The Dealer's Problem: Two-Sided Pricing and Inventory Management in Second-Hand Markets

Mingao Xie *
Georgetown University

September 25, 2026

Abstract

Resale platforms act as intermediaries on both sides of the market; they purchase used goods from individual sellers and sell them to buyers, so the buying and selling prices together shape inventory levels. Using transaction-level data from Déjà Vu, a large used-book platform, I estimate market sizes, demand, and supply, finding that demand is far more price-elastic than supply. I then solve an infinite-horizon dynamic pricing model with inventory as the state variable and compare it with the platform's static pricing. Dynamic pricing raises profits by 0.86% to 12.14%, with larger gains under demand and supply shocks, though it increases profit volatility and does not consistently reduce stockouts. Pricing strategies generated by an AI model underperform the platform's current prices.

JEL codes: L11, L81, C51, D22, D25
Keywords: two-sided platforms, dynamic pricing, inventory management

*Email: mx104@georgetown.edu

1 Introduction

Two-sided resale platforms face a complex pricing problem. Unlike marketplaces such as eBay and Facebook Marketplace, where the platform acts as an intermediary and earns a commission without setting prices itself, two-sided resale platforms purchase used products directly from individual sellers and resell them to individual buyers, exposing the platform to both demand and supply uncertainty. Such platforms also face an inventory management problem, because they lack an upstream manufacturer that can guarantee a relatively stable supply of products; their inventory tends to be low, and stockouts are common. Since both the sale price and the buyout price jointly determine inventory levels, two-sided pricing constitutes a particularly challenging problem.

Pricing under inventory constraints arises in many industries. Hotel rooms, concert tickets, and airline tickets, for example, naturally face inventory or capacity constraints. Hotel room prices are dynamic and typically higher for same-day bookings, while concert and airline ticket prices tend to rise as the event date approaches. However, most of this pricing is single-sided, since these industries cannot readily expand their inventory or capacity. Hotels cannot suddenly build additional rooms, concert venues cannot add extra seats, and airlines cannot flexibly add flights. While applied research has already examined single-sided dynamic pricing in the hotel (Cho and Rust 2010; Cho et al. 2018), concert, and airline ticket industries (Williams 2022; Betancourt et al. 2022; Siegert and Ulbricht 2020), research on two-sided pricing remains largely theoretical, particularly for goods without an expiration date or sale deadline. Little is known about how platforms implement two-sided pricing under inventory constraints.

I address this gap by modeling and estimating both sides of the market. On the supply side, users decide whether to sell their used books to the platform, while on the demand side, users decide whether to buy books and which quality to purchase. Using the estimates from the supply and demand models, I develop an infinite-horizon dynamic pricing model for the platform where inventory serves as the state variable. In this model, the platform sets the buyout price and the sale price simultaneously, accounting for the fact that these two prices jointly determine next period's inventory level. I also build a static pricing model that reflects the platform's current pricing approach, and compare the optimal dynamic pricing strategy against this static benchmark in terms of profits, transaction volume, and stockout rate.

Used books are an ideal product for this research for several reasons. Books are highly standardized, and unlike electronic devices whose components depreciate rapidly, books are far more durable and depreciate slowly. Moreover, demand and supply for books are persistent, in contrast to secondhand clothing, which exhibits seasonal cycles and is largely composed of one-of-a-kind items. I apply my model to Déjà Vu, one of the largest two-sided resale platforms, with more than 15 million users, providing secondhand circulation services for books. Déjà Vu provided all used book transaction data from July 1, 2023, to September 30, 2023. During this period, Déjà Vu employed a simple heuristic pricing strategy; based on the dataset and discussions with their algorithm team, prices were largely static and did not vary with inventory. The platform's team also expressed concern about this pricing challenge, noting, "We don't know if we can lower the buyout prices further, as they are already very low, yet users remain willing to sell their books," and observed that some books exhibit strong substitution between good and medium physical conditions, while others show weak substitution patterns. Understanding the demand and supply structure of this market is therefore essential for designing a profit-maximizing pricing strategy.

The paper proceeds as follows. Section 2 introduces the background of Déjà Vu, its selling and buying process, and the data used in this study. The data provide granular transaction-level detail, including each book's buyout and sale prices, sale and buyout timestamps, physical condition (high or medium quality), and daily inventory levels. Average daily inventory varies widely across books, with a median of less than one. If the platform were managing inventory efficiently, one would expect books with more transactions to carry higher inventory levels; this is not a general pattern observed in the data. The data also reveal a static pricing strategy, with buyout and sale prices remaining relatively stable over the three months despite substantial volatility in inventory. While there were several price adjustments, price drops or increases occurred simultaneously across all books, suggesting platform-wide sales events rather than book-specific dynamic pricing. Finally, I supplement the transaction data with the daily counts of users who clicked on a book and buyers who scanned it to estimate the market size for each book.

In Section 3, I estimate the daily buy-side and sell-side market size distributions, modeling them as a bivariate log-normal distribution that allows the two sides to be correlated. I focus on the five books with the highest transaction volume over the three months, as they provide more data points and greater price variation, allowing me to estimate parameters; they also serve as the primary objects of study throughout the paper. Results show that the sell-side market is four to ten times larger than the buy-side market and more volatile, and that the two are positively correlated; popular books attract both more buyers and more sellers. I then estimate a nested logit model on the sell side, allowing buyers to substitute across quality tiers but not across book titles, and a logit model on the buy side, where sellers face a binary sell/not-sell decision. Parameters are recovered via weighted maximum likelihood estimation, using market sizes as instruments. Results indicate a demand-side price elasticity of about 4, considerably more elastic than the supply-side elasticity of roughly 1. The nesting parameter varies substantially across books; some titles show values near 0, indicating limited substitution across qualities, while others exceed 0.5.

In Section 4, I build a dynamic model alongside a static model that most closely mirrors the platform's actual pricing strategy. Both models capture the platform's core tradeoff between holding costs and lost sales from stockouts. I recover the daily holding cost by assuming the observed static prices are optimal, then simulating inventory dynamics under this assumption to identify the range of holding costs consistent with the empirical inventory level being profit-maximizing; I take the median of this range as the holding cost used in the dynamic pricing model. In the dynamic model, the state variable is inventory across high- and medium-quality books, and given the distribution of buyers and sellers along with the demand- and supply-side estimates, the platform simultaneously sets the buyout and sale prices. Because the model has an infinite horizon and next-period inventory can swing from zero to very high depending on the realized draws of buyers and sellers, allowing prices to adjust with every inventory change creates computation challenge. To make computation tractable, I impose two simplifications, I restrict price adjustments to increments of 1, and I cap inventory at an upper bound of 100 for each book quality, with the platform halting acquisitions once that bound is reached.

Section 5 presents the counterfactual analysis. Using the dynamic pricing policy generated by the model, I compare its performance against static pricing through 200 simulations, each spanning 500 periods. The results show that dynamic pricing increases profits by 0.86% to 12.14%, with the magnitude varying across books, though it also raises the variance of profits. This improvement is either through higher transaction volume or through a wider spread between buyout and sale prices. However, dynamic pricing does not guarantee a lower stockout rate or more stable inventory levels.

1.1 Related Literature

This paper contributes to the growing literature in economics and operations research on pricing dynamics and inventory management. Dynamic pricing takes several forms in the literature. The first is intertemporal price discrimination, common in airlines (Betancourt et al. 2022; Williams 2022; Nair 2006; Siegert and Ulbricht 2020), hotels (Cho and Rust 2010), and event ticketing (Nair 2006), where prices vary with the time remaining before purchase. The second is inventory-driven pricing, in which price depends on remaining capacity; the same industries mentioned above often exhibit this form as well. A third form arises from real-time shifts in demand or supply, such as the surge pricing used by ride-hailing platforms like Uber and Lyft, where prices adjust instantaneously in response to demand spikes (Banerjee, Riquelme, and Johari 2015). Peak-load pricing, as seen in electricity markets (Schittekatte et al. 2024), represents a related form of dynamic pricing tied to fluctuating demand over predictable cycles. More recently, algorithmic dynamic pricing has emerged (Spann et al. 2024; MacKay and Weinstein 2022), in which prices are adjusted automatically based on competitor pricing or user browsing behavior, a practice widely used by online retailers and platforms. This paper focuses on dynamic pricing driven by inventory levels, extending the existing literature in an important way. Most prior work on capacity-constrained pricing, such as concert tickets, hotel rooms, and airline seats, assumes that capacity decreases over time and that there is a fixed deadline by which the inventory must be sold. In contrast, this paper considers a setting where the platform simultaneously buys and sells books, so inventory does not necessarily decline over time, and the problem is set over an infinite time horizon. Additionally, other papers typically focus on one-sided pricing, primarily the prices seen by consumers, whereas this paper studies two-sided pricing, in which the platform acts as both buyer and seller and must simultaneously determine both the sale price and the buyout price.

A substantial theoretical literature on two-sided markets has established the conceptual foundation for subsequent empirical research (Rochet and Tirole 2003; Armstrong 2005; Rochet and Tirole 2006; Hagiu 2009; Rysman 2009; Gao 2018).

This paper also contributes to the literature on firm conduct. Much of the existing work assumes, as a starting point, that firms behave optimally, but in practice, this may not hold. Cho and Rust 2010 document the "flat rental puzzle" in the car rental industry that, rather than offering discounted rates on older vehicles, rental companies tend to offer only relatively new cars at full price. They argue that this pattern represents a departure from optimal firm behavior, because actual rental markets are not fully competitive and firms may not face enough pressure to adopt the profit-maximizing policy. Their model implies that firms should hold onto cars roughly twice as long as they currently do, discounting older vehicles to attract renters, and their experimental results confirm that this strategy increases profits. Romer 2006 finds another puzzle that NFL teams are too conservative and uses NFL fourth-down decisions as a test of whether organizations behave like profit maximizers, and it finds large, systematic departures from the choice that would maximize winning chances. One potential reason for this suboptimal behavior is that managers or coaches rely on intuition or rules of thumb rather than formal optimization. In my paper, each used book has limited inventory, unlike consumer packaged goods in grocery stores, where stock is large and can be easily replenished. Intuitively, when inventory is low, for example below 10 units, we would expect both the sale price and buyout price of the book to rise. However, the static pricing observed in the dataset contradicts this intuition. I therefore build a counterfactual dynamic pricing model to show that the firm may not be pricing optimally.

Finally, this paper also contributes to the growing literature on second-hand markets. Existing research in this area has largely focused on industries such as automobiles, electronics, and clothing, with studies centering on adverse selection (Akerlof et al. 1970; Emons and Sheldon 2002; Peterson and Schneider 2014), information asymmetry (Blundell et al. 2026; Lewis 2011), and trade patterns (Porter and Sattler 1999). In this paper, however, adverse selection is largely absent; used books carry relatively low unit value, and the platform employs a mature, standardized condition-assessment system that, according to discussions with its product team, effectively reduces uncertainty about book quality. Rather than studying strategic interaction between users and the platform, this paper instead focuses on the two-sided pricing problem the platform must solve.

2 Data and Background

2.1 Used Book Market

Books are considered durable goods, but they differ from other durable goods like automobiles or electronic devices, which are used more frequently and whose depreciation is mainly due to physical wear. Once a book is read, its informational or entertainment value is largely consumed by the reader (though it may be reread or referenced later), and the physical copy is typically retained rather than discarded. Because the content of a used book is identical to that of a new copy, many buyers choose second-hand books to save money. The used book market has expanded significantly worldwide (Research and Markets 2024), including both traditional physical used bookstores and online marketplaces such as Amazon, ThriftBooks, and eBay.

Physical used bookstores generally operate within a smaller business radius than online marketplaces, as their customer base is largely confined to the area surrounding the store's physical location. Most physical used bookstores are independently operated rather than franchised, although exceptions exist, such as Book Off in Japan, 2nd & Charles in the U.S, and Aladin in Korea. By contrast, online marketplaces are more organized and operate across a much wider market, enabling booksellers and buyers to trade nationwide (Ellison and Ellison 2018). Online platforms also make it easier to adjust prices, maintain records, and track daily activity, such as the number of users.

2.2 Data

2.2.1 Data Sources

The platform Déjà Vu provides the data used in this research. Déjà Vu is a platform that mainly offers secondhand circulation services for books and durable consumer goods. Déjà Vu adopts a C2B2C model, featuring a self-operated recycling system and standardized pricing mechanisms. It is one of the largest secondhand book companies in China, with more than fifteen million users. Déjà Vu primarily sells online, although it also has some brick-and-mortar stores. However, in 2019, the year of the data used in this research, all products were sold online.

The dataset consists of over two million book transaction records from 2023-07-01 to 2023-09-30. Each transaction contains the following information: book ID [1], author ID, publisher ID, publication date, book rating (sourced from a professional book-rating platform analogous to "Rotten Tomatoes"), book category, book quality, cover price, selling price, sale timestamp, purchase price, acquisition timestamp, buyer ID, seller ID, and the book-quality specific daily inventory. In addition, the platform provides daily engagement metrics for each book title, including the number of user clicks, cart additions, and seller scans, which can be used to estimate daily market demand and supply.

2.2.2 Selling and Buying Process

On the sell side, where users purchase books from the platform, the experience follows a standard online marketplace model. Users can browse featured books displayed on the homepage and click on any title to view detailed information about the book. They can also search by book title to check whether a book is available on Déjà Vu. Note that the same title may exist in multiple editions published by different publishers. For more precise searches, users can also search using a book's ISBN.

On the buy side, Déjà Vu operates a well-established buying system. The store only acquires books it expects to resell at a reasonable price. To determine the acquisition price, users can either scan the barcode on the back of the book, as shown in Figure 1, or manually enter the ISBN into the app. The app will then display an acquisition price similar to the one shown in Figure 2. The final acquisition price depends on the book's physical condition, which is either high quality or medium quality. It is entirely the user's decision whether to place the order. If the user sends in a defective book, they will not be paid and must either

1. Because book titles are hashed, it is not possible to identify whether two books are close substitutes for one another. Given this constraint, this paper assumes that substitution occurs across quality conditions of the same book ID, but not across different book IDs.

cover the return shipping fee to get the book back or allow the platform to donate it to a library or recycle it. For this reason, users should perform a quick self-assessment of the book's physical condition before sending it. In addition, because the platform cannot determine whether a book is high- or medium-quality until it has been evaluated in person, the user will only see a price range when scanning the book, with the high-quality price as the upper bound and the medium-quality price as the lower bound.

Figure 1: ISBN barcode of the book *The Sympathizer* by Viet Thanh Nguyen

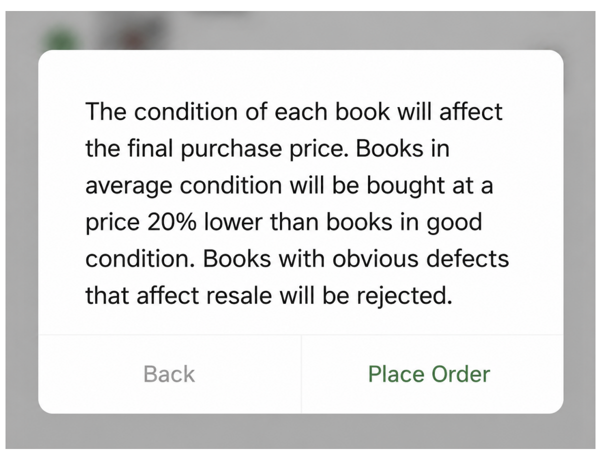

Figure 2: Post-scan Buy-side Pricing Terms

Notes: The photo was taken in May 2026, whereas the dataset used in this study covers July–September 2023. Therefore, the policy stating that the acquisition price for medium-quality units is 20% lower than that for high-quality units does not apply to the data used in this study.

2.2.3 Holding Time

The platform claims that 80 percent of the books can be sold within 8 hours after being listed on the app. Note that the mean and median holding times for a book are 957 hours (~40 days) and 362 hours (~15 days), respectively, both substantially longer than 8 hours. This average holding time is calculated as the duration from when the platform purchases the book until it is sold to another user. The holding time includes the time spent shipping the book to the platform's warehouse from the seller's address, as well as the time for inspection, cleaning, and logging it into the system, along with the preparation for sale and the actual selling process. Detailed summary statistics on sales, prices, and holding times by book category and quality are reported in the Appendix. Holding times vary considerably across categories: Social Sciences, Philosophy and Religion, and Design tend to have shorter holding times, while daily transaction volumes range widely, from over 2,000 for Children's Books to fewer than 200 for Design. If we accept the assumption that different book categories have approximately the same shipping and cleaning times, the holding duration can reflect the selling speed of books.

2.2.4 Inventory

The novelty of the data is that it contains the daily inventory of each book title. According to an interview with the platform's data analyst, the company's warehouses have never reached storage capacity, implying that physical space is not a binding constraint on inventory decisions. However, holding inventory is not costless. Expenses include warehouse rent, shelving, utilities, and climate control, as well as the labor involved in sorting, handling, and preserving books in good condition. We might expect thicker books to have a higher daily holding cost compared with thinner books, but in the data, I do not have information on the number of pages of each book. Also, books are standardized products and do not vary too much in size (it's rare to see a ten-page book and a one-thousand-page book), so it is safe to assume holding costs do not vary too much in absolute terms. The average daily inventory of books is very dispersed, with the median being 1 for high-quality books and 0.5 for medium-quality books. It is noteworthy that 25% of the medium quality books are out of stock every day, and 10% of the high quality books are out of stock and have zero average daily inventory. The out-of-stock could be strategic, resulting from the platform not being interested in selling unprofitable books, or it could be a shortage of books and a flaw in inventory management. And inventory is positively correlated with the popularity of the books. If we focus on the top 100 books that sold the most, they have a much higher daily inventory, with a median of 13 and 11 for high and medium quality, respectively. If the platform is doing optimal inventory management, we would expect the inventory of a book to be correlated with the sell rate, buy rate, sell price, and buy price. However, some popular books have suspiciously high daily inventory that exceeds 160, even though their prices and sell/buy rates are not that different from other popular books that have a daily inventory of 50.

2.2.5 Sell Rate and Buy Rate

Unlike markets for used cars or other high-priced goods, where consumers may prefer medium- or low-priced options, used book prices are relatively compressed and inexpensive. Hence, consumers are unlikely to purchase a book solely because it is cheap. Moreover, a book's price is not tied to its literary quality. For instance, we do not expect Pulitzer Prize-winning books to be priced higher than books that have not won any awards. Table 1 reports the quantiles of sell and buy rates by quality for both the overall population and the popular books. By construction, popular books exhibit higher buy and sell rates. While the sale prices and buyout prices of the 100 most popular books are not higher than those of the overall population, their daily inventories are substantially larger. Notably, in the overall population, the sell and buy rates at the 70th percentile are as low as 1, suggesting that a large share of books are long-tail items. Since these long-tail books record only a single transaction over the sample period, I restrict my analysis to the popular books, whose greater variation in prices and quantities enables the identification of demand and supply.

Table 1: Summary Statistics of Inventory, Prices, and Sell/Buy Rate

All Books

Quantile	sell rate	buy rate	inv h	inv m	sale price	buyout price
Q10	1.00	1.00	0	0	8.72	1.61
Q30	1.00	1.00	0	0	12.20	3.85
Q50	1.00	1.00	1	0	16.79	6.50
Q70	1.00	1.00	1	1	24.05	10.30
Q90	1.12	1.13	3	2	46.00	20.33

Top 100 Popular Books

Quantile	sell rate	buy rate	inv h	inv m	sale price	buyout price
Q10	4.74	3.92	5	2	10.93	1.78
Q30	5.27	4.66	7	4	13.08	4.31
Q50	6.16	5.16	11	13	16.06	6.87
Q70	7.44	6.08	36	58	20.05	9.05
Q90	10.21	7.98	116	157	26.95	13.35

Notes: inv h is inventory for high-quality books; inv m is inventory for medium-quality books.

2.2.6 Platform Pricing Strategy

This platform originated as a small-scale, offline book-collection group, and the very initial inventory and prices were recorded in Excel. As the number of buyers and sellers grew, the founder established the current platform, which operates as a two-sided marketplace for used books. Its pricing strategy has undergone several changes over time. In the early stages, the platform used a very simple rule, buying books at 20% of the cover price and reselling them at 50%, and did not differentiate between book qualities. As the platform matured, books were categorized into high- and medium-quality tiers based on physical condition (for example, whether the interior pages contained writing or markings, and whether the dust jacket was intact), with different prices applied to each tier. Even then, however, the pricing rule remained relatively simple; medium-quality books were priced two dollars below high-quality ones, both for the sale price and, analogously, for the buyout price. Such heuristic rules were easy to implement, and the price gaps were set based on rule-of-thumb judgments. The current pricing strategy is considerably more sophisticated, combining machine learning with heuristic adjustments.

2.2.7 Frequency of Price Adjustments

Unlike bookstores, which purchase books directly from publishers and therefore face considerable certainty over order quantities and receive a fixed wholesale discount, the platform cannot directly control how many books it acquires. It does, however, retain full control over buy and sell prices and can influence buy and sell volumes through price adjustments. Such adjustments can also be used to optimize profits. Dynamic pricing is standard practice in industries such as airlines and concert ticketing, where the scarcity of seats makes intertemporal price variation essential. Books differ in that inventory can be replenished over time, though not instantaneously, owing to acquisition delays. However, the possibility of stockouts still provides a rationale for dynamic pricing. Based on several meetings with the platform's algorithm team, prices are reportedly determined primarily by a machine learning model. I do not have access to the details of this algorithm, which remains a black box. The machine learning model takes historical data as input and returns the buy and sell prices for each book. Basic demand and supply reasoning would suggest that selling prices should rise when demand is high, buying prices should fall when supply is abundant, and prices should be adjusted accordingly whenever feasible. In the data, however, price adjustments are infrequent, and prices typically remain stable for several months at a time. Occasional platform-wide price increases do occur, applying uniformly to all books. Figure 3 and Figure 4 plot the trends of sale price, buyout prices, and inventory.

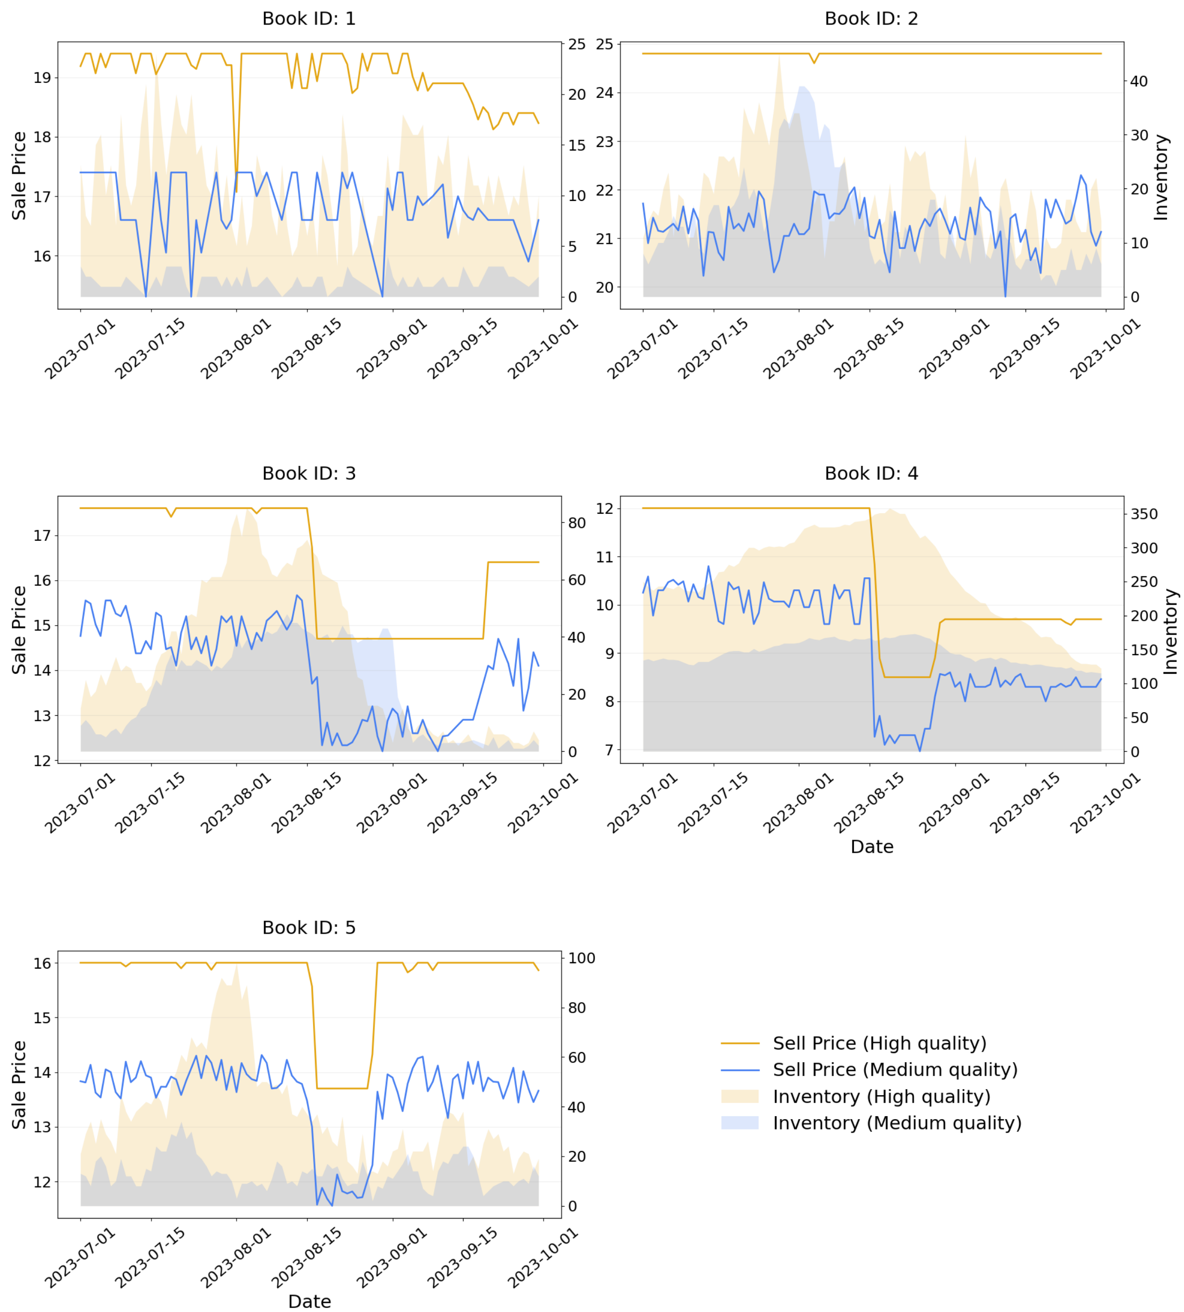

Figure 3: Sale Prices and Inventory Trends by Book and Quality

Notes: Inventory and pricing data are recorded at the daily level. The sale price for medium-quality books represents a weighted average, since the medium-quality category includes two sub-conditions.

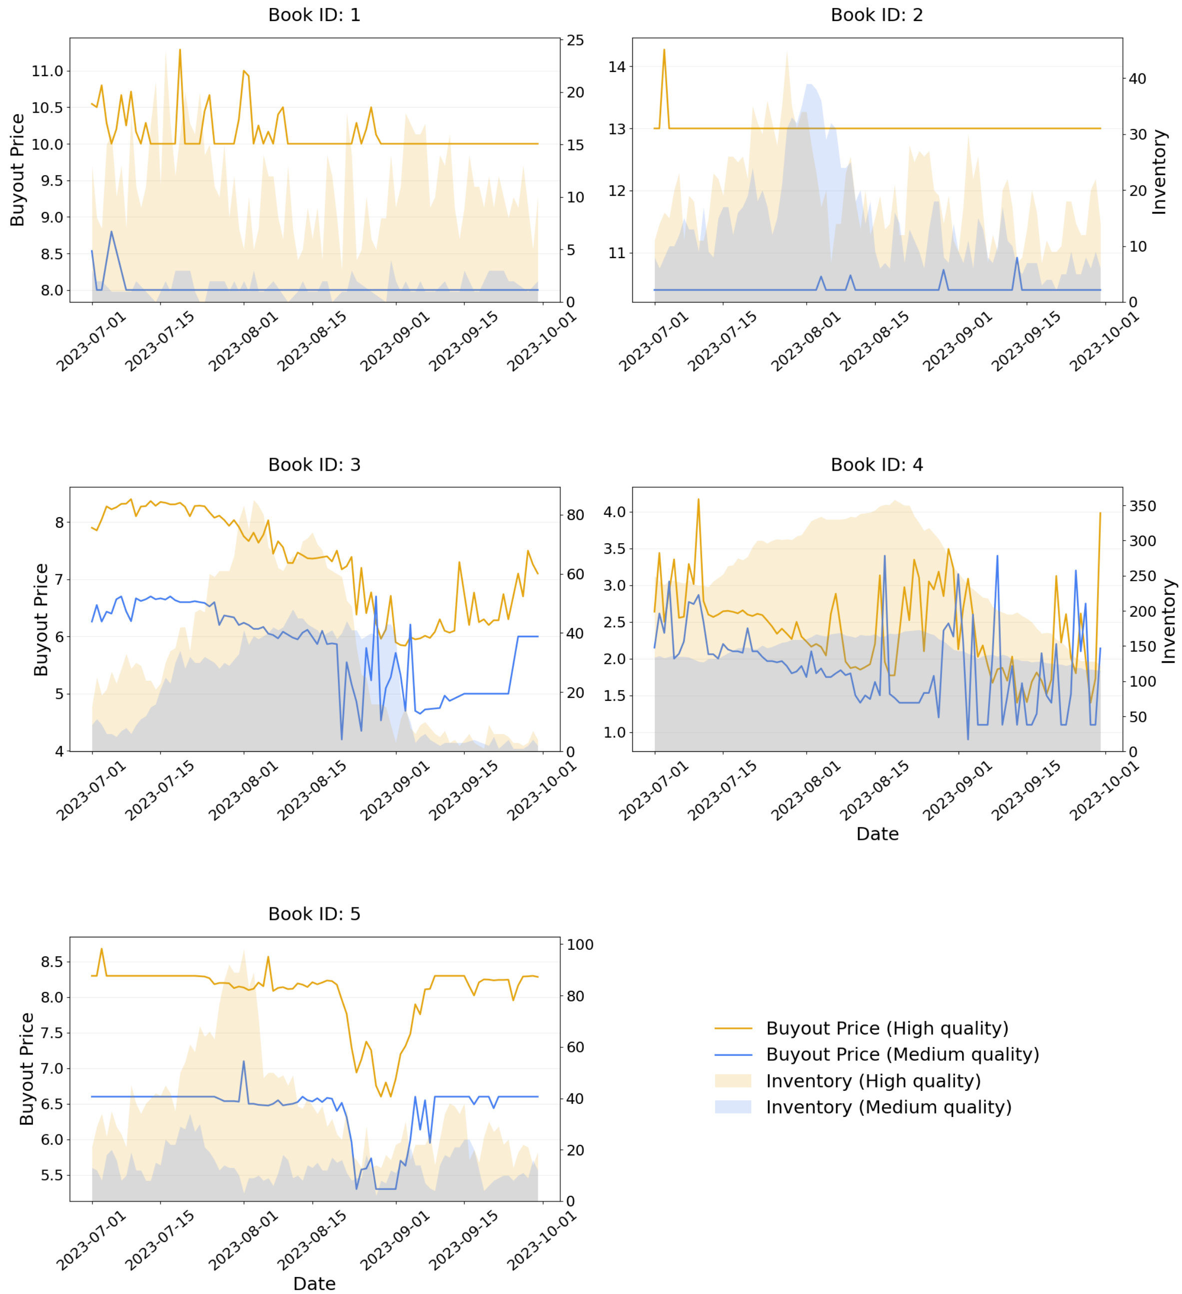

Figure 4: Buyout Prices and Inventory Trends by Book and Quality

Notes: Inventory and pricing data are recorded at the daily level, with inventory reflecting the stock level on the day the book was acquired.

The prices of high-quality books are highly stable, with only two adjustments observed over the three months. The timing of these adjustments coincides across books, suggesting they are platform-wide rather than book-specific. At first glance, medium-quality books appear to exhibit frequent small daily price fluctuations. However, this is misleading. Medium quality encompasses two sub-conditions (one with slight wear and another with internal markings), and their prices differ slightly. The dataset does not distinguish between these two sub-conditions, labeling both simply as medium quality. Upon closer examination of the transaction records, the same set of prices recurs daily, indicating that the platform is not actively adjusting prices for either sub-condition. The apparent fluctuations in the plot are an artifact of averaging across the two sub-conditions at different price points. Moreover, some books exhibit substantial daily inventory fluctuations while their prices remain stable, suggesting that inventory levels may not influence the platform's pricing decisions. A direct comparison of buyout and sell prices is provided in the Appendix.

There are several reasons why the platform maintains relatively static pricing. Frequent price adjustments require ongoing engineering oversight and code maintenance, both of which incur significant costs. Additionally, the platform may not yet have developed a mature dynamic pricing strategy. The algorithm team also noted that overly frequent price changes could lead users to perceive price discrimination, ultimately degrading the user experience.

3 Estimation

3.1 Market Size

The daily sell-side market size for a given book is defined as the maximum of the number of users who clicked on the book and the number of users who added it to their cart. Based on discussions with the company's data scientists, most buyers follow a sequence of clicking the book, viewing the product page, and then adding it to their cart. However, a smaller subset of users adds the book to their cart directly from the app's front page without first clicking on it. Since these two user types cannot be distinguished in the dataset, and the latter group represents only a small proportion, I use the maximum of clicks and add-to-cart actions as the daily sell-side market size. On the buy side, users who wish to sell a book must scan its barcode to receive a quoted price; accordingly, I use the number of scans of a given book as its daily buy-side market size.

No weekly pattern is observed in market size; weekday market size does not differ noticeably from weekend market size. The data spans three months, during which there was at least one major sales event where I observed a platform-wide reduction in sale prices. Overall, the market size remained relatively stable throughout the three months, as shown in Figure 5.

It is worth noting, however, that the market sizes on the two sides may be correlated. For a popular book, high demand (i.e., a large sell-side market size) may coincide with owners being less willing to part with the book, producing a negative correlation between the buy-side

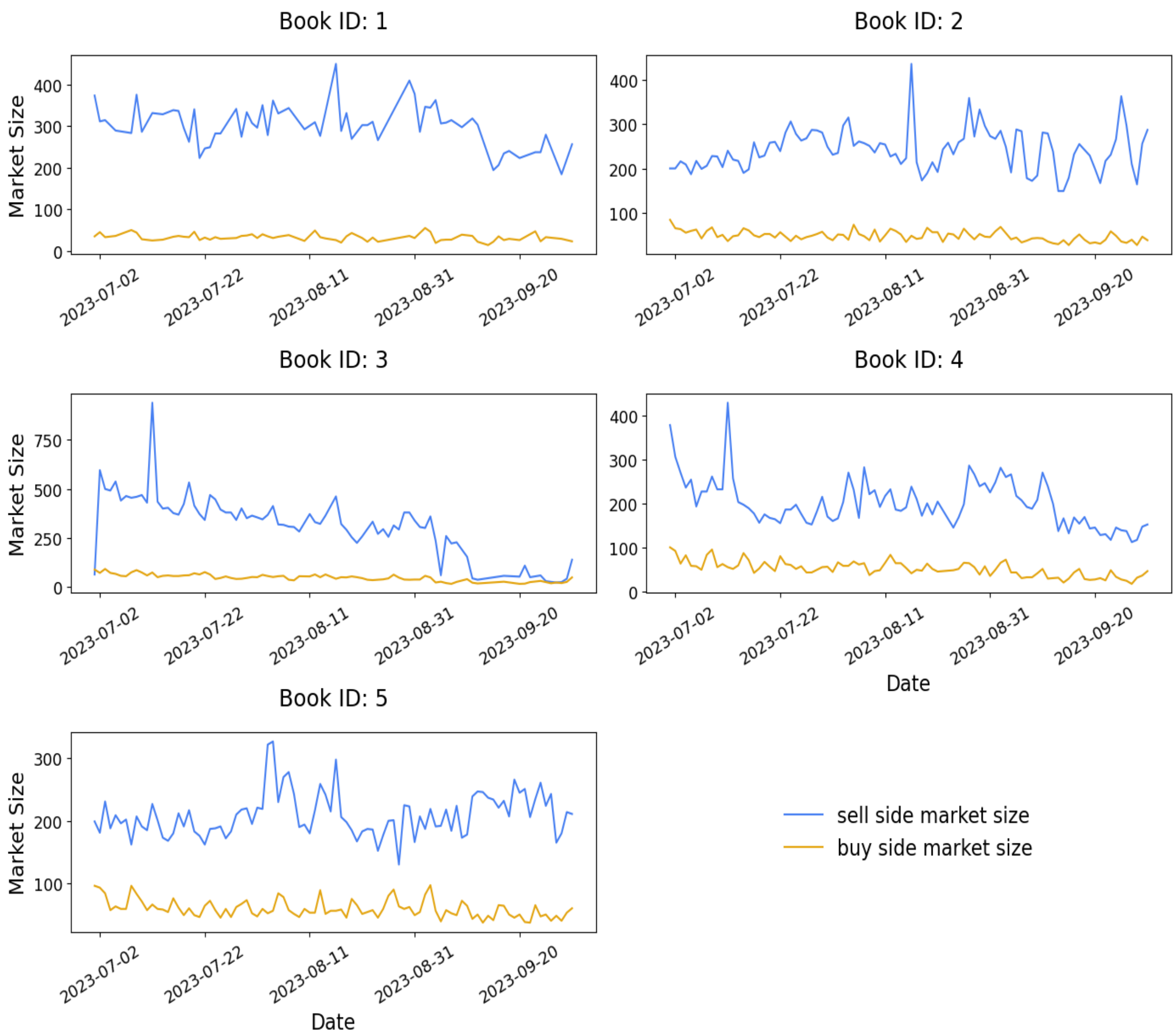

Figure 5: Sell-Side and Buy-Side Market Size Trends

Notes: Sell-side market size is the number of clicks that day from buyers on that book, while buy-side market size is the number of sellers who scanned that book that day

and sell-side market sizes. On the other hand, a book gaining popularity may prompt owners who were previously unaware that the book could be sold to the platform to scan it, thereby inducing a positive correlation between the two market sizes. As shown in Figure 6, the market sizes of the five selected book titles are right-skewed. Combined with the fact that market sizes cannot be negative, this motivates modeling the market sizes on both sides as jointly following a bivariate log-normal distribution. Let S denote the sell-side market size, B denote the buy-side market size, and ρ represent the correlation; then $(\ln S, \ln B) \sim \mathcal{N}_2(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ where

$$\boldsymbol{\mu} = \begin{pmatrix} \mu_S \\ \mu_B \end{pmatrix}, \qquad \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_S^2 & \rho\sigma_S\sigma_B \\ \rho\sigma_S\sigma_B & \sigma_B^2 \end{pmatrix}$$

Table 2 presents the maximum likelihood estimation results for the five books examined in this study. The bivariate log-normal distribution provides a good fit for four of the five books; the exception is one book whose extreme buy-side market size is not well captured by the log-normal distribution. The mean sell-side market size (E[S]) ranges from approximately 203 to 355, whereas the mean buy-side market size (E[B]) is consistently smaller, ranging from approximately 33 to 59. On average, the seller market size is roughly four to ten times larger than the buyer market size across the five books. None of the books exhibit a negative correlation, and three titles display a clear positive correlation between buyer and seller market sizes. Market size volatility varies substantially across the five books, and sell-side volatility is generally far larger than the buy-side volatility. The standard deviation of the sell-side market size ranges from roughly 34 to 51 for four of the five books but spikes to 365 for Book 3. The buy-side market size is more stable, with a standard deviation ranging from 8.6 to 21.9. I also fitted a bivariate normal distribution, but it did not outperform the bivariate log-normal specification.

Table 2: Bivariate log-normal Estimates of Two-sided Market Sizes

Book ID	$\hat{\mu}_{\log s}$	$\hat{\mu}_{\log b}$	$\hat{\sigma}_{\log s}$	$\hat{\sigma}_{\log b}$	$\hat{\rho}_{\log}$
1	5.69	3.45	0.17	0.26	0.28
					[0.04, 0.49]
2	5.47	3.83	0.19	0.24	0.07
					[−0.13, 0.27]
3	5.51	3.79	0.85	0.43	0.72
					[0.59, 0.81]
4	5.28	3.90	0.25	0.35	0.55
					[0.39, 0.68]
5	5.33	4.05	0.16	0.23	−0.11
					[−0.30, 0.10]

Notes: Subscripts b and s denote buy side and sell side. Sample sizes range from $n = 64$ to $n = 92$. Values in brackets are 95% confidence intervals for $\hat{\rho}$ obtained via Fisher's z-transformation: $z = \tanh^{-1}(\hat{\rho})$ with standard error $1/\sqrt{n-3}$, and endpoints back-transformed via $\tanh(\cdot)$.

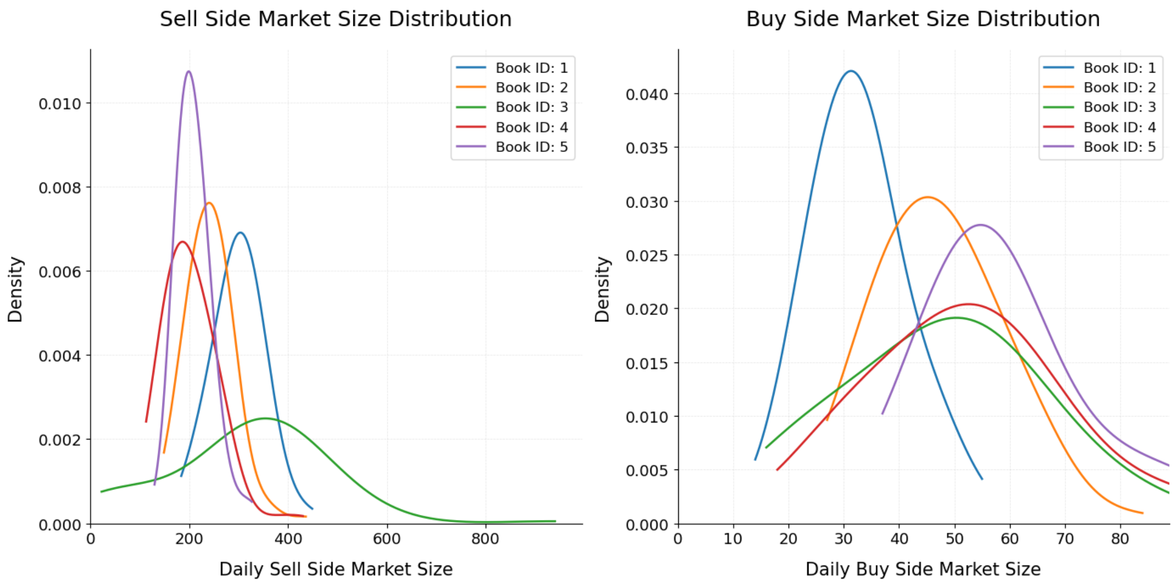

Figure 6: Two-sided Market Sizes Distribution

Notes: Sell-side market size is the number of clicks that day from buyers on that book, while buy-side market size is the number of sellers who scanned that book that day. For visual clarity, the horizontal axis of each marginal density plot is truncated at $x_{\max} = \max_i F_i^{-1}(0.95)$, where F_i^{-1} denotes the quantile function of the lognormal marginal for book i. This removes the extreme right tail driven by the book with the largest σ, while preserving the bulk of every distribution.

3.2 Users

3.2.1 Book Buyers

I assume a nested logit on the sell side. The utility that consumer i derives from choosing product j at time t is given by

$$u_{ijt} = \delta_{jt} + \zeta_{igt} + (1-\sigma)\epsilon_{ijt} \tag{1}$$

where $\delta_{jt} = x_j'\beta + \alpha p_{jt} + \xi_j$. Here ϵ_{ijt} is assumed to be independently and identically distributed extreme value, and x_j denotes the quality fixed effect. σ is the nesting parameter. With $\sigma = 0$, the model collapses to a standard logit.

Conditional on buying, the probability that a consumer chooses product j is:

$$\pi_{j|g=1} = \frac{exp(\frac{\delta_j}{1-\sigma})}{D_g} \tag{2}$$

where

$$D_g = \sum_{k=1,2} exp(\frac{\delta_k}{1-\sigma}) \tag{3}$$

and k=1 denotes high quality while k=2 denotes medium quality.

The unconditional probability that group g is selected is given by

$$\pi_g = \frac{D_g^{(1-\sigma)}}{\sum\limits_g D_g^{(1-\sigma)}} \tag{4}$$

Combining these two expressions, the unconditional purchasing probability of SKU j can be written as

$$\pi_j = \pi_{j|g}\pi_g = \frac{exp(\frac{\delta_j}{1-\sigma})}{D_g}\frac{D_g^{(1-\sigma)}}{\sum\limits_g D_g^{(1-\sigma)}} = \frac{exp(\frac{\delta_j}{1-\sigma})}{D_g^{\sigma}\sum\limits_g D_g^{(1-\sigma)}} \tag{5}$$

The data consist of individual-level transactions. I assume that all books share the same price coefficients for buying and selling, as well as the same quality fixed effects on both sides of the market, while allowing each book to have its own book-specific fixed effect and nesting parameter. This modeling choice aligns with the platform's data scientist's observation that, even when the price gap between high-quality and medium-quality copies is the same across titles, some books show stronger substitution across book qualities than other book titles. Based on sales volume, I select the five most popular book IDs (B=5 in the equation) and estimate the parameters using weighted maximum likelihood

$$\ln \mathcal{L}(\theta) = \sum_{b=1}^{B} w_b \sum_{t=1}^{T}\sum_{i=1}^{N_{bt}}\left(\sum_{j\in\{h,m\}} d_{ibtj}\ln \pi_{ibtj} + (1 - d_{ibth} - d_{ibtm})\ln \pi_{ibt0}\right) \tag{6}$$

where $w_b = \frac{1}{N_b}$, and $N_b = \sum_{t=1}^{T} N_{bt}$ is the total number of transactions for book b. The indicator d_{ibtj} equals 1 if the buyer chooses quality j and 0 otherwise (i.e., either not buying or buying a different quality). Because maximum likelihood estimation naturally places greater weight on books with more transactions, more popular books receive stronger influence in the estimation. To ensure that each book title contributes equally to the estimation, I reweight the observations so that every book title carries the same overall weight.

Prices are endogenous. Using the buy price as an instrument for the sell price seems intuitive because the two prices are highly correlated. However, unobserved demand shocks on the sell side may also influence the buy price, which violates the exogeneity condition. Instead, I use market sizes as instrument variables in both the sell-side and buy-side estimations. On the sell side, market size is defined as the number of users who clicked on a given book in the app on a given day. On the buy side, market size is defined as the number of sellers who scanned that book on that day. The left hand side of Equation (1) is buyer's utility, and it measures the utility of buying a book which should be independent from the number of potential book sellers on that day. However, buy-side market size can influence a book's sale price. A larger buy-side market means inventory is likely to rise in the short run, and to keep inventory at a relatively stable level, the platform may respond by lowering the sale price. Buy-side market size is considered exogenous because it reflects the arrival rate of potential book sellers, a supply shifter, and is therefore uncorrelated with current consumer demand. Similarly, sell-side market size can serve as an instrument, since fluctuations in the number of arriving buyers do not systematically affect the unobserved utility of any individual consumer who has already arrived. The first stage F-statistics are far above 10 (Lee et al. 2022) in both the sell side estimation and buy side estimation as shown in Table 3. The parameters are estimated using control function approach.

3.2.2 Book Sellers

On the buy side, sellers observe only the buy price of their own book and cannot change the book's quality. A nested logit structure is therefore not applicable in this setting, and I instead adopt a simple logit model. Seller i's utility from selling book j is given by

$$u_{ijt} = \alpha p_{jt} + c_j + \epsilon_{ijt}, \tag{7}$$

where p_{jt} is the buy price offered by the platform. The probability that the seller chooses to sell the book is then

$$\Pr(\text{sell}) = \frac{\exp(\alpha p_{jt} + c_j)}{1 + \exp(\alpha p_{jt} + c_j)}$$

Similarly, the parameters of this model are recovered via weighted maximum likelihood estimation, and both sell side market size and buy side market size are used as instrument variables here.

Table 3 presents the estimation results. The demand-side price elasticity is approximately 4, considerably more elastic than the supply-side price elasticity of around 1. This finding is intuitive, as book buyers have many substitutes available to them. If the price of a particular used copy rises, they can easily switch to another seller offering the same title, buy a different edition or quality, purchase from a competing platform, borrow from a library, buy the e-book, or choose a different book. Book sellers face different incentives. They cannot alter the physical quality of their used books, and their outside options are limited to donating the books, discarding them, or listing them on alternative platforms (either selling to individual buyers or to a platform directly). The first two options generate very little utility, while switching to another platform to sell directly to individual buyers requires extra effort (taking pictures of the book, adding a description, creating the listing, communicating with buyers, and arranging shipping). Nesting parameters exhibit substantial variation across book titles, with some titles having nesting parameters close to 0, suggesting limited substitution across qualities, whereas others have nesting parameters exceeding 0.5, indicating stronger substitution across qualities.

Table 3: Demand and Supply Side Estimation Results

	Sell Side	Buy Side
α	-0.266 (0.001)	0.263 (0.080)
Quality FE	-1.424 (0.049)	0.419 (0.136)
Book Fixed Effects		
Book FE_1	1.744 (0.038)	-3.428 (0.801)
Book FE_2	-0.108 (0.039)	-1.894 (0.174)
Book FE_3	0.861 (0.045)	-3.180 (0.593)
Book FE_4	1.909 (0.031)	-2.367 (0.644)
Book FE_5	3.883 (0.036)	-3.664 (1.010)
Nesting Parameters (Sell Side Only)		
σ_1	0.540 (0.041)	
σ_2	0.042 (0.064)	
σ_3	0.000 (0.033)	
σ_4	0.000 (0.051)	
σ_5	0.215 (0.090)	
Implied Price Elasticities		
High Quality	-4.353	1.315
Medium Quality	-3.871	1.076
First-Stage F-Statistics (Instrument Variables)		
First-Stage F	106.931	1108.815

Notes: Standard errors in parentheses, obtained via 100 bootstrap replications. Book FE are book fixed effects. σ_j are book-specific nesting parameters (sell side). Quality FE is the quality fixed effect for medium-quality books. Elasticities are evaluated at the sample mean.

4 Model

4.1 Static Pricing Model

In the static pricing model, the platform sets the buying and selling prices once and does not adjust them over time. Prices are chosen such that the buying rate and selling rate are approximately equal, an assumption motivated by empirical observations from the data. When the two rates are equal, the long-run inventory level follows a random walk.

I model the number of sales and purchases in each period as independent Poisson random variables with rates λ_s and λ_b, respectively, so that the net inventory change follows a Skellam distribution with parameters (λ_s, λ_b).

The random walk behavior of inventory is consistent with what is observed in the data, where inventory levels appear volatile with no clear drift. Since most books have buying and selling rates below 20, the Skellam variance, which scales with $\lambda_s + \lambda_b$, remains relatively small, limiting how far the inventory deviates in practice. Simulation experiments over 1,000 periods with $\lambda_s = \lambda_b = 30$ confirm that inventory levels remain within a reasonable range.

The platform's true objective is not directly observable, whether it maximizes per-period profit or revenue, whether this is done across all book titles jointly, and whether the platform behaves myopically or not. The selling price influences how long a book is expected to remain in inventory; a higher p_s results in a longer holding period and, consequently, higher holding costs. As a reasonable simplification, I assume the platform sets prices to maximize per-period profits, and the objective function is

$$
\begin{aligned}
\max_{p_s, p_b} \quad & \pi(p_s, p_b) = S(p_s)(p_s - p_b)\,[1 - \gamma(p_s, p_b, N)] \\
& \qquad - cN\,\mathbb{E}[\text{holding days} \mid p_s, N] \\
\text{s.t.} \quad & S(p_s) = B(p_b).
\end{aligned} \tag{8}
$$

where $\gamma(p_s, p_b, N)$ is the stock-out rate which depends on two-sided prices p_s, p_b and inventory N.

When a user orders a book, a staff member retrieves it from the shelf by picking randomly from all available copies. This means the inventory level directly influences selection likelihood. The fewer copies on the shelf, the higher the chance any particular copy is picked, but each copy has an equal probability among those present. In other words, this is a random selection process, not a first-in-first-out one. The expected holding time of a book depends on the inventory level, sell rate, and buy rate. Because inventory level follows a random walk, the platform uses the mean historical inventory as a proxy for the overall future inventory level when estimating holding time. It is hard to generate a meaningful closed-form expression for the expected holding time, so I use simulation to recover the holding cost c.

The simulation proceeds as follows. I first select the top 10 book titles from the data and, for each, compute the average daily inventory, buy rate, sell rate, buy price, and sell price; the average daily inventory serves as the baseline initial inventory level in the simulation. I then assume that buys and sells follow independent Poisson processes, with rate parameters set to the empirical average buy and sell rates, and for each book I simulate over 1,000 periods across a range of initial inventory levels around the empirical mean, recording the stockout rate and average holding days under each level. I further assume that the platform sets prices statically in an optimal manner, and calibrate the holding cost so that the empirical mean inventory level is optimal under this assumption. Because small perturbations around the mean (e.g., ±1 unit) produce negligible variation in average holding days and stockout rates, the simulation instead varies the initial inventory level in increments of 10 units to obtain meaningful results. While the simulation cannot pin down a specific value for the holding cost c, it provides a meaningful range $[\underline{c}, \bar{c}]$, which is then used as an input in the dynamic pricing analysis. Higher inventory reduces stockouts and missed sales but increases holding time and total holding costs. At the optimal inventory level N, adding k additional units is unprofitable because the benefit does not exceed the cost:

$$(p_s - p_b)\,\Delta\text{StockoutRate} \;\leq\; c\,[N\,\Delta\mathbb{E}[\text{HoldingDays}] + k\,\mathbb{E}[\text{HoldingDays} \mid N + k]]\,, \tag{9}$$

which yields a lower bound $\underline{c}$ on c. Symmetrically, removing k units is unprofitable because the cost savings do not exceed the loss in revenue:

$$(p_s - p_b)\,\Delta\text{StockoutRate} \;\leq\; c\,[(N-k)\,\Delta\mathbb{E}[\text{HoldingDays}] + k\,\mathbb{E}[\text{HoldingDays} \mid N]]\,, \tag{10}$$

which yields an upper bound $\bar{c}$ on c.

Table 4 presents the calibrated lower and upper bounds of holding costs for the top 5 book titles. In theory, copies of the same book in different quality conditions should share identical holding costs, since they occupy the same physical size. Books with higher inventory levels exhibit relatively lower holding costs. Although the holding cost bounds vary across these titles, the values are quite small in absolute terms. Even when a book is held for over 1,000 periods (nearly three years), the accumulated holding cost remains under one dollar per book.

4.2 Dynamic Pricing Model

The dynamic pricing model allows prices to respond to inventory levels. I assume that for every unsold book, the platform has to pay a per-period, time-invariant holding cost c. The holding cost is incurred at the beginning of the period before any transaction is realized. In the optimization, I adopt the median holding cost calibrated from the static pricing analysis, 0.0002.[2] To assess the sensitivity of the results to this parameter, I also examined alternative values of 0.0001, 0.0003, and 0.001; the optimal pricing solution was unaffected by these variations.

In the dynamic setting, the state variable is the inventory vector by quality $N = (N_h, N_m)$,

2. For each book title, I compute the median of the upper and lower bounds, and then take the mean of these medians across titles.

Table 4: Calibrated Holding Cost Bounds by Book and Quality

Book ID	Quality	Average Inventory	Lower Bound ($\underline{c}$)	Upper Bound ($\bar{c}$)
1	High	11	0.000390	0.000634
1	Medium	2	0.000447	NaN
2	High	19	0.000124	0.000871
2	Medium	14	0.000351	0.000363
3	High	34	0.000072	0.000526
3	Medium	23	0.000180	0.000300
4	High	260	0.000004	0.000009
4	Medium	145	0.000005	0.000006
5	High	38	0.000106	0.000551
5	Medium	14	0.000283	0.000545

Notes: Average inventory refers to the empirical mean inventory level used in the calibration procedure. The calibrated holding cost for each book-quality pair is bounded between $\underline{c}$ and $\bar{c}$.

and the choice variables are the corresponding price vectors on the sell side and buy side, $p^s = (p^s_h, p^s_m)$ and $p^b = (p^b_h, p^b_m)$. In each period, the platform sets buy and sell prices simultaneously. Assuming a book acquired today becomes available for sale the following day, the value function is defined as follows, given the expected sales and purchases:

$$\begin{aligned} V(N_h, N_m) = \max_{p^s_h, p^s_m, p^b_h, p^b_m} \Bigg\{ & \sum_{q\in\{h,m\}} p^s_q \, \mathbb{E}\big[Q^s_q(p^s_q; N_h, N_m)\big] - c(N_h + N_m) \\ & + \delta\, \mathbb{E}\big[V(N'_h, N'_m) \mid N_h, N_m, p^s_q, p^b_q\big] - \sum_{q\in\{h,m\}} \mathbb{E}\big[Q^b_q \mid p^b_q, N_h, N_m\big] \Bigg\}. \end{aligned} \tag{11}$$

The law of motion is given by

$$N'_q \mid \left(N_h, N_m, p^s_q, p^b_q\right) \sim F_{N'_q}\left(\cdot \mid N_h, N_m, p^s_q, p^b_q\right) \tag{12}$$

where $F_{N'_q}$ is induced by the joint distribution of $\left(Q^s_q, Q^b_q\right)$ through

$$N'_q = \max\left\{0, N_q - Q^s_q\right\} + Q^b_q \tag{13}$$

The Bellman equation is solved via policy iteration, which converges faster than value iteration in this setting. The selling price for each title is bounded above by the book's cover price, since the second-hand price cannot exceed the original retail price, and the buyout price is bounded below by 0.1. Because the two-sided pricing problem is solved jointly across the two quality conditions (high and medium), with inventories serving as the state variables, the state space grows quickly. Each inventory level can range from zero to several hundred units, producing a large number of states. To keep the problem tractable, I impose an upper bound on inventory for each book title. The reason is that as inventory grows, both the optimal selling and buyout prices decline, and beyond some threshold, the platform no longer finds it profitable to acquire additional copies. Also, the optimal buyout and sale prices are convexly decreasing and remain relatively stable in high inventory states. And once inventory reaches this cap, the acquisition window shuts down for that title. Hence, to make the price optimization tractable, I introduce several simplifications motivated by the data. I set the price increment to 1, allowing prices to be adjusted each time inventory changes. I also set the inventory upper bound at 100 for both book qualities, so once inventory for a given quality reaches this cap, the platform stops acquiring more copies of that quality.

These simplifications can be relaxed in future work to recover more precise optimal prices. As a result, the solution obtained here is not the exact optimum but should be close to it. Consequently, in the counterfactual analysis that follows, the profit gains relative to the platform's current static pricing should be interpreted as a lower bound on the true gains achievable under fully optimal dynamic pricing.

5 Counterfactual

5.1 Counterfactual A: Dynamic Pricing versus Static Pricing

The most straightforward counterfactual examines whether adopting dynamic pricing instead of the platform's current static pricing strategy would increase profits and by how much. As Table 5 shows, dynamic pricing raises profits across all books, with gains ranging from 0.86 to 12.14 percent, and the magnitude varies considerably by book. Dynamic pricing also increases the variance of profits. Because dynamic pricing allows two-sided prices to adjust with inventory levels, and inventory is in turn linked to market size volatility, profits under dynamic pricing are inherently more volatile. For Book 3, the variance of profits nearly doubled, likely because Book 3 exhibits high market size volatility on both the buy and sell sides. The profit gains arise through two distinct channels. Either lower prices paired with higher trading volume, or higher prices paired with lower volume. In both cases, the net effect is an increase in overall profits. There is, however, no consistent pattern showing that dynamic pricing reduces stockout rates.

The other counterfactual analysis examines the conditions under which the advantage of dynamic pricing over static pricing becomes more pronounced. To ensure a fair comparison,

Table 5: Comparison of Static and Dynamic Pricing

Panel A: Percentage Changes from Static to Dynamic Pricing

Book ID	Δ Profit (%)	Δ Var(Profit) (%)	Δ Sell Volume (%) (High, Low)	Δ Buy Volume (%) (High, Low)	Δ Var(Sell Volume) (%) (High, Low)	Δ Var(Buy Volume) (%) (High, Low)
1	1.76	13.52	(-2.81, -2.46)	(-2.84, -2.83)	(14.31, 28.79)	(0.12, -1.38)
2	0.86	1.09	(2.96, 13.42)	(-3.18, 13.80)	(-6.68, 245.50)	(-0.56, 19.19)
3	12.14	96.57	(16.16, 16.77)	(16.35, 17.15)	(84.99, 73.07)	(31.75, 27.21)
4	7.55	42.49	(31.12, -2.45)	(31.46, -9.95)	(78.89, -1.44)	(52.21, -10.89)
5	2.22	51.67	(16.02, 17.17)	(15.95, 17.23)	(52.29, 37.26)	(29.25, 22.06)

Panel B: Stockout Rates and Inventory Outcomes

Book ID	Static Stockout (%) (High, Low)	Dynamic Stockout (%) (High, Low)	Δ Var(Inventory) (%) (High, Low)	Δ Stockout Rate (%) (High, Low)
1	(8.32, 16.30)	(2.59, 64.42)	(-61.19, -86.57)	(-65.25, 348.98)
2	(1.56, 11.16)	(18.42, 4.29)	(-95.79, -25.02)	(2041.54, -59.72)
3	(15.77, 13.61)	(7.58, 6.77)	(17.29, 17.66)	(-51.20, -49.12)
4	(25.59, 0.16)	(1.54, 7.31)	(210.72, -97.22)	(-93.97, 2056.31)
5	(5.43, 2.89)	(3.49, 3.03)	(-78.41, -61.57)	(-22.85, 100.74)

Notes: High and Low refer to high-quality and low-quality books, respectively. All percentage changes are calculated from static pricing to dynamic pricing. The reported statistics are averages obtained from 200 simulation runs, with each simulation spanning 500 periods.

I held the expected market sizes constant while varying the mean and variance of the bivariate log-normal distribution, and then compared the simulated profits from dynamic versus static pricing under these conditions.

Since optimal static prices may shift as the mean and variance of the market-size distribution change, I used a simulation-based approach. I varied sell-side market size σ from 0.1 to 1.0 while holding the expected market size constant, then iterating over possible combinations of buyout and sale prices. The price increment was set to ¥1 (comparable to a $1 increment in the U.S. currency system). The results indicate that the optimal static prices are largely insensitive to changes in market-size volatility. To the extent they do respond, both the sell-side and buy-side static prices decrease as sell-side market-size volatility increases. Meanwhile, average simulated profit declines as sell-side market-size volatility rises, as shown in Figure 7.

5.2 Counterfactual B: Reaction to Demand or Supply Shocks

In this counterfactual analysis, I examine how static pricing and dynamic pricing respond differently when the market is hit by a demand or supply shock.

I conduct two sets of simulations. In the first set of simulations, I run 100 simulations,

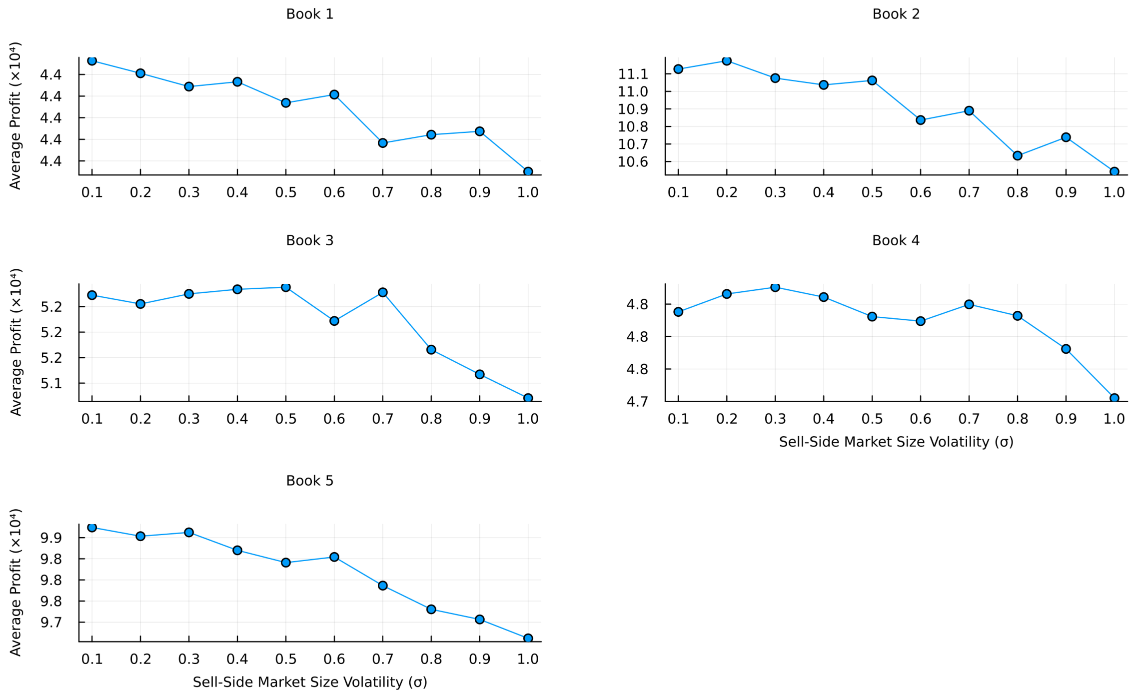

Figure 7: Sell-Side Profits by Market Size Volatility (Static Prices)

Notes: Each panel reports the average profit (in units of 10^4) for Books 1–5 as a function of sell-side market size volatility $\sigma \in \{0.1, \ldots, 1.0\}$. Averages are computed over 60 simulations of 500 periods each. For each σ, the static prices used are the corresponding optimal static prices.

each spanning 500 periods. Starting at t=100, I introduce a positive demand shock[3] that increased the number of daily clicks by roughly 60% and lasted for about 60 periods. I then compare how dynamic pricing and static pricing respond to this shock. Specifically, I compute the percentage change in profit before, during, and after the 60-day shock period, and find that the percentage increase in profit is larger during the positive demand shock. This result highlights the advantage of dynamic pricing. Because the demand shock naturally introduces additional inventory volatility, dynamic pricing lets the firm set higher buy and sale prices when inventory is low and lower prices when inventory is high. As shown in Figure 8, when thereunder a positive demand shock, dynamic pricing raises both the buyout and sale prices, adjusts prices more frequently, and keeps inventory at roughly the same level as before the . Under static pricing, prices remain fixed, so inventory drops sharply during the demand shock period, which can lead to a higher stockout rate. In the second set of simulations, I introduced a negative supply shock that reduced supply by roughly 40% and gradually faded out over 60 days. As shown in Figure 9, When a negative supply shock occurs and fewer booksellers

3. The demand shock follows an AR(1) process.

show up, the platform raises its buyout price to attract supply and also increases the sale price to account for the declining inventory. Comparing the percentage change in profit from static to dynamic pricing, I again find that the profit gain is larger during the negative supply shock period for all books. Together, the two sets of simulations suggest that the advantages of dynamic pricing become more pronounced when an unexpected demand or supply shock occurs, and that dynamic pricing responds to market size shocks more effectively than static pricing.

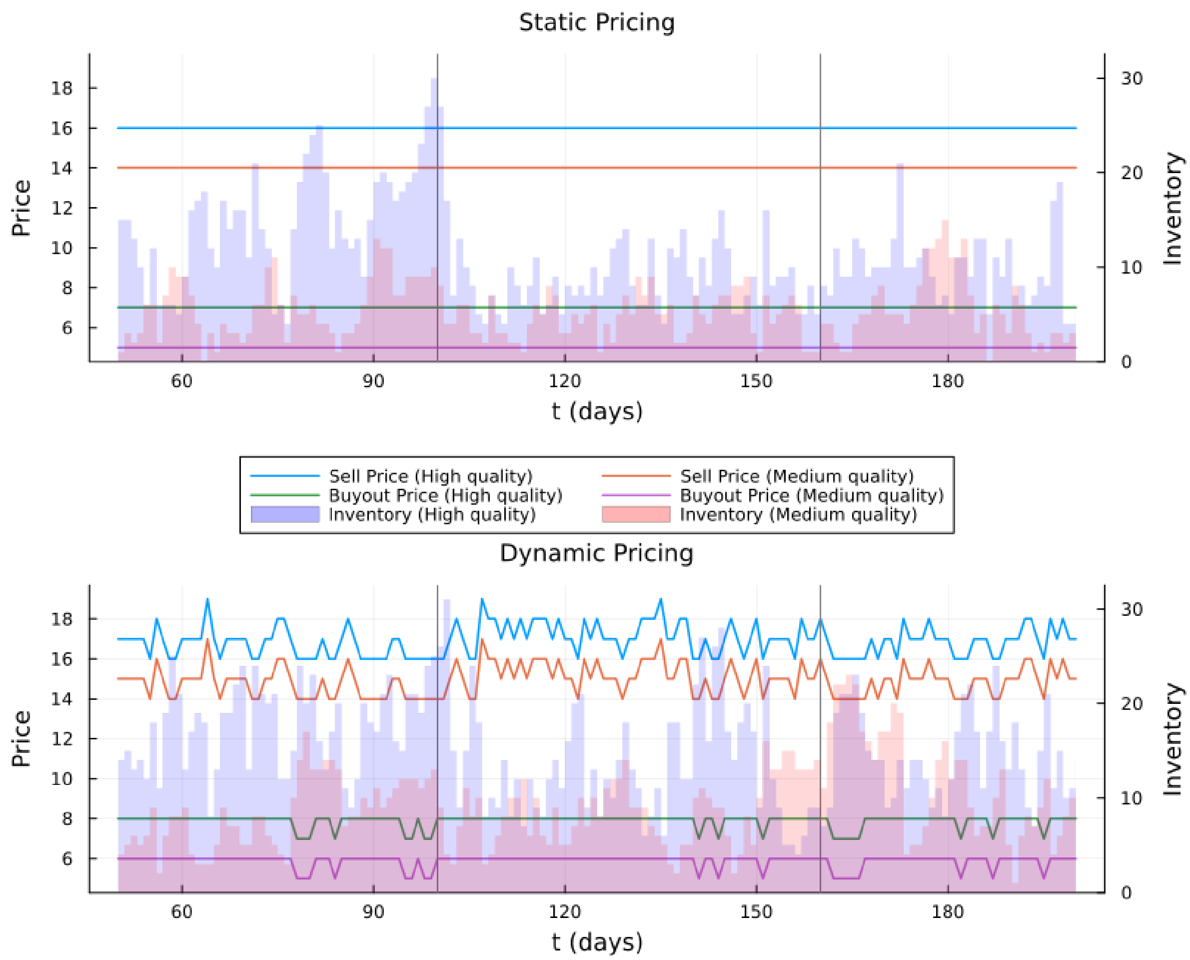

Figure 8: Static vs. Dynamic Pricing Under a Positive Demand Shock (t = 100–160)

Notes: The graph shows the results of a single simulation illustrating how static and dynamic pricing respond to a demand shock. At $t = 100$, I introduce a positive demand shock that raises the number of book buyers by 30%. The shock follows an AR(1) process and decays gradually over the following 60 days, so that by $t = 160$ only 5% of the original shock remains.

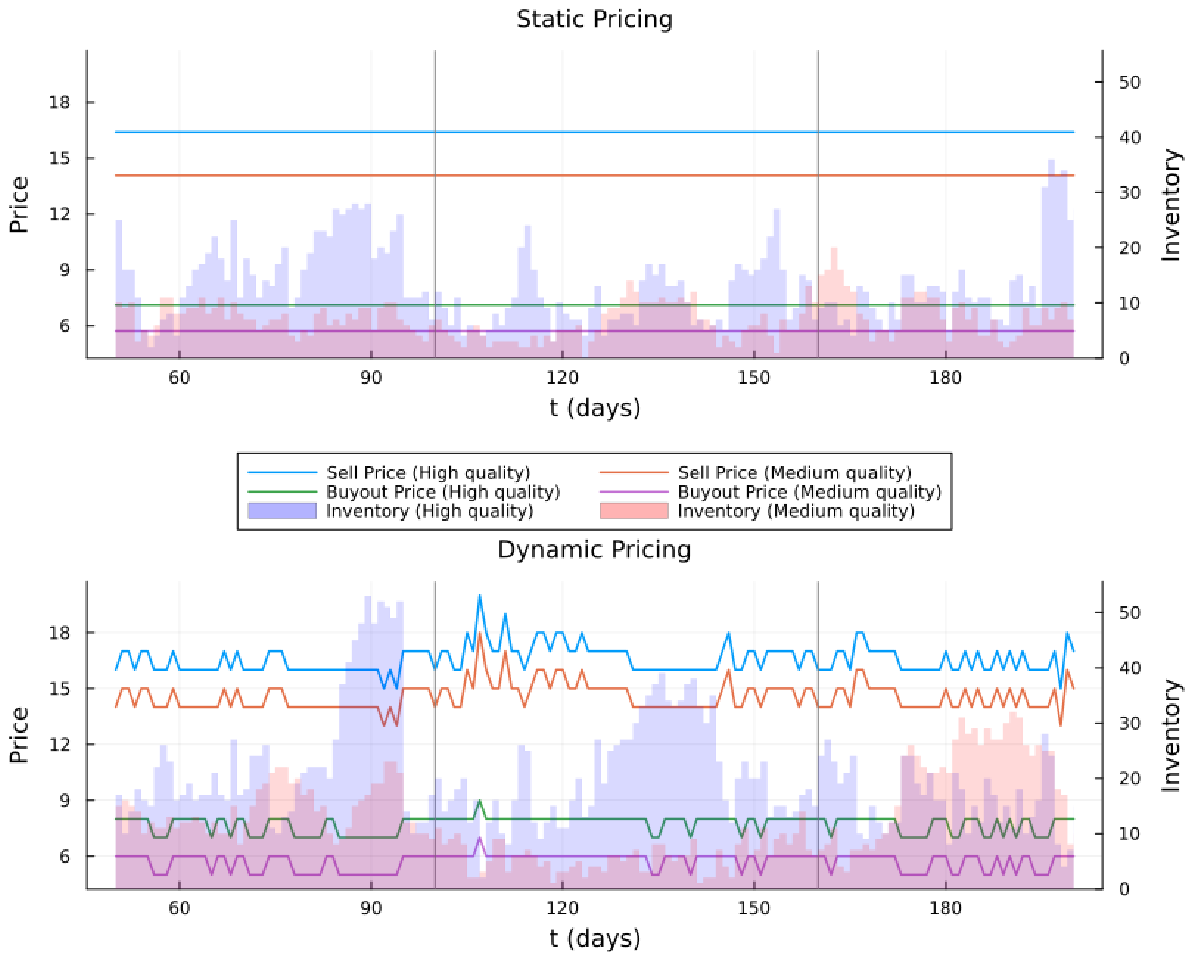

Figure 9: Static vs. Dynamic Pricing Under a Negative Supply Shock (t = 100–160)

Notes: The graph shows the results of a single simulation illustrating how static and dynamic pricing respond to a supply shock. At $t = 100$, I introduce a negative supply shock that reduces the number of book sellers by 40%. The shock follows an AR(1) process and decays gradually over the following 60 days, so that by $t = 160$ only 5% of the original shock remains.

5.3 Counterfactual C: AI-Generated Pricing Strategies

With the growth of AI, it would be interesting to see, given the data, how AI would solve the two-sided pricing problem. To do this, I chose the Claude Opus 5 model and fed it the data I used. Before giving the prompt, I reset Claude's memory to ensure it had no memory from any previous conversations. In the prompt, I explained how the platform operates, described the data structure, and provided the calibrated holding cost while setting the discount factor to 0.98. However, I did not reveal how I would solve the pricing problem.

Even if I did not mention demand or supply functions at all, Claude still suggested fitting a multinomial logit demand model, and because there is limited within-book price variation, Claude pooled the top 20 books together and assumed they share the same price-sensitivity parameters. And on the supply side, Claude fits a binary logit model with book fixed effects and a shared buy price coefficient.

Claude first generated a static pricing strategy with the core idea that buy prices should be set in a way that induces just enough supply to cover the expected sales volume so the platform will not over- or under-acquire. For the static prices Claude provided, it raised the sale price for every book and undercut the buy price for most of them. Using the Claude-suggested static prices, I ran 200 simulations of 500 periods each and compared the results to the platform's current static pricing strategy. Claude's suggested pricing turned out to perform worse, with average profits falling by 35% to 97%. Claude correctly identifies the insight about inventory flow and the relationship between buy price and sale price, but the prices it recommends don't actually lead to higher profits. In fact, because Claude raised sale prices and lowered buy prices, per-book profit margins were higher, but transaction volume dropped sharply and more than offset the margin gains.

Then I asked Claude to generate a dynamic pricing strategy. Because the day-of-week seasonality was too weak in the data, Claude chose inventory level as the dynamic lever and set inventory as the state variable. Claude follows a standard revenue-management feedback rule. It raises buy and sell prices when inventory is low, and does the opposite when inventory is high. Claude also sets a target inventory, and if inventory is above the target, both prices fall, and if inventory is below the target, both prices rise:

$$\begin{aligned} \text{sell price:} \quad p_{st} &= p_s^* + \lambda\,(I^* - I_t)\,, \\ \text{buy price:} \quad p_{bt} &= p_b^* + \mu\,(I^* - I_t)\,. \end{aligned}$$

where p_s^* and p_b^* are the anchor prices derived from the previous static pricing problems. Claude also added price clipping. Using Claude's calibrated parameters, I ran another 200 simulations of 500 periods each and compared the results against the platform's current static pricing. The results show that if using the Claude-suggested dynamic pricing, average profit would decrease by 0.39% to 27.16%, still worse than the platform's current static pricing but better than the static pricing strategy generated by Claude, mostly because the transaction volume is much higher using dynamic pricing.

Given the data and background information, Claude can provide the correct intuition for how to approach the optimal pricing problem. However, it does not perform as well in deriving a precise pricing strategy, and the resulting strategy can even underperform the pricing strategy used in the dataset.

6 Conclusion

This paper studies the two-sided pricing and inventory problem faced by a resale platform that buys used books from individual sellers and resells them to individual buyers. Using transaction-level data from Déjà Vu, I estimate the daily buy-side and sell-side market sizes as a bivariate log-normal distribution, a nested logit model of book buyers, and a logit model of book sellers. The estimates show that the sell-side market is four to ten times larger than the buy-side market and more volatile, that demand (elasticity around 4) is considerably more price-sensitive than supply (elasticity around 1), and that substitution across quality tiers varies substantially across book titles.

Building on these estimates, I develop an infinite-horizon dynamic pricing model in which the platform sets buyout and sale prices jointly as a function of inventory, and I compare it with a static benchmark that mirrors the platform's current practice. Dynamic pricing raises profits by 0.86% to 12.14% across the five books, either through higher transaction volume or through a wider spread between buyout and sale prices, although it also increases profit variance and does not consistently reduce stockout rates. The gains from dynamic pricing become larger when the market experiences unexpected demand or supply shocks, since the platform can adjust both prices to stabilize inventory. In contrast, pricing strategies generated by an AI model capture the correct intuition but underperform the platform's existing static prices.

Because the model restricts prices to integer increments and caps inventory at 100 units per quality, the estimated profit gains should be viewed as a lower bound on what fully optimal dynamic pricing could achieve. Future work could relax these simplifications, extend the analysis beyond the most popular titles, and allow for substitution across book titles.

References

Akerlof, George A, et al. 1970. "The market for 'lemons': quality uncertainty and the market mechanism." *Quarterly journal of economics* 84 (3): 488–500.

Armstrong, Mark. 2005. "Competition in two-sided markets."

Banerjee, Siddhartha, Carlos Riquelme, and Ramesh Johari. 2015. "Pricing in ride-share platforms: A queueing-theoretic approach." *Available at SSRN 2568258.*

Betancourt, Jose M, Ali Hortaçsu, Aniko Oery, and Kevin R Williams. 2022. *Dynamic price competition: Theory and evidence from airline markets.* Technical report. National Bureau of Economic Research.

Blundell, Richard, Ran Gu, Søren Leth-Petersen, Hamish Low, and Costas Meghir. 2026. "Durables and lemons: Private information and the market for cars." *Quantitative Economics* 17 (1): 38–91.

Cho, Sungjin, Gong Lee, John Rust, and Mengkai Yu. 2018. "Optimal dynamic hotel pricing." In *2018 Meeting Papers,* vol. 179.

Cho, Sungjin, and John Rust. 2010. "The flat rental puzzle." *The Review of Economic Studies* 77 (2): 560–594.

Ellison, Glenn, and Sara Fisher Ellison. 2018. *Match quality, search, and the Internet market for used books.* Technical report. National Bureau of Economic Research.

Emons, Winand, and George Sheldon. 2002. "The market for used cars: A new test of the lemons model."

Gao, Ming. 2018. "Platform pricing in mixed two-sided markets." *International Economic Review* 59 (3): 1103–1129.

Hagiu, Andrei. 2009. "Two-sided platforms: Product variety and pricing structures." *Journal of Economics & Management Strategy* 18 (4): 1011–1043.

Lee, David S, Justin McCrary, Marcelo J Moreira, and Jack Porter. 2022. "Valid t-ratio Inference for IV." *American economic review* 112 (10): 3260–3290.

Lewis, Gregory. 2011. "Asymmetric information, adverse selection and online disclosure: The case of eBay motors." *American Economic Review* 101 (4): 1535–1546.

MacKay, Alexander, and Samuel N Weinstein. 2022. "Dynamic pricing algorithms, consumer harm, and regulatory response." *Wash. UL Rev.* 100:111.

Nair, Harikesh. 2006. "Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games."

Peterson, Jonathan R, and Henry S Schneider. 2014. "Adverse selection in the used-car market: evidence from purchase and repair patterns in the Consumer Expenditure Survey." *The RAND Journal of Economics* 45 (1): 140–154.

Porter, Robert H, and Peter Sattler. 1999. *Patterns of trade in the market for used durables: Theory and evidence.*

Research and Markets. 2024. "Second Hand Books Market Trends and Forecasts to 2031: Market Set to Achieve US$45 Billion by 2031, Due to Growing Eco-Consciousness and Increasing Literacy Rates." Yahoo Finance, January. Accessed July 22, 2026. https://finance.yahoo.com/news/second-hand-books-market-trends-144300810.html.

Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform competition in two-sided markets." *Journal of the european economic association* 1 (4): 990–1029.

———. 2006. "Two-sided markets: a progress report." *The RAND journal of economics* 37 (3): 645–667.

Romer, David. 2006. "Do firms maximize? Evidence from professional football." *Journal of Political Economy* 114 (2): 340–365.

Rysman, Marc. 2009. "The economics of two-sided markets." *Journal of economic perspectives* 23 (3): 125–143.

Schittekatte, Tim, Dharik Mallapragada, Paul L Joskow, and Richard Schmalensee. 2024. "Electricity retail rate design in a decarbonizing economy: an analysis of time-of-use and critical peak pricing." *The Energy Journal* 45 (3): 25–56.

Siegert, Caspar, and Robert Ulbricht. 2020. "Dynamic oligopoly pricing: Evidence from the airline industry." *International Journal of Industrial Organization* 71:102639.

Spann, Martin, Marco Bertini, Oded Koenigsberg, Robert Zeithammer, Diego Aparicio, Yuxin Chen, Fabrizio Fantini, Ginger Zhe Jin, Vicki Morwitz, Peter Popkowski Leszczyc, et al. 2024. *Algorithmic pricing: Implications for marketing strategy and regulation.* Technical report. National Bureau of Economic Research.

Williams, Kevin R. 2022. "The welfare effects of dynamic pricing: Evidence from airline markets." *Econometrica* 90 (2): 831–858.

Appendix

A Appendix A

Table A.1: Mapping Between Book IDs and Book Hash IDs

Book ID	Book Hash ID
1	336769107332822685
2	394134705886066823
3	425011221835698110
4	578270280582410845
5	590575417199867722

Notes: For ease of presentation, I assign each book a numerical identifier from 1 to 5. This table reports the mapping between these book IDs and the original book hash IDs. The numerical book IDs are used throughout the remaining figures and tables.

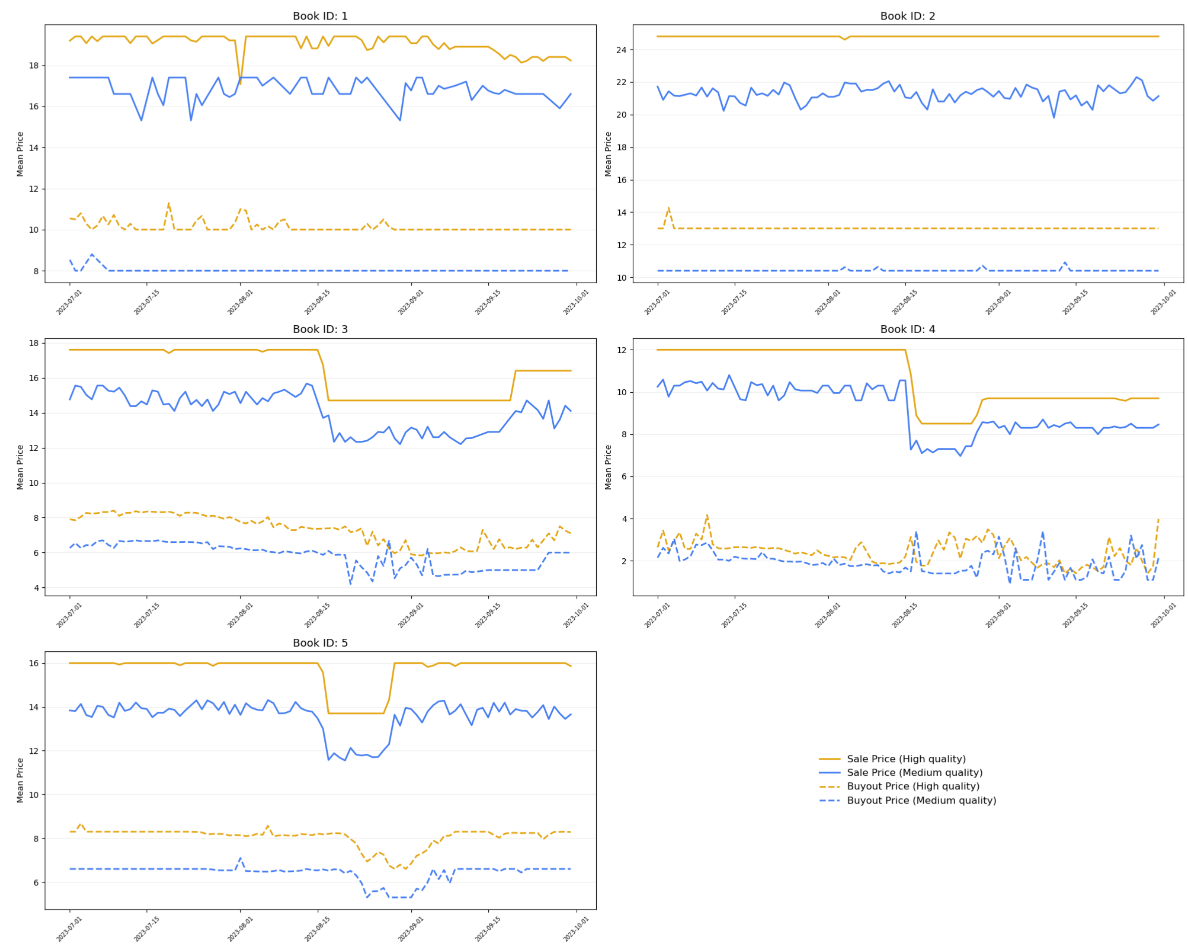

Figure A.1: Two-sided Price Trends by Book and Quality

Table A.2: Summary Statistics by Book Category and Quality

Book Category	Quality	Total Vol.	Sell Price	Buy Price	Hold Time(hr)	Cover Price	Daily Vol.
ACGN	high	41868	24.5	10.3	741	6029.7	356
ACGN	medium	10972	17.3	6.3	898	4556.3	93
Art	high	49169	34.6	15.1	855	8355.6	407
Art	medium	21195	27.0	10.1	1016	6815.3	174
Children's Books	high	217935	18.1	6.3	1041	4864.5	1846
Children's Books	medium	83531	13.7	4.3	1283	4051.8	670
Chinese Literature	high	84586	20.8	8.7	608	4971.4	776
Chinese Literature	medium	48599	15.7	5.6	798	4059.2	436
Design	high	16131	35.5	15.3	1056	8791.1	130
Design	medium	8210	28.2	10.4	1471	7882.9	66
Economy & Business	high	51104	25.0	9.0	1149	6600.7	430
Economy & Business	medium	33200	19.7	6.1	1524	5902.0	271
Foreign Literature	high	108680	20.7	8.5	693	4988.6	961
Foreign Literature	medium	52926	16.0	5.5	857	3976.7	458
Genre Fiction	high	87726	19.9	7.3	824	5095.6	763
Genre Fiction	medium	36277	15.9	5.2	1090	4362.6	311
History & Politics	high	84745	29.9	12.7	786	7043.1	734
History & Politics	medium	35562	23.7	8.6	832	5642.8	310
Humanities & Social Sciences	high	63890	25.2	10.8	696	5859.2	567
Humanities & Social Sciences	medium	36198	20.4	7.6	862	4991.0	319
Life Topics	high	100901	20.0	7.4	912	5205.8	876
Life Topics	medium	61983	16.1	5.0	1362	4685.2	519
Lifestyle	high	45064	23.8	9.7	860	6154.9	375
Lifestyle	medium	16602	18.8	6.7	1231	5638.3	135
Original Editions	high	20827	41.9	21.0	1298	9319.9	153
Original Editions	medium	10245	33.0	14.7	1544	8839.1	73
Other Literary Forms	high	133922	22.3	9.1	749	5350.4	1169
Other Literary Forms	medium	71558	17.3	6.0	981	4336.0	612
Philosophy & Religion	high	43888	24.1	10.4	586	5266.8	397
Philosophy & Religion	medium	26348	20.3	7.6	680	4488.4	235
Science & Popular Science	high	73085	26.1	11.0	750	6524.9	632
Science & Popular Science	medium	35313	21.2	7.6	918	5631.0	301
Textbooks & Reference	high	70486	22.9	8.1	1120	6494.6	622
Textbooks & Reference	medium	93087	16.6	5.0	2010	5750.0	840