j,2nd}) - v_{ib,1st} + v_{ib,2nd}, z_{ib,1st} \geq \alpha_i \Delta_i P_{j,1st} - v_{ib,1st}].
\end{aligned}
$$

The probability of keeping the second item only is:

$$
\begin{aligned}
&\mathbb{P}(keep\ 2nd)_{ib} = \mathbb{P}(V'(keep\ 2nd)_{ib} \geq \max(V'(keep)_{ib}, V'(keep\ 1st)_{ib}, V'(keep\ none)_{ib}))\\
=&\mathbb{P}(v_{ib,2nd} + z_{ib,2nd} - \alpha_i \delta_i P_{j,2nd}\\
\geq&\max(v_{ib,1st} + v_{ib,2nd} + v_{ib} + z_{ib,1st} + z_{ib,2nd} + z_{ib} - \alpha_i P_b, v_{ib,1st} + z_{ib,1st} - \alpha_i \Delta_i P_{j,1st}, 0))\\
=&\mathbb{P}(z_{ib,1st} + z_{ib} \leq \alpha_i (P_b - \Delta_i P_{j,2nd}) - v_{ib,1st} - v_{ib},\\
&z_{ib,2nd} - z_{ib,1st} \geq \alpha_i \Delta_i (P_{j,2nd} - P_{j,1st}) - v_{ib,2nd} + v_{ib,1st},\\
&z_{ib,2nd} \geq \alpha_i \Delta_i P_{j,2nd} - v_{ib,2nd})
\end{aligned}
$$

and the expected value of keeping the second item only:

$$
\begin{aligned}
&\mathbb{E}[V'(keep\, 2nd)_{ib} \mid keep\, 2nd] \\
=&v_{ib,2nd} - \alpha_i \Delta_i P_{j,2nd} + \mathbb{E}[(z_{ib,2nd} \mid z_{ib,1st} + z_{ib} \le \alpha_i (P_b - \Delta_i P_{j,2nd}) - v_{ib,1st} - v_{ib}, \\
&z_{ib,2nd} - z_{ib,1st} \ge \alpha_i \Delta_i (P_{j,2nd} - P_{j,1st}) - v_{ib,2nd} + v_{ib,1st}, z_{ib,2nd} \ge \alpha_i \Delta_i P_{j,2nd} - v_{ib,2nd}].
\end{aligned}
$$

Lastly,

$$
\mathbb{E}[V'(keep\, none)_{ib} \mid keep\, none]\mathbb{P}(keep\, none)_{ib} = 0
$$

since the expected value of returning both items is zero.

3.11 Appendix II: Insights from Simulations of the Model

Building upon the model developed above, we aim to gain insights into customer purchase and return behaviors through simulation exercises. Our simulations reveal that the post-purchase net surplus of a basket is a critical determinant in the decision to purchase and subsequently keep the products. Furthermore, when a retailer offers bundle discounts, sales tend to increase in scenarios where customers display higher levels of inattentiveness in the market. However, this increase in sales is often accompanied by a significantly higher number of returns, which suggests that the benefit of increased sales due to inattentive customer purchases may be offset by the corresponding rise in returns.

3.11.1 Expected Value of a Basket.

The expected value of a basket, $\mathbb{E}[V(Basket)]_{ib}$ (expressed in Eq. (3.5) and Eq. (3.6)) largely determines a customer's purchase decision. We analyze how it is influenced by factors such as a customer's initial valuation (v), subsequent valuation update (z), and their level of attentiveness (δ). For simplicity, our analysis is limited to two products: Item 1 and Item 2. We examine the dynamics when Item 1 is constantly in the basket, and we explore how the inclusion of Item 2 alters the basket's expected value. This exploration involves varying Item 2's valuation and valuation update, as well as the customer's level of attentiveness. For the purposes of the simulation, we assume both Item 1 and Item 2 have an original price of $p = 10$. When purchased together, they are available at a bundle discount price of 16. We fix a customer's initial valuation of Item 1 at $v_1 = 8$ and set her valuation update as a random variable drawn from a normal distribution $N(0,2)$.

In Figure 3.4, we illustrate how the expected value of the basket changes under three different conditions: (1) varying v_2, the valuation of Item 2, (2) varying the mean μ_2 of the valuation update z_2 for Item 2, where z_2 is drawn from $N(\mu_2, 2)$, and (3) varying δ, representing the customer's level of attentiveness to price details. This analysis aims to provide a nuanced understanding of how these variables interact and influence a customer's expected value of a basket, which determines her purchase decision.

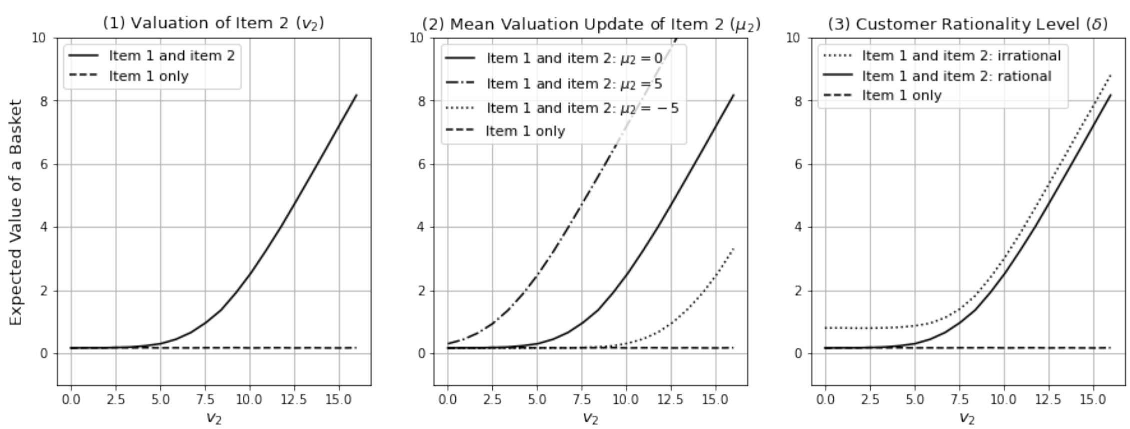

Figure 3.4: Expected Value of a Basket ($\mathbb{E}[V(Basket]_{ib}$) under Bundle Discount.

In Plot (1) of Figure 3.4, we set $z_2 \sim N(0,2)$ and $\delta = 1$, indicating that the valuation update for Item 2 has a mean of zero and customers are fully aware of price details. As illustrated in this plot, adding Item 2 to a basket would at least weakly increase its expected value, which is intuitive because customers always have the option to return without incurring any cost. As the valuation of Item 2 (v_2) increases, so does the expected value of the basket. In Plot (2), while keeping $\delta = 1$, we introduce scenarios with positive and negative mean valuation updates. When the valuation update has a positive mean, the basket's expected value is consistently larger compared to the scenario with a zero-mean update. Conversely, a negative mean valuation update results in a lower expected value of the basket. These observations affirm that the basket's expected value heightens with increasing product valuation, irrespective of whether this increase occurs at the purchase stage or the return stage.

Plot (3) of Figure 3.4 explores how different levels of attentiveness (δ) impact the expected value of a basket, under the condition that the mean valuation update remains zero. Here, we find that the expected value curve for an inattentive customer ($\delta = 0$) is consistently higher across all valuations of Item 2 than that for an attentive customer ($\delta = 1$). This suggests that an inattentive customer tends to ascribe higher value to the basket than an attentive customer, making her more inclined to proceed with the purchase.

These simulations on the expected values of baskets play a crucial role in determining a customer's likelihood of purchasing a basket or exiting the market without a purchase. In the subsequent section, we will conduct further simulations to demonstrate how these expected values translate into actual purchase and return behaviors.

3.11.2 Impact of Customer Inattentiveness on Purchases and Returns.

We analyze sales and return quantities under our model and compare the performance of bundle discount with single-item discount in terms of sales and net sales. As shown in Section 3.11.1 above, both a higher product valuation (v) and a higher return stage valuation update (z) lead to a higher expected value of a basket, which leads to a higher probability of purchases. As v and the realized z go linearly into a customer's return decision, it is also apparent that higher v and z will lead to less returns. However, the influence of customer attentiveness (δ) on purchase and return behavior is more nuanced and merits closer examination. We explore this aspect under varying levels of valuation update z.

We find that customer attentiveness is a key determinant in both purchase and return decisions. While customers may initially be drawn to two-item baskets under bundle discounts, they reevaluate and recognize the true value of their purchase at the return stage. Interestingly, despite higher sales volumes observed with less attentive customers, the net sales – accounting for returns – align closely with those involving more attentive customers. This suggests that the perceived benefit of increased sales from inattentiveness behavior is offset when considering the impact of returns, highlighting a critical factor that retailers need to consider in designing discount strategies.

In these simulations, we consider a market scenario with a single product priced at $p = 10$. The retailer introduces a bundle discount: purchasing two units of this product in one order costs 16, effectively a 20% discount for buying two. However, if a customer buys just one item, the price remains $p = 10$. Customers' valuations of the product are modeled as independent and identically distributed draws from a normal distribution, $v \sim N(8, 2)$. Additionally, their post-purchase valuation updates are drawn from another independent and identically distributed normal distribution, $z \sim N(\mu_z, 2)$. We allow μ_z to range from -50% to 50% of the mean of v to assess the impact of attentiveness under different scenarios of customer return stage valuation update.

At the purchase stage, customers have three options: to buy one item, to buy two items (thus benefiting from the bundle discount), or not to buy at all. Opting not to buy is assumed to yield a surplus of 0. During the return stage, customers re-evaluate their surplus for each item by factoring in their valuation updates, which can differ between the two items purchased. For comparative purposes, we also simulate a scenario where a single-item discount of 20% reduces the individual product price to 8. The valuation distributions and the prices are designed to mirror retail situations where pricing strategies aim to extract extra consumer surplus at the time of purchase. Consequently, in our simulations, positive post-purchase valuation updates lead to positive surpluses, while negative updates result in negative surpluses, allowing us to explore the effects of various consumer post-purchase valuation scenarios.

Our investigation is particularly focused on whether the presence of attentive customers ($\delta = 1$) versus inattentive customers ($\delta = 0$) leads to discernible differences in the effectiveness of bundle discount strategies. As the distribution of customer valuation updates is less predictable in reality, we are also keen to investigate how varying the mean valuation update (μ_z) alters customer decisions. It is noteworthy that in scenarios with single-item discounts, customer attentiveness level remains constant and does not affect purchase or return decisions. Thus, we use single-item discount as the baseline to examine the performance of bundle discounts in terms of both sales and net sales.

Figure 3.5 visualizes the results of our simulation analysis. The blue shaded areas in the first two columns represent the conditions under which a retailer experiences higher sales (first column) and net sales (second column) with bundle discounts compared to single-item discounts, suggesting that the bundle discount is the superior strategy in these regions. The third column illustrates return volumes, with the blue area indicating a higher number of returns resulting from the bundle discount than from the single-item discount. This column is instrumental in elucidating the shift from the blue/red regions in the sales column to the red/blue regions in the net sales column, offering insights into how return volumes influence the overall efficacy of the discount strategies.

Upon examining Figure 3.5, it is clear that a retailer's optimal strategy can differ markedly when focusing on net sales instead of just sales, as shown by the contrast be-

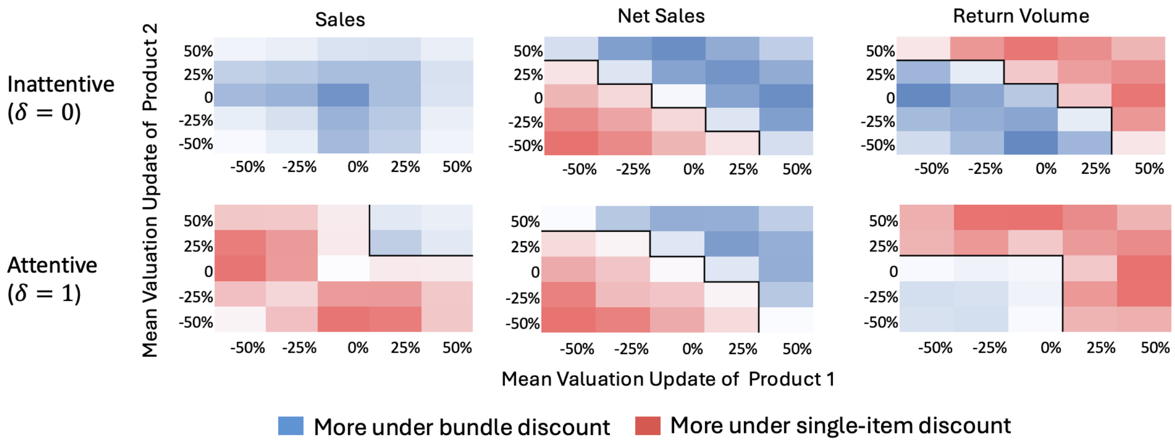

Figure 3.5: Bundle Discount vs. Single-item Discount. The charts displayed visualize the comparative performance of bundle discounts against single-item discounts, examining sales quantities, net sales, and return volumes. The axes of each plot correspond to the valuation updates for two items of the product. Scenarios with varied post-purchase valuation updates $z \sim N(\mu_z, 2)$ were simulated, where the mean update μ_z spans from -50% to +50% of the initial purchase stage valuation mean μ_v. Each scenario represents the behavior of 10,000 simulated customers as they navigate through purchase and return decisions, contingent on their attentiveness levels.

tween the first and second columns. Additionally, the comparison between the behaviors of inattentive and attentive customers indicates that customer base attentiveness significantly affects the retailer's choice of superior strategy. When considering sales alone, the bundle discount is favored for inattentive customers across all scenarios. For attentive customers, however, the bundle discount is preferable only in scenarios with positive mean valuation updates for both items.

When analyzing net sales, which take returns into account, the disparity in behavior between attentive and inattentive customers narrows. This pattern suggests that inattentive customers tend to return more products, especially when faced with negative valuation updates, ultimately resulting in net sales volumes that are similar to those of attentive customers. Yet, upon closer inspection, differences still persist between attentive and inattentive customers. For inattentive customers, a non-negative post-purchase surplus is sufficient for bundle discounts to outperform single-item discounts. In contrast, attentive customers demand a positive post-purchase surplus for the bundle discount to be considered the superior strategy.

The simulation exercises conducted in this section demonstrate two critical factors for the effectiveness of bundles as a discount strategy in terms of enhancing net sales. Firstly, constructing bundles that offer customers non-negative post-purchase net surpluses. Such surpluses make bundles more appealing and increase the likelihood of customers opting to keep them at the return stage. Secondly, by ensuring that customers understand the price tags of bundle discounts, retailers can significantly reduce the volume of returns.

The simulations presented in this section highlight two key elements that contribute to the success of bundle discounts as a strategy to improve net sales. First, the creation of bundle offers that result in non-negative post-purchase net surpluses for customers is essential. These positive surpluses enhance the appeal of bundles and raise the probability that customers will choose to retain their purchases rather than return them. Second, retailers can benefit from having more attentive customers. By providing clear and comprehensible pricing information for bundle discounts, retailers can substantially decrease product returns with little impact on sales.

In the upcoming section, we will return to the sales and return data from the Retailer to address a series of empirical questions. We aim to determine the sources of consumer post-purchase surpluses: Are they primarily derived from the initial purchase stage valuation, from the valuation updates at the return stage, or from the complementary values between bundled products? Additionally, we seek to ascertain the typical level of customer attentiveness to prices in a real-world fashion retail context. Understanding these aspects will provide valuable insights into customer behavior and help the design of discount strategies.

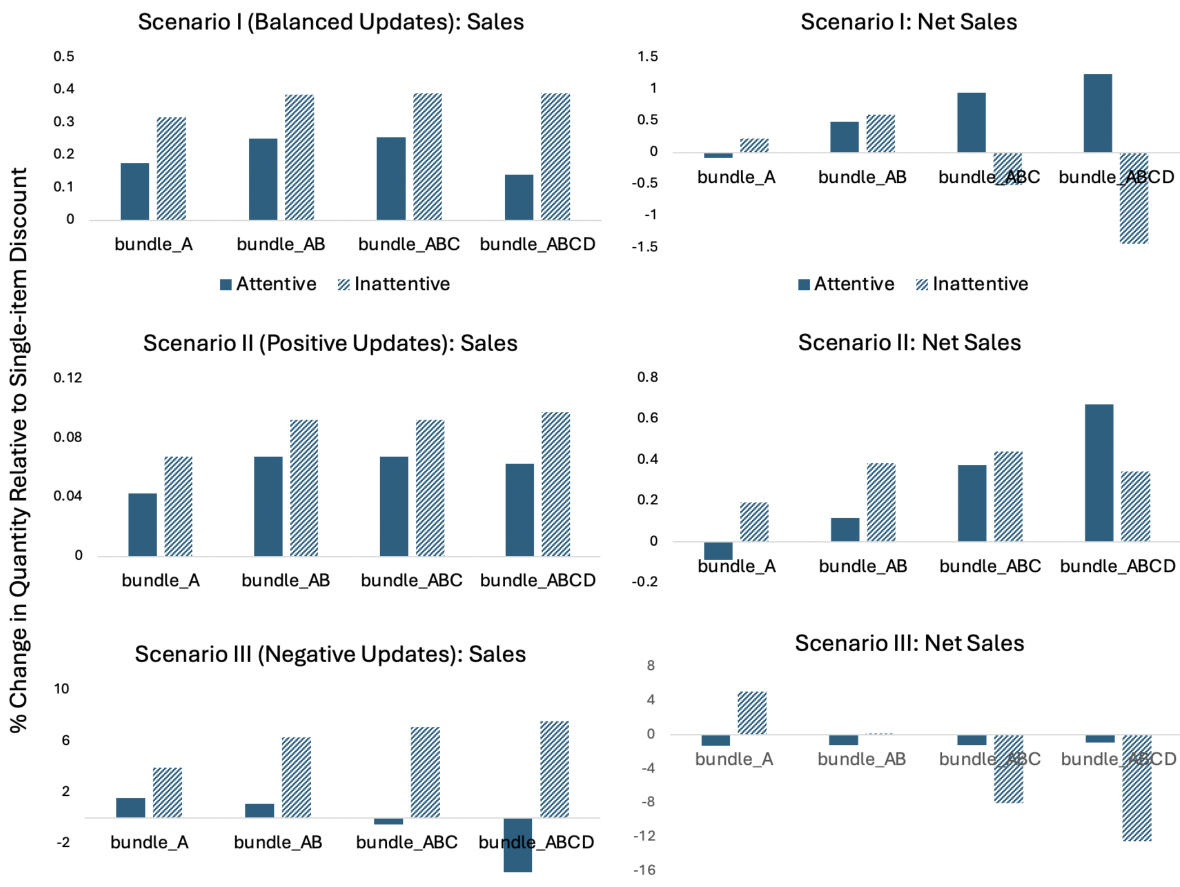

Figure 3.6: Quantities Sold under Bundle Discounts Compared to Single-item Discount. The numbers displayed are the percentage change in quantity of each bundle discount relative to that under single-item discount.

3.12 Appendix III: Simulations with Generic Categories

Our simulations were conducted across three scenarios as described in Table 3.8, each consisting of 10 trials with 10,000 simulated customers making purchasing and returning decisions. The standard price for each product category was set at \$100 with a corresponding discount price of \$80. The initial valuation at the purchase stage for each customer was independently drawn from a normal distribution $v \sim N(80, 20^2)$. The post-purchase valuation update was similarly drawn from a normal distribution $z \sim N(\mu_z, 10^2)$, with the mean μ_z varying according to the product category, as elaborated in Table 3.8. For clarity, no complementarity or substitutivity is considered. The parameter δ denotes the level of customer attentiveness to pricing, with $\delta = 1$ for fully attentive customers and $\delta = 0$ for those who are completely inattentive.

Figure 3.6 displays the results of simulations analyzing a retailer's sales and net sales under different discount strategies with attentive and inattentive customer bases. The optimal strategy in each plot is summarized in Figure 3.3.

Bibliography

[1] Alberto Abadie. Using synthetic controls: Feasibility, data requirements, and methodological aspects. *J. Econom. Literature*, 59(2):391–425, 2021.

[2] Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *J. Amer. Statist. Assoc.*, 105(490):493–505, 2010.

[3] Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from us labor markets. *Journal of political economy*, 128(6):2188–2244, 2020.

[4] William James Adams and Janet L Yellen. Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics*, 90(3):475–498, 1976.

[5] Ajay Agrawal, Joshua Gans, and Avi Goldfarb. The simple economics of machine intelligence. *Harvard Business Review*, 2016.

[6] Mehmet Sekip Altug and Tolga Aydinliyim. Counteracting strategic purchase deferrals: The impact of online retailers' return policy decisions. *Manufacturing & Service Operations Management*, 18(3):376–392, 2016.

[7] Muhammad Amjad, Vishal Misra, Devavrat Shah, and Dennis Shen. mrsc: Multidimensional robust synthetic control. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(2):1–27, 2019.

[8] Muhammad Amjad, Devavrat Shah, and Dennis Shen. Robust synthetic control. *J. Machine Learning Res.*, 19(1):802–852, 2018.

[9] Eric T Anderson, Karsten Hansen, and Duncan Simester. The option value of returns: Theory and empirical evidence. *Marketing Science*, 28(3):405–423, 2009.

[10] Volodymyr Babich and Gilles Hilary. OM forum—distributed ledgers and operations: What operations management researchers should know about blockchain technology. *Manufacturing & Service Oper. Management*, 22(2):223–240, 2020.

[11] Jie Bai. *Melons as lemons: Asymmetric information, consumer learning and seller reputation*. Center for International Development at Harvard University, 2021.

[12] Jie Bai, Maggie X Chen, Jin Liu, Xiaosheng Mu, and Daniel Yi Xu. Stand out from the millions: Market congestion and information friction on global e-commerce platforms. 2022.

[13] Mihai Banciu, Esther Gal-Or, and Prakash Mirchandani. Bundling strategies when products are vertically differentiated and capacities are limited. *Management Science*, 56(12):2207–2223, 2010.

[14] Heski Bar-Isaac, Steven Tadelis, et al. Seller reputation. *Foundations and Trends in Microeconomics*, 4(4):273–351, 2008.

[15] Marcus Bartley Johns, Paul Brenton, Massimiliano Cali, Mombert Hoppe, and Roberta Piermartini. The role of trade in ending poverty. *World Trade Organization*, 2015.

[16] Alexis Bateman and Leonardo Bonanni. What supply chain transparency really means. *Harvard Business Review*, 2019.

[17] Niloofar Bayat, Cody Morrin, Yuheng Wang, and Vishal Misra. Synthetic control, synthetic interventions, and covid-19 spread: Exploring the impact of lockdown measures and herd immunity. *arXiv preprint arXiv:2009.09987*, 2020.

[18] Ruth Beer, Hyun-Soo Ahn, and Stephen Leider. Can trustworthiness in a supply chain be signaled? *Management Sci.*, 64(9):3974–3994, 2018.

[19] Marco Bertini and Luc Wathieu. Research note—attention arousal through price partitioning. *Marketing Science*, 27(2):236–246, 2008.

[20] Bext360. Complete transparency, no greenwashing, 2023.

[21] Hemant K Bhargava. Mixed bundling of two independently valued goods. *Management Science*, 59(9):2170–2185, 2013.

[22] Philippe Blaettchen, Andre P Calmon, and Georgina Hall. Traceability technology adoption in supply chain networks. *arXiv preprint arXiv:2104.14818*, 2021.

[23] Wolfgang Bock, Dominic Field, Paul Zwillenberg, and Kristi Rogers. The growth of the global mobile internet economy. *Boston Consulting Group*, 2015.

[24] Gary E Bolton, David J Kusterer, and Johannes Mans. Inflated reputations: Uncertainty, leniency, and moral wiggle room in trader feedback systems. *Management Sci.*, 65(11):5371–5391, 2019.

[25] Leah Platt Boustan, Jiwon Choi, and David Clingingsmith. Automation after the assembly line: Computerized machine tools, employment and productivity in the united states. Technical report, National Bureau of Economic Research, 2022.

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

[27] Robert L Bray, Juan Camilo Serpa, and Ahmet Colak. Supply chain proximity and product quality. *Management Sci.*, 65(9):4079–4099, 2019.

[28] Ryan W Buell and Basak Kalkanci. How transparency into internal and external responsibility initiatives influences consumer choice. *Management Science*, 67(2):932–950, 2021.

[29] Ryan W Buell, Tami Kim, and Chia-Jung Tsay. Creating reciprocal value through operational transparency. *Management Science*, 63(6):1673–1695, 2017.

[30] Hongbin Cai, Ginger Z Jin, Chong Liu, and Li-An Zhou. More trusting, less trust? An investigation of early e-commerce in China. Technical report, National Bureau of Economic Research, 2013.

[31] Qingning Cao, Kathryn E Stecke, and Jun Zhang. The impact of limited supply on a firm's bundling strategy. *Production and Operations Management*, 24(12):1931–1944, 2015.

[32] Robert E Carter and David J Curry. Transparent pricing: theory, tests, and implications for marketing practice. *Journal of the Academy of Marketing Science*, 38:759–774, 2010.

[33] Christian Catalini and Joshua S Gans. Some simple economics of the blockchain. *Comm. ACM*, 63(7):80–90, 2020.

[34] Judith Chevalier and Austan Goolsbee. Are durable goods consumers forward-looking? evidence from college textbooks. *The Quarterly Journal of Economics*, 124(4):1853–1884, 2009.

[35] Andrew T Ching, Tülin Erdem, and Michael P Keane. Learning models: An assessment of progress, challenges, and new developments. *Marketing Sci.*, 32(6):913–938, 2013.

[36] Jiri Chod, Nikolaos Trichakis, Gerry Tsoukalas, Henry Aspegren, and Mark Weber. On the financing benefits of supply chain transparency and blockchain adoption. *Management Sci.*, 66(10):4378–4396, 2020.

[37] Bo Cowgill and Catherine E Tucker. Economics, fairness and algorithmic bias. *preparation for: Journal of Economic Perspectives*, 2019.

[38] Yao Cui and Vishal Gaur. Supply chain transparency using blockchain: Benefits, challenges, and examples. In *Global Logistics and Supply Chain Strategies for the 2020s: Vital Skills for the Next Generation*, pages 307–326. Springer, 2022.

[39] Yao Cui, Vishal Gaur, and Jingchen Liu. Supply chain transparency and blockchain design. *Available at SSRN 3626028*, 2020.

[40] Yao Cui, Ming Hu, and Jingchen Liu. Value and design of traceability-driven blockchains. *Available at SSRN 3291661*, 2019.

[41] Andrew M Davis, Vishal Gaur, and Dayoung Kim. Consumer learning from own experience and social information: An experimental study. *Management Sci.*, 67(5):2924–2943, 2021.

[42] Scott Davis, Michael Hagerty, and Eitan Gerstner. Return policies and the optimal level of "hassle". *Journal of Economics and Business*, 50(5):445–460, 1998.

[43] Prabuddha De, Yu Hu, and Mohammad S Rahman. Product-oriented web technologies and product returns: An exploratory study. *Information Systems Research*, 24(4):998–1010, 2013.

[44] Chrysanthos Dellarocas and Charles A Wood. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Sci.*, 54(3):460–476, 2008.

[45] Yiting Deng, Christopher S Tang, Wei Wang, and Onesun Steve Yoo. Can third-party sellers benefit from a platform's entry to the market? *Service Science*, 2023.

[46] Lingxiu Dong, Puping Jiang, and Fasheng Xu. Impact of traceability technology adoption in food supply chain networks. *Management Sci.*, 2022.

[47] David Dranove and Ginger Zhe Jin. Quality disclosure and certification: Theory and practice. *J. Econom. Literature*, 48(4):935–63, 2010.

[48] Amitava Dutta, Hau L Lee, and Seungjin Whang. RFID and operations management: Technology, value, and incentives. *Prod. Oper. Management*, 16(5):646–655, 2007.

[49] Daria Dzyabura, Siham El Kihal, John R Hauser, and Marat Ibragimov. Leveraging the power of images in managing product return rates. *Marketing Science*, 2023.

[50] Benjamin Enke and Florian Zimmermann. Correlation neglect in belief formation. *The Review of Economic Studies*, 86(1):313–332, 2019.

[51] Tülin Erdem and Michael P Keane. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing science*, 15(1):1–20, 1996.

[52] Tülin Erdem, Michael P Keane, and Baohong Sun. A dynamic model of brand choice when price and advertising signal product quality. *Marketing Sci.*, 27(6):1111–1125, 2008.

[53] Necati Ertekin and Anupam Agrawal. How does a return period policy change affect multichannel retailer profitability? *Manufacturing & Service Operations Management*, 23(1):210–229, 2021.

[54] Necati Ertekin, Jeffrey D Shulman, and Haipeng Chen. On the profitability of stacked discounts: Identifying revenue and cost effects of discount framing. *Marketing Science*, 38(2):317–342, 2019.

[55] Emmanuel Farhi, Josh Lerner, and Jean Tirole. Fear of rejection? Tiered certification and transparency. *RAND J. Econom.*, 44(4):610–631, 2013.

[56] Olumurejiwa Fatunde, Andre Calmon, Joann F de Zegher, and Gonzalo Romero. The value of long-term relationships when selling to informal retailers—evidence from India. *Rotman School of Management Working Paper*, (3792639), 2021.

[57] Apostolos Filippas, John J. Horton, and Joseph M. Golden. Reputation inflation. *Marketing Sci.*, 41(4):733–745, 2022.

[58] Jason Gale, Lydia Mulvany, and Monte Reel. How antibiotic-tainted seafood from China ends up on your table, 2016.

[59] Vishal Gaur and Abhinav Gaiha. Building a transparent supply chain. *Harvard Business Review*, 2020.

[60] Igal Hendel and Aviv Nevo. Intertemporal price discrimination in storable goods markets. *American Economic Review*, 103(7):2722–2751, 2013.

[61] Dorothée Honhon and Xiajun Amy Pan. Improving profits by bundling vertically differentiated products. *Production and Operations Management*, 26(8):1481–1497, 2017.

[62] Yan Huang, Param Vir Singh, and Kannan Srinivasan. Crowdsourcing new product ideas under consumer learning. *Management Sci.*, 60(9):2138–2159, 2014.

[63] Xiang Hui, Ginger Zhe Jin, and Meng Liu. Designing quality certificates: Insights from eBay. Technical report, National Bureau of Economic Research, 2022.

[64] Xiang Hui, Zekun Liu, and Weiqing Zhang. Mitigating the cold-start problem in reputation systems: Evidence from a field experiment. *Available at SSRN 3731169*, 2020.

[65] Xiang Hui, Maryam Saeedi, Giancarlo Spagnolo, and Steven Tadelis. Raising the bar: Certification thresholds and market outcomes. *Amer. Econom. J.: Microeconomics*, 2021.

[66] Kyle Hyndman, Santiago Kraiselburd, and Noel Watson. Aligning capacity decisions in supply chains when demand forecasts are private information: Theory and experiment. *Manufacturing & Service Oper. Management*, 15(1):102–117, 2013.

[67] IBM Food Trust. Ibm supply chain intelligence suite: Food trust, 2023.

[68] Karl Inderfurth, Abdolkarim Sadrieh, and Guido Voigt. The impact of information sharing on supply chain performance under asymmetric information. *Prod. Oper. Management*, 22(2):410–425, 2013.

[69] Danielle Inman. 17 most worrying textile waste statistics facts, 2022.

[70] Garud Iyengar, Fahad Saleh, Jay Sethuraman, and Wenjun Wang. Economics of permissioned blockchain adoption. *Management Sci.*, 2022.

[71] JD Chain. JD blockchain application white paper. Technical report, JD Technology, 2020.

[72] Zhaohui Jiang, Yan Huang, and Damian R Beil. The role of feedback in dynamic crowdsourcing contests: A structural empirical analysis. *Management Science*, 68(7):4858–4877, 2022.

[73] Ginger Zhe Jin, Zhentong Lu, Xiaolu Zhou, and Lu Fang. Flagship entry in online marketplaces. Technical report, National Bureau of Economic Research, 2021.

[74] Ginger Zhe Jin, Michael Luca, and Daniel Martin. Complex disclosure. *Management Science*, 68(5):3236–3261, 2022.

[75] Kyogo Kanazawa, Daiji Kawaguchi, Hitoshi Shigeoka, and Yasutora Watanabe. Ai, skill, and productivity: The case of taxi drivers. Technical report, National Bureau of Economic Research, 2022.

[76] N Bora Keskin, Chenghuai Li, and Jing-Sheng Jeannette Song. The blockchain newsvendor: Value of freshness transparency and smart contracts. *Available at SSRN 3915358*, 2021.

[77] Tobias J Klein, Christian Lambertz, and Konrad O Stahl. Market transparency, adverse selection, and moral hazard. *J. Political Econom.*, 124(6):1677–1713, 2016.

[78] Botond Kőszegi. Behavioral contract theory. *Journal of Economic Literature*, 52(4):1075–1118, 2014.

[79] Tim Kraft, León Valdés, and Yanchong Zheng. Consumer trust in social responsibility communications: The role of supply chain visibility. *Prod. Oper. Management*, 2022.

[80] Kailin Kroetz, Gloria M Luque, Jessica A Gephart, Sunny L Jardine, Patrick Lee, Katrina Chicojay Moore, Cassandra Cole, Andrew Steinkruger, and C Josh Donlan. Consequences of seafood mislabeling for marine populations and fisheries management. *Proc. Nat. Academy Sci.*, 117(48):30318–30323, 2020.

[81] Stephen Leahy. Revealed: seafood fraud happening on a vast global scale, 2021.

[82] Hau Lee and Özalp Özer. Unlocking the value of RFID. *Prod. Oper. Management*, 16(1):40–64, 2007.

[83] Hau L Lee and Seungjin Whang. Higher supply chain security with lower cost: Lessons from total quality management. *Internat. J. Prod. Econom.*, 96(3):289–300, 2005.

[84] Christopher G Leggett. Environmental valuation with imperfect information the case of the random utility model. *Environmental and Resource Economics*, 23(3):343–355, 2002.

[85] Retsef Levi, Somya Singhvi, and Yanchong Zheng. Economically motivated adulteration in farming supply chains. *Management Sci.*, 66(1):209–226, 2020.

[86] Jun Li, Nelson Granados, and Serguei Netessine. Are consumers strategic? structural estimation from the air-travel industry. *Management Science*, 60(9):2114–2137, 2014.

[87] Lingfang Li, Steven Tadelis, and Xiaolan Zhou. Buying reputation as a signal of quality: Evidence from an online marketplace. *RAND J. Econom.*, 51(4):965–988, 2020.

[88] Guilherme Liberali, Glen L Urban, and John R Hauser. Competitive information, trust, brand consideration and sales: Two field experiments. *Internat. J. Res. Marketing*, 30(2):101–113, 2013.

[89] Jane Alice Liu, Margaux Friocourt, and Yves-Marie Origlia. The problem with seafood supply chains, 2022.

[90] Steve Lohr. Economists pin more blame on tech for rising inequality, 2022.

[91] Steve Markenson and David Orgel. Transparency in an evolving omnichannel world, 2022.

[92] Dina Mayzlin, Yaniv Dover, and Judith Chevalier. Promotional reviews: An empirical investigation of online review manipulation. *Amer. Econom. Rev.*, 104(8):2421–55, 2014.

[93] R Preston McAfee, John McMillan, and Michael D Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, 104(2):371–383, 1989.

[94] Kevin F McCardle, Kumar Rajaram, and Christopher S Tang. Bundling retail products: Models and analysis. *European Journal of Operational Research*, 177(2):1197–1217, 2007.

[95] Jorge Mejia, Gloria Urrea, and Alfonso J Pedraza-Martinez. Operational transparency on crowdfunding platforms: Effect on donations for emergency response. *Production and Operations Management*, 28(7):1773–1791, 2019.

[96] Bhavya Mohan, Ryan W Buell, and Leslie K John. Lifting the veil: The benefits of cost transparency. *Marketing Science*, 39(6):1105–1121, 2020.

[97] Ken Moon, Prashant Loyalka, Patrick Bergemann, and Joshua Cohen. The hidden cost of worker turnover: Attributing product reliability to the turnover of factory workers. *Management Sci.*, 68(5):3755–3767, 2022.

[98] Leela Nageswaran, Soo-Haeng Cho, and Alan Scheller-Wolf. Consumer return policies in omnichannel operations. *Management Science*, 66(12):5558–5575, 2020.

[99] Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Available at SSRN 4375283*, 2023.

[100] Özalp Özer, Upender Subramanian, and Yu Wang. Information sharing, advice provision, or delegation: What leads to higher trust and trustworthiness? *Management Sci.*, 64(1):474–493, 2018.

[101] Özalp Özer and Yanchong Zheng. Establishing trust and trustworthiness for supply chain information sharing. In *Handbook of Information Exchange in Supply Chain Management*, pages 287–312. Springer, 2017.

[102] Özalp Özer, Yanchong Zheng, and Kay-Yut Chen. Trust in forecast information sharing. *Management Sci.*, 57(6):1111–1137, 2011.

[103] Özalp Özer, Yanchong Zheng, and Yufei Ren. Trust, trustworthiness, and information sharing in supply chains bridging China and the United States. *Management Sci.*, 60(10):2435–2460, 2014.

[104] Hubert Pun, Jayashankar M Swaminathan, and Pengwen Hou. Blockchain adoption for combating deceptive counterfeits. *Prod. Oper. Management*, 30(4):864–882, 2021.

[105] Arabella Ruiz. 17 most worrying textile waste statistics facts, 2023.

[106] Nachiketa Sahoo, Chrysanthos Dellarocas, and Shuba Srinivasan. The impact of online product reviews on product returns. *Information Systems Research*, 29(3):723–738, 2018.

[107] Jeffrey D Shulman, Anne T Coughlan, and R Canan Savaskan. Managing consumer returns in a competitive environment. *Management science*, 57(2):347–362, 2011.

[108] Eirini Spiliotopoulou, Karen Donohue, and Mustafa Çagri Gürbüz. Information reliability in supply chains: The case of multiple retailers. *Prod. Oper. Management*, 25(3):548–567, 2016.

[109] Stefan Stremersch and Gerard J Tellis. Strategic bundling of products and prices: A new synthesis for marketing. *Journal of marketing*, 66(1):55–72, 2002.

[110] Xuanming Su. Consumer returns policies and supply chain performance. *Manufacturing & Service Operations Management*, 11(4):595–612, 2009.

[111] Dmitrii Sumkin, Sameer Hasija, and Serguei Netessine. Does blockchain facilitate responsible sourcing? An application to the diamond supply chain. 2021.

[112] Tianshu Sun, Zhe Yuan, Chunxiao Li, Kaifu Zhang, and Jun Xu. The value of personal data in internet commerce: A high-stakes field experiment on data regulation policy. *Management Science*, 2023.

[113] Steven Tadelis. Reputation and feedback systems in online platform markets. *Ann. Rev. Econom.*, 8:321–340, 2016.

[114] Catherine Tucker. Algorithmic exclusion: The fragility of algorithms to sparse and missing data. 2023.

[115] R Venkatesh and Wagner Kamakura. Optimal bundling and pricing under a monopoly: Contrasting complements and substitutes from independently valued products. *The Journal of business*, 76(2):211–231, 2003.

[116] Seungjin Whang. Timing of RFID adoption in a supply chain. *Management Sci.*, 56(2):343–355, 2010.

[117] Chenguang Wu, Chen Jin, and Ying-Ju Chen. Managing customer search via bundling. *Manufacturing & Service Operations Management*, 24(4):1906–1925, 2022.

[118] Hao Ying, Xiaosong Peng, Xiande Zhao, and Zhong Chen. The effects of signaling blockchain-based track and trace on consumer purchases: Insights from a quasi-natural experiment. *Production and Operations Management*, 2023.

[119] Feng Zhu. Friends or foes? examining platform owners' entry into complementors' spaces. *Journal of Economics & Management Strategy*, 28(1):23–28, 2019.