Competitive Effects of Minimum Resale Price Maintenance: Evidence from the Publishing Industry*

Kohei Kawaguchi[†] Jeff Qiu[‡] Yi Zhang[§]

March 22, 2026

Abstract

Minimum resale price maintenance (RPM) is treated strictly in most jurisdictions, although in the United States it is subject to the rule of reason. Although minimum RPM can, in theory, improve consumer surplus, whether it does so empirically remains unclear. This paper evaluates minimum RPM in a setting where its procompetitive justification is theoretically relevant. Specifically, we study whether minimum RPM in the Japanese publishing industry improves consumer surplus by inducing retailers to hold more inventory under demand uncertainty. We develop a model of bookstores and a publisher under RPM and estimate its parameters using comprehensive book delivery and sales data. Counterfactual simulations show that minimum RPM decreases consumer surplus whereas improves supply-side welfare. A parameter region in which minimum RPM improves consumer surplus does exist, but it occupies only a limited part of the plausible parameter range we consider. Based on these results, we argue that the procompetitive justification for minimum RPM is better confined to exemption criteria than used in illegality judgments.

*Keywords: Resale price maintenance; Return policy; Wholesale model; Vertical relationship; Antitrust law; Demand uncertainty; Inventory; Product variety; Publishing industry; Bookstores. JEL Codes: L11; L13; L42; L81; L82; M31; K21. This study was financially supported by GRF 16501622, Hong Kong and JST ERATO JPMJER2301, Japan. We use data provided by one of the book wholesalers in Japan. We did not receive any financial support from them. Before the data was provided, we had a contract that prohibits the data provider from intervening in the editorial process.

[†]Department of Economics, School of Business and Management, Hong Kong University of Science and Technology. e-mail: kkawaguchi@ust.hk

[‡]Department of Economics and Finance, Lang School of Business, University of Guelph. e-mail: yjqiu@uoguelph.ca

[§]Department of Marketing, Guanghua School of Management, Peking University. e-mail: yizhang@gsm.edu.pku.cn

1 Introduction

Resale price maintenance (RPM) remains contested in competition law. In the United States, both maximum and minimum RPM are judged under the rule of reason at the federal level, while some states still apply a stricter rule. In the EU, RPM is still treated strictly in practice even after *Super Bock* required attention to competitive effects. In Japan, RPM remains an unfair trade practice absent justification, although copyrighted physical goods, including books, receive a special exemption. Across these legal regimes, the same question arises: when a manufacturer controls resale prices, should competition law treat the practice mainly as a restriction on downstream competition, or can it sometimes improve consumer welfare?

That question is harder for minimum RPM. A maximum price ceiling can be justified more easily because it limits retail markups and may mitigate double marginalization. A minimum price floor is harder to justify because it directly restricts discounting. Economic theory offers several reasons why minimum RPM may improve welfare, including service provision, free riding, and inventory under demand uncertainty. But it is not clear whether these effects are important in practice. Courts, competition authorities, and policymakers therefore need to know whether the main procompetitive defenses of minimum RPM identify economically important cases or only rare ones. The answer matters for both illegality standards and the design of exemptions in sectors where RPM is legally permitted.

To bridge this gap between theory and empirics, we empirically study one setting in which minimum RPM may improve consumer welfare: the inventory channel (Deneckere et al., 1996, 1997; Butz, 1997; Flath and Nariu, n.d.; Krishnan and Winter, 2007). Retailers choose inventory before demand is realized, but they compete in prices only after demand is realized. At that stage, inventory is already in place, so the wholesale payment for those units is sunk and the effective marginal cost of selling an additional unit is low. Retailers therefore have an incentive to cut prices to take sales from rival retailers selling the same title. From the viewpoint of the vertical structure, this price competition can be excessive because it mainly reallocates sales across retailers rather than creating comparable gains for the supply side. If firms expect this ex post margin erosion in low-demand states, they may choose lower inventory ex ante. Minimum RPM can then have a procompetitive effect by limiting ex post undercutting, increasing inventory incentives, and lowering realized retail prices in high-demand states when higher inventory reduces the incidence of binding stock constraints. The same price floor can also reduce consumer welfare in low-demand states by keeping prices too high when the floor binds. Whether minimum RPM improves consumer welfare on net is therefore an empirical question. Maximum RPM raises a different question because it can mitigate double marginalization by limiting downstream markups.

We study the Japanese publishing industry because it is a useful setting for testing this theory.

RPM is common in this market. Demand uncertainty for new titles is high, so inventory choice matters. Returns are also common, which makes the inventory channel important because the supply side can tolerate overstock ex ante while still facing the effects of downstream pricing ex post. This combination of uncertainty, returns, and resale-price control is central in earlier theory as well (Marvel and Peck, 1995; Flath and Nariu, n.d.). The data include deliveries, sales, returns, and prices at a detailed level. Observing both endogenous prices and endogenous deliveries is important for identification because it allows us to separately recover two supply-side parameters: marginal costs and subjective demand uncertainty. We develop a structural model of publishers and bookstores in which inventory is chosen before demand is realized and estimate consumer demand and supply-side uncertainty using title-bookstore-month data.

In the observed institution, bookstore-level delivery is centralized on the wholesaler side and chosen before demand is realized. In the analysis, we treat the publisher and wholesaler as integrated for this delivery decision. The vertical relationship is governed by a revenue-sharing rule rather than a wholesale-pricing contract. The upstream firm sets the retail price before demand is realized, so price is committed in advance rather than adjusted to realized demand, and that price is uniform across bookstores, so the same title sells for the same posted price at every store. The status quo therefore combines centralized ex ante delivery choice, revenue sharing, upstream firm pricing, ex ante price commitment, and uniform pricing across stores.

The counterfactual analysis changes these features one at a time. We first allow the upstream firm to set different prices across bookstores while still choosing those prices before demand is realized. We then allow the upstream firm to set those bookstore-specific prices after demand is realized. Next, we transfer ex post pricing authority from the upstream firm to competing bookstores while keeping the centralized ex ante delivery choice fixed. Starting from that decentralized-pricing benchmark, we then examine what changes when the upstream firm can also impose either a minimum resale price or a maximum resale price on bookstore pricing. We also conduct comparative statics for the comparison between minimum RPM and no RPM by varying demand uncertainty and the intensity of downstream competition. Throughout these main counterfactuals, we hold fixed the centralized delivery system and the revenue-sharing rule; counterfactual changes in those margins are possible, but they are not the focus of this paper.

Relative to the benchmark in which the upstream firm can set bookstore-specific prices after demand is realized, the observed fixed-price system performs worse because uniform pricing and ex ante price commitment distort both deliveries and retail prices. When ex post pricing authority is transferred from the upstream firm to competing bookstores, expected retail prices fall and consumer surplus rises slightly, but deliveries and total surplus fall because downstream competition weakens inventory incentives. Adding a minimum resale price to that decentralized-pricing benchmark partly restores deliveries and supply-side welfare, but it also raises average retail prices and

lowers consumer surplus on average. Thus, the pro-competitive justification of minimum RPM based on the inventory channel fails in this specific context. Adding a maximum resale price has little effect in the estimated environment because the ceiling is not binding.

The comparative statics show that the effect of minimum RPM on consumer surplus depends on demand uncertainty and the intensity of downstream competition. Minimum RPM improves consumer surplus only in a limited region of the parameter range we consider, when demand uncertainty is moderate and downstream competition is strong. We interpret that range economically as follows: demand uncertainty, σ_0, runs from no prior title-level uncertainty to the estimated level, and the price-sensitivity parameter, α, spans values that imply roughly own-price elasticities from -1 to -10 at observed prices and shares. The decomposition clarifies the mechanism behind this result by separating the consumer-surplus effect into a low-demand-state loss and a high-demand-state gain. Consumer surplus rises only where the high-demand-state gain dominates the low-demand-state loss. When uncertainty is low, there is little inventory distortion to correct, so the high-demand-state gain is small. When uncertainty is very high, the price floor distorts prices in too many low-demand states, so the low-demand-state loss dominates.

These findings suggest that minimum RPM in the Japanese publishing industry is not well justified through the inventory channel we study. Comparative statics based on the estimated model also indicate that a pro-competitive justification for minimum RPM is not impossible, but it is not robust to key parameters, especially the degree of demand uncertainty and the market power of bookstores. This distinction matters for legal treatment. Minimum RPM restricts price in a way that tends to reduce consumer surplus unless sufficiently strong countervailing pro-competitive effects are present, so such effects are better treated as grounds for exemption rather than as a basis for legality. By contrast, maximum RPM differs because the price restraint itself can improve consumer surplus by mitigating double marginalization.

To the best of our knowledge, this is the first empirical structural paper to test whether minimum RPM can improve consumer surplus in a setting where theory predicts that it can, namely through the inventory channel under demand uncertainty. We answer this question by developing a structural model of a manufacturer and retailers under RPM that separates the role of minimum RPM from that of maximum RPM and allows us to decompose the channels through which minimum RPM affects consumer surplus. The analysis yields clear implications for the legal assessment of minimum RPM, in contrast to maximum RPM.

The remainder of the paper is organized as follows. Section 2 discusses the legal status of RPM and explains why decomposing fixed prices into minimum- and maximum-RPM components is useful for the legal and economic questions we study. Section 3 describes the Japanese publishing industry and the wholesaler data. Section 4 develops a stylized model of the pro-competitive inventory channel. Section 5 presents the empirical model, Section 6 reports the estimation, and

Section 7 presents the counterfactual welfare analysis.

1.1 Novelty and Contribution

Deneckere et al. (1996, 1997) theoretically showed that when demand is uncertain and inventory must be chosen before demand is realized, RPM can also prevent price wars among retailers in low-demand states and support higher inventory holdings. Closely related papers by Butz (1997) and Flath and Nariu (n.d.) also emphasize that uncertain demand can make minimum RPM privately or socially valuable by protecting ex post margins and encouraging ex ante distribution. Marvel and Peck (1995) shows that returns policies play a related coordination role under demand uncertainty, which is especially relevant in publishing. In a closely related vertical-control framework, Krishnan and Winter (2007) study downstream retailers that choose both price and inventory under uncertain demand and show that price floors or buyback policies can restore supply-chain coordination. Unlike these theoretical papers, our model matches the publishing institution by centralizing inventory upstream, focusing on the minimum-RPM channel rather than the broader contract-design problem, and abstracting from stockout-based substitution across bookstores, while sharing their core insight that minimum RPM can have pro-competitive effects by supporting inventory incentives under demand uncertainty.

The theoretical literature identifies several possible effects of RPM. For minimum RPM, Telser (1960) and Jullien and Rey (2007) argue that it can facilitate upstream collusion by making deviations from an upstream cartel less profitable and easier to detect, thereby reducing social welfare. Minimum RPM can also deter upstream entry by making retailers less willing to accommodate entrants (Asker and Bar-Isaac, 2014). On the other hand, minimum RPM can mitigate the "free-riding problem," in which a discount retailer free rides on the pre-sale investments of full-service retailers and then attracts consumers through low prices, thereby strengthening retailers' incentives to provide services (Telser, 1960; Marvel and McCafferty, 1984). For maximum RPM, many studies have argued that it can serve to eliminate the double-marginalization problem (Spengler, 1950; Tirole, 1988; Klein, 1999; Blair and Lafontaine, 1999). When the demand of retailers is positively correlated, it can also force retailers to internalize the positive externality and reduce prices (Blair and Lafontaine, 1999).

The welfare effects of minimum RPM therefore remain theoretically ambiguous. This paper contributes to the empirical RPM literature by evaluating RPM in an environment in which it may have a procompetitive effect through inventory choices made before demand is realized. Despite the policy importance of RPM, there are still relatively few empirical studies of its effects (MacKay and Smith, 2017). One strand of the literature studies the competitive effects of RPM using changes in RPM practice or policy (Gilligan, 1986; Ornstein and Hanssens, 1987; Bailey and Leonard, 2010)

or evidence from court cases (Ippolito and Overstreet, 1996). Another strand develops structural models of manufacturer-retailer relationships and evaluates RPM relative to alternative vertical contracts. The closest paper to ours is Xia (2025), which estimates a structural model of downstream and upstream markets in the Chinese pharmaceutical industry. He finds that maximum RPM improves welfare, although retail-price coordination across retailers attenuates the gain. He also uses an antitrust enforcement episode to provide model-free evidence on the procompetitive effects of RPM. We differ by analyzing both the minimum-RPM and maximum-RPM components of the conduct and by microfounding the inventory channel under demand uncertainty. Another related paper is Bonnet and Dubois (2010), which estimates a structural model of the bottled-water market under alternative vertical contracts: linear pricing, two-part tariffs, and two-part tariffs with RPM. They find that the two-part tariff with RPM best fits the data and that banning RPM would improve consumer welfare. Unlike that setting, we directly observe that RPM and revenue sharing prevail. RPM also has no procompetitive role in their model. In our model, by contrast, RPM can generate both procompetitive and anticompetitive effects through inventory decisions, market structure, and model parameters.

This paper also contributes to the literature on RPM in publishing. Tosdal (1915) describes the history of RPM in the German and US book markets. Bittlingmayer (1988) argues that RPM in the German book market functioned as a bookseller cartel. Flath and Nariu (n.d.) provides an early theory account tying demand uncertainty, returns, and book-like RPM institutions to Japan. Fishwick (2008) and Ball et al. (2008) study the abolition of RPM in the UK in 1995. Beck (2004) and Beck (2008) study the German book market around the abolition of RPM in 2008, and Williams (2024) studies the effects of abolition across the EU. Because RPM in publishing remains legal in Japan, our strategy does not rely on before-and-after policy variation. Instead, we use detailed institutional information and rich transaction data to recover the model primitives needed to evaluate the welfare implications of RPM in the Japanese publishing industry.

This paper is also related to the literature on the structural analysis of the publishing industry. Li (2021) estimates consumer preferences over book genres and reading formats, including e-books, print books from online retailers, and from offline bookstores, to investigate the optimal multichannel price strategy. Li (2019) estimates consumers' preferences in a dynamic discrete choice model of e-reader adoption and e-book purchases to study the optimal joint intertemporal price discrimination strategy. Daljord (2021) exploits the abolition of fixed price policy in the Norwegian book market to estimate consumers' preferences for books and the discount factor in a dynamic discrete choice model of books to study the implications for optimal durable-goods pricing. Our paper differs from them by highlighting a different aspect of publishers' pricing strategy, RPM and its competitive effects.

2 Legal Status of RPM

2.1 Classification and Burden of Proof

For legal analysis, the threshold issue is not only whether RPM can ever generate procompetitive effects, but also where those effects matter in the legal inquiry. If RPM is treated as effectively per se unlawful, the authority need only show that the practice falls within the prohibited category. If RPM is assessed under a rule-of-reason or effects-based framework, by contrast, the decision-maker must confront competitive effects more directly. This distinction matters for the present paper because our contribution is to quantify one setting in which RPM may have a procompetitive rationale in theory, namely the inventory channel under demand uncertainty.

The treatment of RPM differs across jurisdictions. In the United States, maximum RPM moved to rule-of-reason treatment in *State Oil v. Khan* (1997),[1] and minimum RPM followed in *Leegin Creative Leather Products v. PSKS* (2007).[2] Some states, however, still maintain a stricter approach to RPM. In the EU, RPM was long treated as a "by object" restriction under Article 101(1) TFEU,[3] but the Court of Justice in *Super Bock* required an assessment of actual or potential competitive effects.[4] Even there, defendants invoking Article 101(3) TFEU must satisfy demanding cumulative conditions.[5] In Japan, RPM remains an unfair trade practice absent justifiable grounds under the Antimonopoly Act,[6] and the JFTC's Guidelines place the burden on defendants to establish the required justification.[7][8]

This paper's legal motivation becomes clearer once fixed resale prices are decomposed into their upper- and lower-price components. Maximum RPM can be defended more naturally because it caps downstream prices and can mitigate double marginalization. Minimum RPM is harder to

[1] State Oil Co. v. Khan, 522 U.S. 3 (1997). Available at: https://supreme.justia.com/cases/federal/us/522/3/, accessed on July 17, 2025.

[2] Leegin Creative Leather Products, Inc. v. PSKS, Inc., 551 U.S. 877 (2007). Available at: https://supreme.justia.com/cases/federal/us/551/877/, accessed on July 17, 2025.

[3] Treaty on the Functioning of the European Union, Article 101. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:12016E101, accessed on July 17, 2025.

[4] Super Bock Bebidas SA and Others v Autoridade da Concorrência, Case C-211/22, Court of Justice of the European Union, 29 June 2023. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:62022CJ0211, accessed on July 17, 2025.

[5] Treaty on the Functioning of the European Union, Article 101(3). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:12016E101, accessed on July 17, 2025.

[6] Act on Prohibition of Private Monopolization and Maintenance of Fair Trade (Act No. 54 of April 14, 1947, as amended). Available at: https://www.jftc.go.jp/en/legislation_gls/amended_ama09/, accessed on July 17, 2025.

[7] Japan Fair Trade Commission, Guidelines Concerning Distribution Systems and Business Practices Under the Antimonopoly Act (June 23, 2017). Available at: https://www.jftc.go.jp/en/legislation_gls/imonopoly_guidelines_files/DistributionSystemsAndBusinessPractices.pdf, accessed on July 17, 2025.

[8] Designation of Unfair Trade Practices (Fair Trade Commission Notification No. 15 of June 18, 1982, as amended). Available at: https://www.jftc.go.jp/en/legislation_gls/unfairtradepractices.html, accessed on July 17, 2025.

defend because it directly restricts discounting. A persuasive defense of minimum RPM therefore requires both a concrete mechanism through which the restraint improves market performance along another margin and, under a consumer-welfare standard, evidence that those gains outweigh the direct price distortion in the states where the floor binds.

2.2 RPM in Publishing

Successful RPM defenses remain rare in practice. Recent enforcement actions in the EU and Japan continue to treat RPM strictly,[9][10][11][12] and, even in the United States, successful defenses have been limited. After *Leegin*, closer scrutiny tends to arise precisely in cases where RPM appears to facilitate collusion or where market power is a concern.[13][14][15][16]

The publishing industry is therefore a particularly informative case study. In Japan, copyrighted works in physical media are exempted from the general prohibition on RPM. In a 2001 notice, the JFTC recommended phasing out this exemption in the long term, but also confirmed its continuation subject to regular monitoring (Japan Fair Trade Commission, 2001). As a result, fixed book pricing remains widespread in Japan. Several EU member states also maintain fixed-price laws for books,

[9] European Commission, Cases AT.40465 (Asus), AT.40469 (Denon & Marantz), AT.40181 (Philips), and AT.40182 (Pioneer), July 24, 2018. Press release available at: `https://europa.eu/rapid/press-release_IP-18-4601_en.htm`, accessed on July 17, 2025.

[10] Autorité de la concurrence, Decision 24-D-11 of December 19, 2024 regarding practices in the household appliances sector. Available at: `https://www.autoritedelaconcurrence.fr/en/article/household-appliances-eu611-million-fines-imposed-12-companies-vertical-price-fixing`, accessed on July 17, 2025.

[11] Japan Fair Trade Commission, Cease and Desist Orders against Aprica Children's Products G.K. (July 1, 2019) and Combi Co., Ltd. (July 24, 2019). Available at: `https://www.jftc.go.jp/en/pressreleases/yearly-2019/July/190701.html` and `https://www.jftc.go.jp/en/pressreleases/yearly-2019/July/190724.html`, accessed on July 17, 2025.

[12] Japan Fair Trade Commission, Approval of the Commitment Plan submitted by ICHIRAN Inc., May 19, 2022. Available at: `https://www.jftc.go.jp/en/pressreleases/yearly-2022/May/220519.html`, accessed on July 17, 2025.

[13] PSKS Inc. v. Leegin Creative Leather Products Inc., 615 F.3d 412 (5th Cir. 2010). Available at: `https://www.courtlistener.com/opinion/173496/psks-inc-v-leegin-creative-leather-products/`, accessed on July 17, 2025.

[14] Maryland v. Johnson & Johnson Vision Care Inc., No. 02-C-16-002271 (Cir. Ct. Baltimore Cty. Feb. 29, 2016), settlement announced March 30, 2017. See Maryland Daily Record, "Md. settles price-fixing lawsuit against Johnson & Johnson" (March 30, 2017), available at: `https://thedailyrecord.com/2017/03/30/attorney-general-frosh-johnson-and-johnson/`, accessed on July 17, 2025.

[15] In re Online Travel Co. (OTC) Hotel Booking Antitrust Litigation, No. 3:12-CV-3515-B, 2014 WL 5460450 (N.D. Tex. Oct. 27, 2014). For analysis, see K&L Gates, "No Rule of Reason Here: State AG Reminds Manufacturers That Minimum Retail Price Agreements Are Illegal Per Se in Maryland with Latest Suit" (March 21, 2016), available at: `https://www.klgates.com/` (search for article title), accessed on July 17, 2025.

[16] New York et al. v. Herman Miller, Inc., No. 08 CV 02977 (S.D.N.Y. 2008). See National Association of Attorneys General, multistate case information available at: `https://www.naag.org/multistate-case/new-york-et-al-v-herman-miller-inc/`, accessed on July 17, 2025.

including Germany's Buchpreisbindungsgesetz[17] and France's Lang Law.[18]

What makes Japanese publishing especially useful for this paper is that the observed institution bundles together both a ceiling and a floor on retail prices. The counterfactual analysis can therefore decompose fixed RPM into its maximum-RPM and minimum-RPM components and evaluate their welfare implications separately. That decomposition matches the legal question emphasized above: even if some form of RPM can be justified in principle, does the minimum-RPM component improve consumer welfare in a market where one possible procompetitive justification is the inventory channel under demand uncertainty?

3 Publishing Industry in Japan

3.1 Industry Background

Japan is a useful setting for studying RPM in physical books because the market is large, the institutional rules are stable, and brick-and-mortar retailing still matters quantitatively during our sample period. Total book sales amounted to 742 billion yen[19], and average per capita expenditure on books was 5,838 yen in 2015 (Japan Book Publishers Association, 2017)[20]. Although online bookstores were growing, traditional brick-and-mortar bookstores still accounted for 64.6% of total book and magazine sales in 2015, while online sales accounted for only 9.6%.[21] E-book sales were estimated at 150.2 billion yen in 2015, with e-comics comprising 76.5% of that total. This paper focuses exclusively on physical books, and henceforth the term "book" refers only to physical copies.

Three features of market structure are especially important for the analysis that follows. First, retail competition is local rather than national: publisher concentration is limited, whereas wholesaling is concentrated, and bookstore competition varies substantially across counties. In our data, 43% of counties have only one bookstore. This geographic variation motivates the model's emphasis on downstream competition at the county level. Second, newly published books account for most sales: books published within one month account for 20% of sales, and books published within six months account for 67%. This motivates our focus on the months immediately after publication, when pricing and delivery choices are most consequential. Third, literature is the single largest genre, accounting for 39% of sales, which helps motivate the literature-focused estimation

[17] Buchpreisbindungsgesetz vom 2. September 2002 (BGBl. I S. 3448), available at: `https://www.gesetze-im-internet.de/buchprg/BJNR344810002.html`, accessed on July 17, 2025.

[18] Loi n° 81-766 du 10 août 1981 relative au prix du livre, available at: `https://www.legifrance.gouv.fr/loda/id/JORFTEXT000000517179/`, accessed on July 17, 2025.

[19] Roughly 7 billion US dollars.

[20] We focus on the period 2015-2017, which corresponds to the coverage of our dataset.

[21] Other channels include convenience stores, direct sales to libraries, and sales at station kiosks.

sample introduced later.

The first distinctive institutional feature is that publishers set the retail prices of new physical books under RPM. Consumers therefore pay the same price for a given title regardless of whether they purchase it from a neighborhood bookstore, a large chain, or an online retailer. RPM is enforced through two price-maintenance agreements, one between publishers and wholesalers and another between wholesalers and bookstores. RPM does not apply to e-books or used books. For our purposes, the key implication is that the relevant counterfactual is not whether publishers can coordinate retail prices under RPM, but what happens when ex post pricing authority shifts away from the publisher.

The second distinctive institutional feature is that bookstores have limited control over inventory and assortment. Book wholesalers determine the allocation of new titles to bookstores with which they have exclusive supply contracts.[22] Wholesalers bear the delivery costs from publishers as well as warehousing and delivery costs to bookstores. Bookstores earn revenue by selling the assigned titles, typically retaining a margin of 22% of the fixed retail price, wholesalers receive 8%, and the remaining 70% goes to publishers and authors (Japan Book Publishers Association, 2017). Bookstores may return unsold books to their exclusive wholesalers within a specified period, usually six months, at the original price.[23] For the model, the important implication is that the publisher-wholesaler side effectively controls both price and delivery under the status quo, while the return policy sharply reduces retailers' inventory risk. The next subsection therefore focuses on demand uncertainty, returns, and stockouts as the empirical margins most relevant for the model.

These institutional features are not independent; they are linked by RPM itself. Enforcing RPM in this industry is easier when unsold copies can be returned, because retailers can comply with the fixed price without bearing the full risk of overstocking. In turn, a large-scale return system is easier to operate when transactions are organized through revenue sharing at a fixed percentage of the retail price and through a centralized wholesaling network that handles delivery, warehousing, and reverse logistics. This linkage between returns and vertical price control is also emphasized in the earlier theory literature (Marvel and Peck, 1995; Flath and Nariu, n.d.). The relevant economic environment is therefore not a plain wholesale model with decentralized inventory decisions, but a broader institutional bundle in which RPM, returns, revenue sharing, and centralized distribution reinforce one another. That bundled environment is the object of analysis in publishing. At the same time, the main mechanism we study, namely that RPM can soften ex post price competition and thereby affect inventory incentives under demand uncertainty, does not depend on every detail

[22] Because the wholesaler can effectively force bookstores to accept bundled deliveries, one possible interpretation is that the contract serves a strategic role similar to full-line forcing in the movie-rental industry (Ho et al., 2012). We abstract from that possibility because it is not explicit in the contract and because adding that margin to the current model would substantially complicate the analysis.

[23] In practice, wholesalers often accept returns even after the specified period.

of the baseline institution.

3.2 Wholesaler Data

We obtain data on brick-and-mortar bookstores from one of the largest book wholesalers in Japan. The dataset includes records of the number of deliveries and returns at the title, bookstore, and month level, covering the period from 2015 to 2017. For bookstores equipped with point-of-sale (POS) systems, we also observe sales volume. These POS bookstores account for 90% of the total net deliveries (i.e., deliveries minus returns). We treat bookstores that transact with other large book wholesalers and non-POS bookstores as the outside option.[24]

We observe publisher, author, price, category, and publication date at the title level. Prices are set by publishers and are uniform across bookstores for a given title because of RPM. The category variable is obtained from the Open BD project[25], which provides a detailed classification by genre, edition, and target audience. We also construct inventory data from transaction records and impute sales for missing observations. Appendix A provides the details of the category data and the data construction procedure. The resulting dataset is a detailed title-bookstore-month panel.

3.3 Demand Uncertainty and Inventory Control

Using the full nationwide random sample, we provide suggestive evidence on three points that motivate the empirical model. First, publishers and wholesalers face substantial demand uncertainty in the first few months after a title's publication. Second, conditional on that uncertainty, inventory outcomes are much more likely to involve overstock than stockout. Third, the uncertainty resolves relatively quickly once the book is published. The patterns documented in this section motivate the model's focus on ex ante inventory choice under uncertain demand and help clarify which inventory margin is empirically most relevant in this market.

First, the regular and substantial volume of book returns provides strong evidence that the supply side faces ex ante demand uncertainty. Overstock by itself need not imply uncertainty: even under perfect foresight or rational expectations, a seller may optimally carry excess inventory because of delivery frictions, fixed ordering costs, or other nonconvexities. By contrast, copies that are eventually returned for scrap are much harder to reconcile with perfect foresight of realized demand at the maintained price. If publishers and wholesalers had such foresight, they would have little reason to deliver copies that are never sold and are ultimately destroyed. According to the dominant wholesaler Tohan, industry-wide book return rates have remained around 35–40%,

[24] The behaviors of the bookstores in the outside option are held fixed throughout counterfactual simulations.

[25] See `https://openbd.jp/`

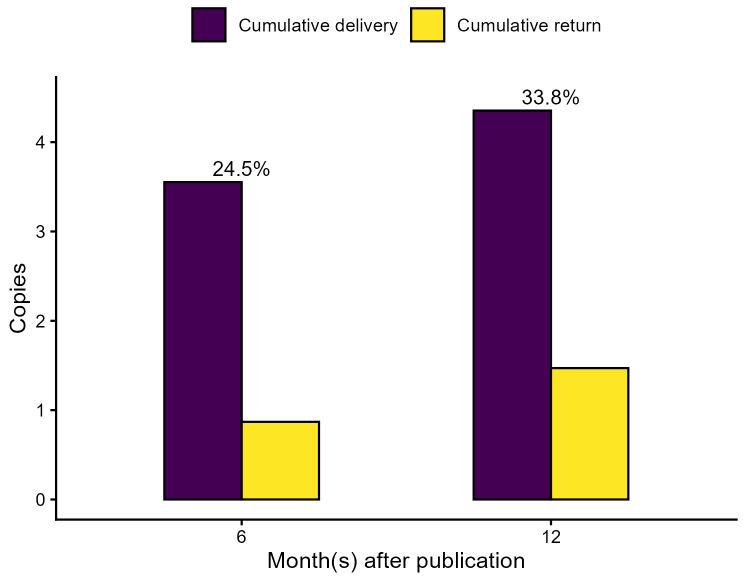
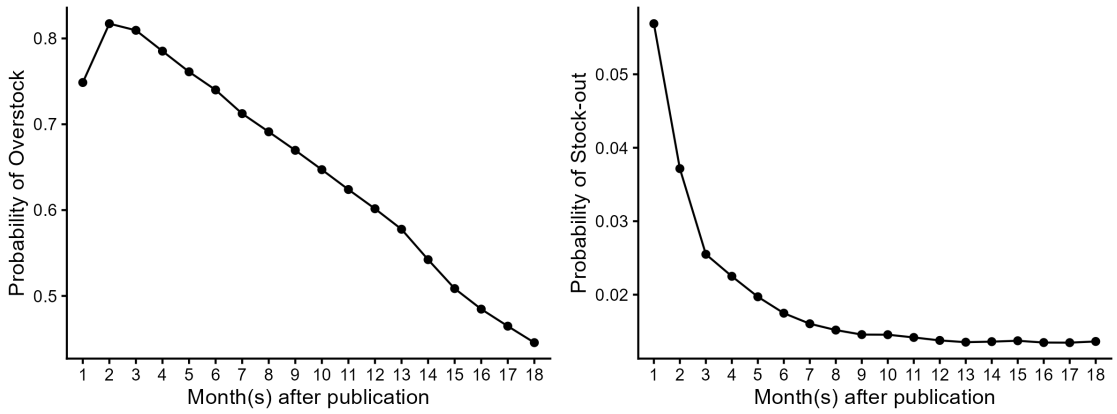


Figure 1: Book Delivery and Return

Note: the ratio on top of the purple bar is the return ratio, defined as the cumulative return divided by the cumulative delivery.



(a) Overstock (b) Stockout

Figure 2: Book Overstock and Stockout

Note: The overstock probability is defined as the percentage of title-bookstore pairs where sales are less than initial inventory of the month. The stockout probability is defined as the percentage of title-bookstore pairs that sell through their initial inventory of the month, conditional on having positive inventory.

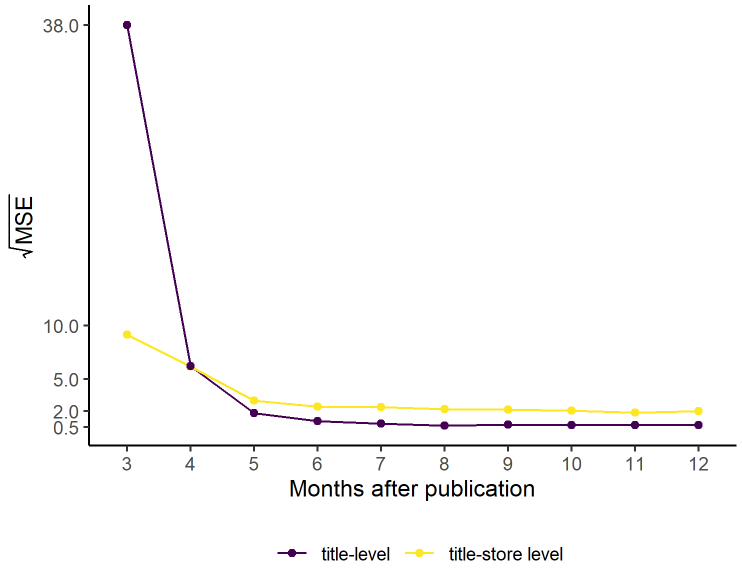


Figure 3: The accuracy of one-month ahead prediction of sales

Note: the y axis is the average square root of the mean square error across books, normalized again the average sales of the book, where the error is defined as the difference between the actual sales at month $t+1$ and the predicted sales at month $t+1$ based on month t sales.

with the most recent figure being 36.5% in 2023, as reported in their financial statement.[26] This is consistent with the nationwide data, as shown in Figure 1: up to the sixth month after a book's publication, on average, 3.55 copies are delivered and 0.87 copies are returned, leading to a return ratio of 24.5%. This ratio increases to 33.8% by the end of the twelfth month.

Second, given that demand is uncertain, the realized inventory outcomes are much more likely to involve overstock than stockout. Figure 2 shows that 81.7% of title–bookstore pairs are overstocked (defined as selling fewer copies than the initial inventory) in the second month. In contrast, stockouts are rare: only 5.7% of title–bookstore pairs sell through their initial inventory in the first month, and this rate falls to 1.4% by the twelfth month. Thus, the relevant empirical question is not whether uncertainty exists, but how it maps into inventory outcomes. In this market, uncertainty is associated primarily with overstock rather than frequent stockout.[27]

Lastly, the demand uncertainty, measured by the root mean squared prediction error of month $t+1$ sales based on month t sales, is highest in the first three months after publication and declines rapidly thereafter. After six months, the uncertainty becomes relatively stable at both the title and title-store levels. Figure 3 plots the root mean squared prediction errors, normalized by the mean

[26] See `https://www.tohan.jp/wp/wp-content/themes/tohan/pdf/20240531_kessan2023.pdf`, accessed on August 7, 2025.

[27] The overstock and stockout probabilities are calculated conditional on having positive inventory at the beginning of each month after publication. These two values do not sum to one because the remaining probability corresponds to cases where the initial inventory—and thus sales—is zero.

sales. The mean is estimated from a Poisson model: $\log(\mathbb{E}[\text{sales}_{jbt}]) = \alpha_j + \beta_j t$, where j indexes the title, b the bookstore, and t the number of months since publication. Because the wholesaler observes the total sales and shares the information across bookstores and publishers, we expect demand uncertainty to resolve in a similar fashion across the distribution chain.

4 A Pro-Competitive Justification

This section develops a stylized analytical model that isolates the inventory channel behind the procompetitive justification for minimum RPM. Its purpose is conceptual rather than quantitative. The model builds on the core insight of Deneckere et al. (1996, 1997); Butz (1997); Flath and Nariu (n.d.) and Krishnan and Winter (2007) that, under demand uncertainty, ex post downstream price competition can weaken ex ante inventory incentives and that a minimum price floor can partly offset that distortion. Our setting differs from these theoretical papers in three respects. First, inventory is centralized upstream to match the publishing institution, rather than chosen by downstream retailers. Second, we focus on the minimum-RPM channel rather than the broader contract-design problem. Third, we abstract from stockout-based substitution across bookstores. Even with these simplifications, the core insight is preserved: if low-demand-state price competition reduces the expected return to carrying inventory, then a minimum price floor can still raise ex ante inventory incentives even when the inventory decision is centralized upstream. The section is not intended to reproduce every institutional detail of the structural model used later in the paper.

We use a simple linear-demand model in which downstream retailers compete in prices under differentiated Bertrand competition. The analytical model therefore follows the same competitive protocol as the empirical model. The simplification is that retailers are symmetric and horizontal differentiation takes a reduced-form linear form that preserves positive markups and yields a closed-form equilibrium. The key mechanism does not rely on any special feature of this linear specification. What matters is that stronger downstream competition lowers the ex post payoff from carrying additional inventory in some demand states, especially intermediate-demand states in which extra inventory mainly intensifies price competition and erodes inframarginal margins rather than creating valuable additional sales.

Inventory must be chosen before demand is realized, while sales and prices are determined after demand is realized. This timing creates the basic trade-off that drives the rest of the paper: ex ante inventory is costly, but ex post it acts like capacity. Whether firms want to hold more inventory therefore depends on how profitable extra sales are in the realized demand state.

We assume uncertain demand in a market-size shock form. There are $N \geq 2$ symmetric retail-

ers, and ex post demand for retailer i is

$$q_i(\theta, \boldsymbol{p}) = \theta - p_i + \gamma \sum_{k \neq i} p_k, \tag{1}$$

where $\theta \sim F$ has support $[1, \bar{\theta}]$ with $\bar{\theta} > 1$ and density f, and $\gamma \in \left(0, \frac{1}{N-1}\right)$ indexes substitutability across retailers. A higher γ means that retailers are closer substitutes, so downstream competition is more intense.

We study two benchmark institutions, vertical integration and a decentralized revenue-sharing model, to isolate how downstream competition distorts inventory incentives relative to the integrated-industry benchmark.

Vertical integration A monopolist wholesaler is integrated with downstream retailing. The integrated firm chooses per-store inventory $d \geq 0$ before θ is realized, at per-unit cost c. After θ is realized, it chooses a common retail price across stores to maximize total revenue subject to the per-store capacity constraint.

Revenue-sharing with downstream competition A monopolist wholesaler supplies N identical retailers. Before θ is realized, the wholesaler chooses inventories $d_i \geq 0$ for each retailer at per-unit cost c. After θ is realized, retailers compete in prices subject to inventory constraints:

$$0 \leq q_i(\theta, \boldsymbol{p}) \leq d_i.$$

Revenues are shared: the wholesaler receives fraction $\rho \in [0, 1]$ of realized retail revenue, and retailers collectively receive $1 - \rho$. Equivalently, each retailer's objective is $(1 - \rho)p_i q_i$, which scales payoffs but does not change the Bertrand best responses. The parameter ρ matters for the wholesaler's inventory incentives because it determines how much of downstream revenue the wholesaler internalizes.

4.1 Inventory Distortion under Downstream Competition

Under both institutions, inventory is chosen before θ is realized, so the ex post outcome is characterized by a threshold $\theta_0(d)$ that splits the state space into two regions.

Fix a symmetric precommitted inventory level d. When demand is low, firms would like to sell less than d even if they could freely choose prices after observing θ. In that region, extra inventory has zero marginal benefit because it remains unsold while still incurring its marginal cost. When demand is high, the unconstrained optimum exceeds d, inventory binds, and realized sales equal

available inventory. In that region, increasing d raises sales and therefore has a positive marginal benefit.

Under vertical integration, the integrated firm solves

$$\max_{p} \; Np\left(\theta - (1-\gamma(N-1))p\right),$$

which yields the unconstrained monopoly price

$$p^{VI}(\theta) = \frac{\theta}{2(1-\gamma(N-1))}$$

and per-store quantity

$$q^{VI}(\theta) = \frac{\theta}{2}.$$

Thus inventory binds iff $\theta \geq \theta_0^{VI}(d) \equiv 2d$.

Under the revenue-sharing model with N identical Bertrand retailers, the symmetric unconstrained ex-post price solves

$$\max_{p_i} \; p_i \left(\theta - p_i + \gamma \sum_{k \neq i} p_k\right),$$

which yields

$$p^{RS}(\theta) = \frac{\theta}{2-\gamma(N-1)}$$

and per-store quantity

$$q^{RS}(\theta) = \frac{\theta}{2-\gamma(N-1)}.$$

Thus inventory binds iff

$$\theta \geq \theta_0^{RS}(d) \equiv (2-\gamma(N-1))d.$$

Because $\gamma > 0$, we have $\theta_0^{RS}(d) < \theta_0^{VI}(d)$ for any given d: downstream competition makes the inventory-constrained region start at a lower demand state.

We can therefore write the first-order condition for inventory under each institution by equating marginal benefit and marginal cost. Under vertical integration, the optimal inventory d^{VI} satisfies

$$\int_{2d^{VI}}^{\bar{\theta}} \frac{\theta - 2d^{VI}}{1-\gamma(N-1)} \, dF(\theta) = c. \tag{2}$$

Extra inventory is valuable only in states in which the firm is inventory constrained and additional sales still have positive marginal revenue. Under the revenue-sharing symmetric equilibrium, when

inventory binds the market-clearing price is

$$p^c(\theta, d) = \frac{\theta - d}{1 - \gamma(N-1)},$$

so per-store profit is $d(\theta - d)/(1 - \gamma(N-1))$. Therefore, we have

$$\begin{aligned} &\rho \int_{(2-\gamma(N-1))d^{RS}}^{\bar{\theta}} \frac{\theta - 2d^{RS}}{1 - \gamma(N-1)}\, dF(\theta) = c \iff \\ &\underbrace{\int_{(2-\gamma(N-1))d^{RS}}^{2d^{RS}} \frac{\theta - 2d^{RS}}{1 - \gamma(N-1)}\, dF(\theta)}_{\leq 0} + \int_{2d^{RS}}^{\bar{\theta}} \frac{\theta - 2d^{RS}}{1 - \gamma(N-1)}\, dF(\theta) = \frac{c}{\rho}. \end{aligned} \tag{3}$$

Relative to vertical integration, the decentralized revenue-sharing model contains two distinct distortions. The decomposition in Equation (3) makes both visible.

The first distortion is *reduced internalization of revenue*. Because the wholesaler receives only fraction ρ of retail revenue while bearing the full inventory cost, the effective marginal cost on the right-hand side is scaled up from c to c/ρ. This force reduces inventory even if ex post product-market conduct were otherwise unchanged.

The second distortion is *competition-induced erosion of the marginal return to inventory*. The additional non-positive term captures intermediate-demand states in which extra inventory mainly intensifies downstream price competition and depresses inframarginal margins rather than creating valuable additional sales. By contrast, an integrated monopolist internalizes this margin effect and avoids expanding inventory in states where marginal revenue is low. This force becomes stronger when downstream competition is stronger, because a higher γ lowers $\theta_0^{RS}(d)$ and expands the set of states in which additional inventory has low or negative marginal value.

4.2 Effects of Resale Price Maintenance

Because the distortion of interest here comes from downstream competition reducing inventory incentives, and not from double marginalization, the analytically relevant RPM instrument is a minimum price floor rather than a price ceiling. Let $\underline{p} \geq 0$ denote a retail price floor imposed in the revenue-sharing model, so that the ex-post retail price cannot fall below $\underline{p}$. For a given precommitted inventory d, the floor becomes relevant in the constrained region when

$$\frac{\theta - d}{1 - \gamma(N-1)} < \underline{p} \qquad \iff \qquad \theta < d + (1 - \gamma(N-1))\underline{p}. \tag{4}$$

The price floor therefore truncates the low-price region and changes the threshold separating the abundant- and constrained-inventory regions:

$$\theta_0^{RS}(d;\underline{p}) \;\equiv\; \max\left\{(2-\gamma(N-1))d,\; d+(1-\gamma(N-1))\underline{p}\right\}, \tag{5}$$

When the floor is nontrivial, so that $(2-\gamma(N-1))d < d+(1-\gamma(N-1))\underline{p}$, we have $\theta_0^{RS} = d+(1-\gamma(N-1))\underline{p}$. The first-order condition for inventory then becomes

$$\begin{aligned}
&\rho\int_{d^f+(1-\gamma(N-1))\underline{p}}^{\bar\theta}\frac{\theta-2d^f}{1-\gamma(N-1)}\,dF(\theta)=c \iff\\
&\underbrace{\int_{d^f+(1-\gamma(N-1))\underline{p}}^{2d^f}\frac{\theta-2d^f}{1-\gamma(N-1)}\,dF(\theta)}_{\geq\int_{(2-\gamma(N-1))d^f}^{2d^f}\frac{\theta-2d^f}{1-\gamma(N-1)}\,dF(\theta)}+\int_{2d^f}^{\bar\theta}\frac{\theta-2d^f}{1-\gamma(N-1)}\,dF(\theta)=\frac{c}{\rho}.
\end{aligned} \tag{6}$$

Raising $\underline{p}$ increases θ_0^{RS} and therefore shortens, or may eliminate, the interval $\left[\theta_0^{RS},\,2d\right]$ over which marginal revenue is negative. Intuitively, by preventing prices from falling too much in intermediate-demand states, a price floor softens ex post competition and reduces the margin erosion caused by extra inventory. This raises the marginal value of inventory and strengthens delivery incentives.

Effects on consumer surplus The same price floor that raises inventory incentives has an ambiguous effect on consumer surplus. Under linear demand, consumer surplus at retailer i is $\frac{1}{2}q_i^2$. Under symmetry, aggregate consumer surplus is therefore $\frac{N}{2}q^2$. Assuming that the floor is nontrivial and efficient rationing, expected consumer surplus with a price floor is

$$\mathbb{E}\big[CS\mid d,\underline{p}\big]=\int_1^{d+(1-\gamma(N-1))\underline{p}}\frac{N}{2}\left(\theta-(1-\gamma(N-1))\underline{p}\right)^2\,dF(\theta)\;+\;\int_{d+(1-\gamma(N-1))\underline{p}}^{\bar\theta}\frac{N}{2}d^2\,dF(\theta). \tag{7}$$

Compared with the pure revenue-sharing model,

$$\mathbb{E}[CS\mid d]=\int_1^{(2-\gamma(N-1))d}\frac{N}{2}\left(\frac{\theta}{2-\gamma(N-1)}\right)^2\,dF(\theta)\;+\;\int_{(2-\gamma(N-1))d}^{\bar\theta}\frac{N}{2}d^2\,dF(\theta). \tag{8}$$

At a given inventory level d, a nontrivial price floor lowers consumer surplus in low- and intermediate-demand states because it keeps prices above the competitive level. But by increasing equilibrium inventory, it can raise consumer surplus in high-demand states by relaxing inventory constraints and lowering realized prices there. The overall effect is therefore ambiguous and de-

pends on both demand uncertainty and the intensity of downstream competition.

4.3 Channel Decomposition and Comparative-Statics Predictions

Let $cs^f(\theta; d, \underline{p})$ denote the state-contingent consumer surplus under minimum RPM and $cs^0(\theta; d)$ the corresponding object under no RPM:

$$cs^f(\theta; d, \underline{p}) = \begin{cases} \frac{N}{2}\left(\theta - (1-\gamma(N-1))\underline{p}\right)^2, & \theta \leq d + (1-\gamma(N-1))\underline{p}, \\ \frac{N}{2}d^2, & \theta > d + (1-\gamma(N-1))\underline{p}, \end{cases}$$

$$cs^0(\theta; d) = \begin{cases} \frac{N}{2}\left(\frac{\theta}{2-\gamma(N-1)}\right)^2, & \theta \leq (2-\gamma(N-1))d, \\ \frac{N}{2}d^2, & \theta > (2-\gamma(N-1))d. \end{cases}$$

These are the integrands in (7) and (8). Let d^0 denote equilibrium inventory under no RPM, d^f equilibrium inventory under minimum RPM, and let

$$\theta^f \equiv d^f + (1-\gamma(N-1))\underline{p}$$

be the cutoff below which the floor is relevant under minimum RPM. Under the nontrivial-floor case studied above, this is exactly the minimum-RPM stockout cutoff evaluated at the minimum-RPM equilibrium inventory, that is, $\theta^f = \theta_0^{RS}(d^f; \underline{p})$. We then define low-demand states as $\theta \leq \theta^f$ and high-demand states as $\theta > \theta^f$.

The primitive welfare object is the equilibrium comparison

$$\Delta CS = \mathbb{E}[CS \mid d^f, \underline{p}] - \mathbb{E}[CS \mid d^0].$$

This total effect compares minimum RPM at its own equilibrium inventory, d^f, with no RPM at its own equilibrium inventory, d^0. To understand the mechanism behind this equilibrium comparison, we evaluate both pricing rules at the minimum-RPM inventory d^f and then add and subtract $\mathbb{E}[CS \mid d^f]$. This turns the total effect into a price-floor component, holding inventory fixed at d^f, plus an inventory component, holding the no-RPM pricing rule fixed. Using the minimum-RPM inventory as the reference point also makes the state partition align with the first stockout state under minimum RPM, $\theta^f = \theta_0^{RS}(d^f; \underline{p})$.

Using these objects, we divide demand states into two regions, low-demand states $\theta \leq \theta^f$ and high-demand states $\theta > \theta^f$, and decompose the total effect within each region into a direct price-cost component and an availability component. The four terms below therefore do not represent four mutually exclusive demand-state classifications. Rather, they are two state regions (L and H)

crossed with two mechanisms (P and A):

$$\Delta CS = \Delta CS^{L,P} + \Delta CS^{L,A} + \Delta CS^{H,P} + \Delta CS^{H,A},$$

where

$$\begin{aligned}
\Delta CS^{L,P} &\equiv \int_{1}^{\theta^f} \left[cs^f(\theta; d^f, \underline{p}) - cs^0(\theta; d^f)\right] dF(\theta), \\
\Delta CS^{L,A} &\equiv \int_{1}^{\theta^f} \left[cs^0(\theta; d^f) - cs^0(\theta; d^0)\right] dF(\theta), \\
\Delta CS^{H,P} &\equiv \int_{\theta^f}^{\bar{\theta}} \left[cs^f(\theta; d^f, \underline{p}) - cs^0(\theta; d^f)\right] dF(\theta), \\
\Delta CS^{H,A} &\equiv \int_{\theta^f}^{\bar{\theta}} \left[cs^0(\theta; d^f) - cs^0(\theta; d^0)\right] dF(\theta).
\end{aligned}$$

Here, L and H denote the two demand-state regions defined by the minimum-RPM cutoff θ^f, while P and A denote the two channels within each region: the direct price-cost effect and the inventory-availability effect. This decomposition is defined entirely in terms of the state-contingent surplus objects already derived above.

In the stylized model, two of these four mechanism-specific terms are negligible. In low-demand states, inventory is slack under d^f, so the availability effect $\Delta CS^{L,A}$ is approximately zero. In high-demand states, the floor is no longer the first-order distortion, so the direct price-cost term $\Delta CS^{H,P}$ is also approximately zero. The dominant terms are therefore the low-demand-state price-cost term, $\Delta CS^{L,P} < 0$, and the high-demand-state availability term, $\Delta CS^{H,A} > 0$. It is then useful to re-aggregate the decomposition back to the two demand-state channels

$$\Delta CS^{L} = \Delta CS^{L,P} + \Delta CS^{L,A}, \qquad \Delta CS^{H} = \Delta CS^{H,P} + \Delta CS^{H,A},$$

so that, in the stylized model, $\Delta CS^{L} \approx \Delta CS^{L,P} < 0$ and $\Delta CS^{H} \approx \Delta CS^{H,A} > 0$.

This decomposition yields comparative-static predictions even though it does not determine the sign of the net effect on its own. Stronger downstream competition, captured here by a higher γ, tends to strengthen the high-state gain by making the no-RPM inventory distortion more severe. Higher demand uncertainty expands the range of low- and intermediate-demand states in which the floor distorts price. Consumer surplus can therefore rise under minimum RPM only when the high-state gain dominates the low-state loss. This condition is most likely when downstream competition is strong enough to distort inventory incentives but uncertainty is not so high that the floor binds over too wide a range of states. These are the same comparative-static insights that guide the empirical analysis.

This stylized model provides the conceptual framework for the structural analysis that follows. The objects here map naturally into the empirical model: θ plays the role of the realized title-level demand state, d corresponds to ex ante delivery or inventory, and the ex post competition stage maps directly to the retail-pricing stage in the counterfactual analysis. What the structural model adds is richer demand heterogeneity across bookstores and counties, publisher belief formation, and store-level truncation of realized demand by inventory. The structural model is then used to quantify the low- and high-state channels in the actual publishing environment.

5 An Empirical Model of the Publishing Industry

This section presents the empirical model used in the estimation and counterfactual analysis. For each title, the upstream firm first observes a noisy signal about demand, then chooses a uniform retail price and bookstore-level deliveries for the first six months after publication. Consumers subsequently decide which bookstore to visit, and realized sales are determined by the interaction between store choice and inventory availability. We begin with demand and then turn to the upstream firm's problem.

5.1 Demand

We assume that books are highly differentiated and that substitution across titles is negligible, so we model demand for one title at a time. A market is defined as book j in county l over the first six months after publication. In each market, a mass M of consumers has unit demand and chooses one bookstore in the county or the outside option. A consumer can purchase the book only if the chosen bookstore has positive inventory; otherwise, the consumer leaves without buying. Appendix B shows that stockout-based substitution is limited in our data, so we abstract from consumers searching across bookstores after a stockout. We therefore distinguish between *potential demand*, which captures store choice before inventory constraints are imposed, and *realized demand*, which captures sales after inventory constraints are imposed.

We restrict attention to the first six months after publication for three reasons. First, most sales occur within this window, and the evidence in Section 3.3 shows that demand uncertainty declines substantially over it. Second, Appendix B shows little evidence that stockouts in the first six months shift demand to later months. Third, inventory holding costs are limited because bookstores can return unsold copies. A dynamic model would therefore add substantial complexity while offering limited additional insight for the question we study.

Potential Demand We model potential demand using a nested logit with two nests: an inside nest containing all bookstores in county l, and an outside nest containing the no-purchase option. The nesting parameter ω governs substitution within the inside nest relative to substitution between the inside and outside nests. A higher ω implies that bookstores are closer substitutes for one another within the county market and that downstream competition is more intense.

Let $\mathcal{B}_l$ denote the set of bookstores in county $l \in \mathcal{L}$. For a bookstore option $b \in \mathcal{B}_l$, the indirect utility of consumer i from purchasing book j at bookstore b in county l is

$$\begin{aligned} u_{ijbl} &\coloneqq x_j'\beta + \alpha p_{jbl} + x_{jbl}'\psi + \xi_j + \zeta_{ij,\text{in}} + (1-\omega)\epsilon_{ijbl}, \\ &\coloneqq \delta_{jbl} + \zeta_{ij,\text{in}} + (1-\omega)\epsilon_{ijbl}, \end{aligned} \tag{9}$$

where x_j and x_{jbl} are vectors of observed exogenous characteristics of the book and bookstore b in county l, respectively.[28] Here, p_{jbl} is the retail price of book j at bookstore b in county l, and ξ_j is an unobserved (to the econometrician) title-specific demand shifter. The term $\zeta_{ij,\text{in}} + (1-\omega)\epsilon_{ijbl}$ follows the generalized extreme-value structure that yields the nested logit model with nesting parameter ω (Cardell, 1997). We normalize the mean utility of the outside option to zero.

Let $\boldsymbol{p}_{jl} \coloneqq [p_{jbl}]'_{b\in\mathcal{B}_l}$ denote the vector of bookstore prices for book j in county l. Then the probability that a consumer selects bookstore b is

$$q_{jbl}(\boldsymbol{p}_{jl}, \xi_j) \coloneqq \frac{\exp(\frac{\delta_{jbl}}{1-\omega})}{\left[\sum_{d\in\mathcal{B}_l}\exp(\frac{\delta_{jdl}}{1-\omega})\right]^{\omega}\left\{1+\left[\sum_{d\in\mathcal{B}_l}\exp(\frac{\delta_{jdl}}{1-\omega})\right]^{(1-\omega)}\right\}}, \tag{10}$$

and the probability of choosing the outside option is

$$q_{j0l}(\boldsymbol{p}_{jl}, \xi_j) \coloneqq \frac{1}{1+\left[\sum_{d\in\mathcal{B}_l}\exp(\frac{\delta_{jdl}}{1-\omega})\right]^{(1-\omega)}}. \tag{11}$$

Potential demand at bookstore b is therefore $Mq_{jbl}(\boldsymbol{p}_{jl}, \xi_j)$.

Realized Demand Realized demand equals potential demand truncated by inventory. Let n_{jbl} denote the inventory of book j at bookstore b in county l. Then realized demand, or sales, is

$$\tilde{Q}_{jbl}(\boldsymbol{p}_{jl}, n_{jbl}, \xi_j) \coloneqq \min\{n_{jbl}, Mq_{jbl}(\boldsymbol{p}_{jl}, \xi_j)\}, \tag{12}$$

[28] In estimation, we include publisher dummies, year-of-publication dummies, and a new-author dummy in x_j. We include the total number of books in the bookstore, county-level employment, and the county-level total number of libraries in x_{jbl}.

which implies that stockouts matter only through the inventory constraint. Under the current RPM regime, the publisher sets a uniform retail price for title j across all bookstores and counties, so $p_{jbl} = p_j$ for all $b \in \mathcal{B}_l$ and $l \in \mathcal{L}$. We therefore write the supply-side pricing decision in terms of the scalar p_j.

Consumer Surplus Our welfare object must account for the possibility that a consumer visits a bookstore but cannot buy because the title is out of stock. Under the maintained assumptions of no further search after a stockout and random rationing among consumers who visit the same bookstore, we use a stockout-adjusted nested-logit surplus index. The adjustment scales the contribution of each bookstore option by the probability that a consumer who visits that bookstore is actually served. Under random rationing, this service probability is

$$\mathbb{P}(i \text{ can buy from } b) = \min\left\{1, \frac{n_{jbl}}{M q_{jbl}(\boldsymbol{p}_{jl}, \xi_j)}\right\}.$$

For the outside option, we normalize mean utility to zero and set the service probability equal to one. Using this adjustment, the expected consumer surplus for book j in county l can be written as

$$-\alpha^{-1} \sum_{b \in \{0\} \cup \mathcal{B}_l} \mathbb{P}(i \text{ chooses option } b)\mathbb{E}_\epsilon(u_{ijbl} \mid i \text{ chooses option } b)\mathbb{P}(i \text{ can buy from } b). \tag{13}$$

Equivalently, this object equals the standard nested-logit surplus contribution from each option multiplied by the service probability at that option. It is therefore the relevant consumer-surplus index under our maintained stockout assumptions, rather than the frictionless nested-logit surplus that would apply if consumers could always purchase after their initial visit. Appendix C derives the corresponding closed-form expression in our notation.

5.2 Supply

Given the structure of the Japanese publishing industry, we treat the publisher and wholesaler as a single upstream decision-maker. Under the status quo, this upstream firm chooses the retail price and the delivery to each bookstore, while bookstores passively sell the delivered units at the mandated price. We aggregate all supply decisions made during the first six months after publication into a single period.

Inventory at bookstore b is therefore the number of copies delivered to that bookstore, so we use inventory and delivery interchangeably. Retail revenue is split according to the fixed industry shares: the upstream firm receives share ρ and bookstores receive share $1 - \rho$. Because bookstores do not choose prices or deliveries under the status quo, the relevant optimization problem is entirely

the upstream firm's problem.

Demand uncertainty The upstream firm faces uncertainty about title-level demand, summarized by the demand shifter ξ_j. Before observing any title-specific information, it holds the prior belief

$$\xi_j \sim N(\mu_{0,\xi}, \sigma^2_{0,\xi}).$$

Before setting price and deliveries, it observes a noisy signal

$$\xi^*_j = \xi_j + \sigma_{j,\xi} e_j, \qquad e_j \sim N(0,1),$$

so that

$$\xi^*_j \mid \xi_j \sim N(\xi_j, \sigma^2_{j,\xi}).$$

The signal noise varies across books according to

$$\sigma_{j,\xi} = \sigma_\xi + \beta_\xi \cdot NewAuthor_j,$$

where $NewAuthor_j$ is an indicator equal to one if book j is written by a new author (no publication before book j) and zero otherwise. This specification captures the idea that the upstream firm has less precise information about demand for books by new authors.

After observing ξ^*_j, the upstream firm updates its belief about ξ_j using Bayes' rule. The posterior distribution remains normal:

$$\xi_j \mid \xi^*_j \sim N\left(\mu^{post}_{j,\xi},\, \sigma^{2,post}_{j,\xi}\right),$$

with posterior variance

$$\sigma^{2,post}_{j,\xi} = \frac{\sigma^2_{0,\xi}\sigma^2_{j,\xi}}{\sigma^2_{0,\xi} + \sigma^2_{j,\xi}},$$

and posterior mean

$$\mu^{post}_{j,\xi} = \frac{\sigma^2_{j,\xi}}{\sigma^2_{0,\xi} + \sigma^2_{j,\xi}}\mu_{0,\xi} + \frac{\sigma^2_{0,\xi}}{\sigma^2_{0,\xi} + \sigma^2_{j,\xi}}\xi^*_j.$$

The posterior variance $\sigma^{2,post}_{j,\xi}$ captures the residual demand uncertainty at the time price and delivery are chosen. It is therefore a key object governing the strength of the inventory channel through which RPM affects outcomes.

Marginal cost of production We assume that the marginal cost of printing and delivering each copy is

$$MC_{jbl} = \lambda' w_{jbl} + \varepsilon_{jbl}, \tag{14}$$

where w_{jbl} is a subset of exogenous characteristics drawn from x_j and x_{jbl}[29], and ε_{jbl} is a bookstore-level marginal-cost shock distributed as $\mathcal{N}(0, \sigma_\varepsilon)$.

Although one might consider fixed costs of delivery and return, we treat them as negligible at the title and period level. This is because deliveries and returns between bookstores and the wholesaler's warehouse typically involve consolidated shipments of multiple titles. Similarly, storage costs arising from the shadow price of the bookstore capacity constraint are also negligible near the equilibrium level of storage, as bookstores can always clear space by returning unsold books. While the fixed costs of creating a new book (e.g., plate fees and editorial labor) are non-negligible, we abstract from this margin by conditioning on the set of books observed in the data.

5.2.1 Publisher's Problem

After observing the signal but not the true realization of ξ_j, the upstream firm chooses the uniform retail price p_j and the vector of bookstore-level deliveries $\boldsymbol{d}_j \coloneqq [d_{jbl}]'_{b\in\mathcal{B}_l,\, l\in\mathcal{L}}$, where delivery equals inventory, $d_{jbl} = n_{jbl}$. Its problem is to maximize expected upstream profit:

$$\max_{p_j, \mathbf{d}_j} \quad \mathbb{E}_{\xi|\xi_j^*}\left\{\sum_{l\in\mathcal{L}}\sum_{b\in\mathcal{B}_l}\left[\rho p_j \tilde{Q}_{jbl}(p_j, d_{jbl}, \xi_j) - MC_{jbl} d_{jbl}\right]\right\}, \tag{15}$$

subject to

$$p_j \geq 0,$$
$$d_{jbl} \geq 0$$

5.2.2 Optimal Decisions

Pricing decision Fix deliveries $\{d_{jbl}\}_{l,b}$ and define expected upstream revenue net of delivery cost as

$$\Pi_j(p_j; \mathbf{d}_j, \xi_j^*) \coloneqq \mathbb{E}_{\xi|\xi_j^*}\left\{\sum_{l\in\mathcal{L}}\sum_{b\in\mathcal{B}_l} \rho p_j \min\{d_{jbl}, Mq_{jbl}(p_j, \xi_j)\}\right\}. \tag{16}$$

The cost term does not depend on p_j, so the pricing problem is $\max_{p_j\geq 0}\Pi_j(p_j; \mathbf{d}_j, \xi_j^*)$. For a given bookstore-state pair (b, l, ξ_j), let

$$m_{jbl}(p_j, \xi_j) \coloneqq Mq_{jbl}(p_j, \xi_j).$$

[29] In estimation, publisher dummies are incorporated into w_{jbl}. Consequently, in the implementation we employ w_j.

For later use, define the stockout indicator as

$$I_{jbl}^{Stockout}(p_j, d_{jbl}, \xi_j) \coloneqq \mathbb{1}_{\{m_{jbl}(p_j,\xi_j) > d_{jbl}\}}.$$

Because $\tilde{Q}_{jbl}(p_j, d_{jbl}, \xi_j) = \min\{d_{jbl}, m_{jbl}(p_j, \xi_j)\}$, the contribution of bookstore b in county l to expected revenue is

$$g_{jbl}(p_j, \xi_j) \coloneqq \rho p_j \min\{d_{jbl}, m_{jbl}(p_j, \xi_j)\}.$$

Hence

$$g_{jbl}(p_j, \xi_j) = \begin{cases} \rho p_j d_{jbl}, & \text{if } m_{jbl}(p_j, \xi_j) > d_{jbl}, \\ \rho p_j m_{jbl}(p_j, \xi_j), & \text{if } m_{jbl}(p_j, \xi_j) < d_{jbl}. \end{cases}$$

Therefore, whenever $m_{jbl}(p_j, \xi_j) \neq d_{jbl}$, the derivative is

$$\frac{\partial g_{jbl}(p_j, \xi_j)}{\partial p_j} = \begin{cases} \rho d_{jbl}, & \text{if } m_{jbl}(p_j, \xi_j) > d_{jbl}, \\ \rho \left[m_{jbl}(p_j, \xi_j) + p_j \frac{\partial m_{jbl}(p_j, \xi_j)}{\partial p_j} \right], & \text{if } m_{jbl}(p_j, \xi_j) < d_{jbl}. \end{cases}$$

At a kink, where $m_{jbl}(p_j, \xi_j) = d_{jbl}$, the left and right derivatives differ:

$$\frac{\partial^- g_{jbl}(p_j, \xi_j)}{\partial p_j} = \rho d_{jbl}, \qquad \frac{\partial^+ g_{jbl}(p_j, \xi_j)}{\partial p_j} = \rho \left[d_{jbl} + p_j \frac{\partial m_{jbl}(p_j, \xi_j)}{\partial p_j} \right].$$

Aggregating across bookstores and counties and taking posterior expectations yields the left and right derivatives of Π_j:

$$\begin{aligned} \Pi_j^{-\prime}(p_j; \mathbf{d}_j, \xi_j^*) = \mathbb{E}_{\xi|\xi_j^*} \Bigg[\sum_{l \in \mathcal{L}} \sum_{b \in \mathcal{B}_l} \rho \Bigg(& \mathbb{1}_{\{m_{jbl}(p_j,\xi_j) \geq d_{jbl}\}} d_{jbl} \\ & + \mathbb{1}_{\{m_{jbl}(p_j,\xi_j) < d_{jbl}\}} \left[m_{jbl}(p_j, \xi_j) + p_j \frac{\partial m_{jbl}(p_j, \xi_j)}{\partial p_j} \right] \Bigg) \Bigg], \end{aligned} \tag{17}$$

$$\begin{aligned} \Pi_j^{+\prime}(p_j; \mathbf{d}_j, \xi_j^*) = \mathbb{E}_{\xi|\xi_j^*} \Bigg[\sum_{l \in \mathcal{L}} \sum_{b \in \mathcal{B}_l} \rho \Bigg(& \mathbb{1}_{\{m_{jbl}(p_j,\xi_j) > d_{jbl}\}} d_{jbl} \\ & + \mathbb{1}_{\{m_{jbl}(p_j,\xi_j) \leq d_{jbl}\}} \left[m_{jbl}(p_j, \xi_j) + p_j \frac{\partial m_{jbl}(p_j, \xi_j)}{\partial p_j} \right] \Bigg) \Bigg]. \end{aligned} \tag{18}$$

A necessary condition for an interior local optimum $p_j^* > 0$ is therefore the subgradient condition

$$\Pi_j^{-\prime}(p_j^*; \mathbf{d}_j, \xi_j^*) \geq 0 \qquad \text{and} \qquad \Pi_j^{+\prime}(p_j^*; \mathbf{d}_j, \xi_j^*) \leq 0. \tag{19}$$

Under concavity of $\Pi_j(\cdot; \mathbf{d}_j, \xi_j^*)$, the same condition is also sufficient. If no bookstore-state pair is exactly at the kink with positive posterior probability, then $\Pi_j^{-\prime}(p_j^*; \mathbf{d}_j, \xi_j^*) = \Pi_j^{+\prime}(p_j^*; \mathbf{d}_j, \xi_j^*)$, and

the condition reduces to the usual equality $\Pi_j'(p_j^*; \mathbf{d}_j, \xi_j^*) = 0$. In the implementation, we solve this kinked pricing problem store by store within the pricing best-response step by searching over two candidate regions: an unrestricted interior optimum within the smooth region and a boundary point at which $m_{jbl}(p_j, \xi_j) = d_{jbl}$ for the focal bookstore-state pair. We then select the candidate with the higher objective value. Economically, when a bookstore stocks out, realized sales are locally fixed at delivery and a higher price raises revenue on inframarginal units. When it does not stock out, the standard marginal-revenue tradeoff applies. The kink arises because price changes can move bookstore-state pairs between these two regimes.

Delivery decision Holding the retail price fixed, the derivative of upstream profit with respect to delivery to bookstore b is nonzero only in demand states in which the bookstore would otherwise stock out. The optimal delivery d_{jbl}^* therefore satisfies

$$\mathbb{E}_{\xi|\xi_j^*}\left\{I_{jbl}^{Stockout}(p_j, d_{jbl}^*, \xi_j)(\rho p_j) - (\lambda' w_{jbl} + \varepsilon_{jbl})\right\} = 0. \tag{20}$$

The condition equates the marginal cost of sending one more copy to bookstore b with its marginal expected benefit. An additional copy is valuable only when it relaxes a stockout, in which case it generates upstream revenue equal to the upstream firm's revenue share, ρp_j.

For any fixed price p_j, the upstream profit function is concave in each delivery choice because $\min\{d_{jbl}, Mq_{jbl}(p_j, \xi_j)\}$ is concave in d_{jbl} state by state, and expectation preserves concavity. Under the continuous posterior distribution implied by the normal prior and signal structure, the marginal benefit of delivery decreases continuously with d_{jbl}, so the optimal delivery is unique. By contrast, uniqueness of the pricing decision is not guaranteed by the current assumptions. The inventory cap creates regime changes between stockout and non-stockout states, and because a common price must apply across heterogeneous bookstores, aggregating these regime changes can make expected profit non-concave in price. In practice, however, we did not encounter numerical difficulties or evidence of multiple local optima when solving for prices. In the implementation, we therefore solve for the delivery vector conditional on price and then choose the profit-maximizing price numerically.

6 Estimation

This section describes the estimation sample and the two-step estimation strategy. We first define the sample used in the structural analysis. We then recover demand using aggregate first-period sales, which remain informative in the presence of stockouts, and estimate the supply-side parameters by matching simulated pricing and delivery moments to the data.



(a) Location of Shiga (b) Counties in Shiga

Figure 4: Map of Shiga

Note: Panel (a) shows the location of Shiga prefecture (highlighted in purple) within Japan. Panel (b) shows the county boundaries within Shiga prefecture, where the book market data is collected.

6.1 Estimation sample

A substantial share of sales comes from newly published books, top authors, and literature titles. We therefore restrict the estimation sample to literature books written by the top 1000 literature authors, observed within the first 6 months after publication. We also restrict the geographic scope to Shiga prefecture, where our wholesaler has a significant presence, holding a market share of 65.9%. Shiga is a prefecture of Japan located in the Kansai region of Honshu. Figure 4 shows the location of Shiga, which is divided into 24 counties. By 2015, Shiga had a population of 1,412,916 and a geographic area of 4,017 km^2. We use the 6-month window because it matches the horizon of the structural model and covers the period in which demand uncertainty is most important. To focus on demand uncertainty, we exclude reprints of older titles because demand for new books is typically uncertain at launch and becomes clearer over time as sales are realized. We then restrict the sample to books for which at least one bookstore has overstock in the first period. This is an invertibility requirement of the title-level demand inversion, not a discretionary sample restriction, as discussed in Section 6.2. Because 99.62% of books satisfy this condition, the excluded share is negligible and too small to materially affect the results. Finally, we randomly sample 100 books for estimation to alleviate the computational burden.

Table 1 reports summary statistics for the estimation sample. The average county population is 92,327.5, with an average of 3.7 bookstores per county. The average bookstore stocks 1,069.7 distinct titles, and the average book price is 1,299.1 JPY. At the title-store level, a bookstore sells 2.3 copies of a given title over the first 6 months after publication and receives an initial inventory of 3.1 copies on average. At the title-store-month level, each bookstore holds, on average, 1.2

Table 1: Summary Statistics of the Estimation Sample

	N	Min	Mean	Max	Sd
County Level					
Population	14	7,566.0	92,327.5	337,634.0	80,869.1
Num. store	14	1.0	3.7	14.0	3.5
Store Level					
Store size (copy)	52	11.0	1,069.7	1,957.9	576.1
Title Level					
Price (JPY)	4,262	500.0	1,299.1	7,200.0	354.1
Title-store level					
Aggre. sales (copy)	87,579	0.0	2.3	701.0	9.0
Initial inventory (copy)	87,579	0.0	3.1	413.0	7.1
Title-store-month level					
Sales (copy)	1,971,985	0.0	0.2	389.0	1.3
Inventory (copy)	1,971,985	0.0	1.2	216.0	2.8

Note: Store size is measured by the average number of book titles in the bookstore across months. Aggregate sales at the title-store level are aggregate sales of a title in a bookstore over the first 6 months since publication. Initial inventory is the initial delivery (adjusted for negative inventory) of a title to a bookstore when the title first becomes available there.

copies of a given title in inventory and sells 0.2 copies per month.

6.2 Demand Estimation

Stockouts invalidate the standard store-level share inversion because realized sales are censored by inventory. This empirical difficulty is well known in demand estimation with incomplete product availability (Conlon and Mortimer, 2013). We therefore estimate demand using only the first six-month window and invert demand at the title level rather than at the store level. This choice is consistent with the evidence in Section 3.3 that demand uncertainty is concentrated early in a title's life cycle and with Appendix B, which finds little evidence of intertemporal substitution and only limited stockout-based substitution across stores or titles.

For each title j, the estimator matches observed first-period realized sales aggregated across all bookstores carrying that title. Let δ_j denote a common title-level demand shifter, which corresponds to the unobserved component of utility shared across bookstores for title j. Conditional on the nesting parameter ω and the coefficients on store- and region-level demand shifters, predicted

aggregate realized sales are

$$\tilde{Q}_j(p_j, \mathbf{n}_j, \delta_j) = \sum_{l \in \mathcal{L}} \sum_{b \in \mathcal{B}_l} \min \left\{ n_{jbl},\ M_l q_{jbl}(p_j, \delta_j; \psi, \omega) \right\}, \tag{21}$$

where M_l is the market size of county l, n_{jbl} is first-period inventory at bookstore b, and $q_{jbl}(p_j, \delta_j; \psi, \omega)$ is the nested-logit store-choice probability implied by the demand system. The inversion targets the scalar δ_j so that predicted total realized sales in Equation (21) equal observed total realized sales for title j in the first period.

This aggregation is necessary because stockouts break the monotone mapping between store-level shares and mean utility. By contrast, if at least one bookstore carries excess inventory for title j, aggregate realized sales remain strictly increasing in the common title-level shifter δ_j. A higher δ_j raises potential demand at every bookstore, and at least one unconstrained bookstore translates that increase into higher realized sales. This is why we restrict the estimation sample to titles for which at least one bookstore has excess inventory in the first period. The title-level inversion is also consistent with the model, in which the unobserved demand shock ξ_j is common across bookstores while observed store and region heterogeneity enters through bookstore-specific covariates in the store-choice probabilities.

The implemented estimator directly searches over the nesting parameter ω together with the coefficients on the observed store- and region-level demand shifters that enter the store-choice probabilities. Although these coefficients enter utility linearly, they affect aggregate realized sales nonlinearly once demand is aggregated to the title level and censored by inventory, so they cannot be concentrated out after the inversion. Conditional on a candidate value of these parameters, the inversion recovers δ_j title by title by matching observed and predicted aggregate realized sales. The remaining linear parameters are then recovered from the inverted title-level shifter. Specifically, after obtaining $\hat{\delta}_j(\omega, \psi)$, we project it onto price and title-level observables using IV/GMM:

$$\hat{\delta}_j(\omega, \psi) = \alpha p_j + x_j' \beta + \xi_j,$$

where x_j includes title-level observables and fixed effects. The resulting residual $\hat{\xi}_j(\omega, \psi)$ enters the orthogonality conditions.

The moment conditions combine standard IV/GMM orthogonality restrictions with an additional stockout-matching moment. The orthogonality conditions are

$$\mathbb{E}\left[z_j \hat{\xi}_j(\omega, \psi) \right] = 0, \tag{22}$$

where z_j includes included demand shifters and excluded instruments. In the implementation, the included controls comprise title-level fixed effects and observed store- and region-level demand

shifters that enter predicted store choice. The excluded instruments include the number of bookstores in the market, which shifts the structure of inside-option competition and helps discipline the nesting parameter ω, and physical book characteristics such as the number of pages and volume, which are treated as supply shifters that affect price through cost while being excluded from utility conditional on the included title controls. In addition, the estimator matches the average stockout rate in the data to the average stockout rate predicted by the model. This additional moment helps discipline the extent of censoring induced by inventory constraints; it is used as a model-fit moment rather than as a separate identification argument.

Estimation therefore proceeds in two steps for each candidate value of the nonlinear parameter vector (ω, ψ). First, the estimator inverts a one-dimensional title-level shifter δ_j for each title by matching observed and predicted aggregate realized sales in the first period. Second, conditional on the recovered $\hat{\delta}_j(\omega, \psi)$, it estimates the remaining linear parameters by IV/GMM and evaluates the objective function using both the orthogonality moments and the stockout-rate moment. The estimate $(\hat{\omega}, \hat{\psi})$ minimizes this GMM criterion, after which the corresponding linear parameters are recovered from the second-step IV/GMM projection.

6.3 Supply Estimation

The supply-side parameters are estimated by simulated minimum distance. Let $\theta_3 := \{\lambda, \sigma_\xi, \beta_\xi, \sigma_\varepsilon\}$ denote the vector of supply parameters, where λ collects the coefficients in the marginal-cost function, including publisher effects and the continuous cost shifters in the number of pages and volume. The table reports only the coefficients on pages and volume, while the publisher fixed effects are suppressed for brevity. The parameter σ_ξ is the baseline noise in the publisher's signal about the title-level demand shock, and β_ξ is the additional signal noise for new-author titles. The parameter σ_ε is the standard deviation of the marginal-cost shock. Its mean is normalized to zero because the deterministic cost index already contains saturated publisher effects and observed title characteristics, so a nonzero mean would not be separately identified from the overall cost level.

We focus on the first period, when demand uncertainty is greatest, and treat the demand estimates from Section 6.2 as given. The belief parameters describe the publisher's and wholesaler's information about demand at the time of the initial pricing and delivery decision. This is a limitation of the current setup: in the counterfactual analysis, we also assume that bookstores share the same belief. For a candidate θ_3, the signal-noise scale for title j is

$$\sigma_{\text{signal},j} = \sigma_\xi + \beta_\xi \mathbf{1}\{\text{new author}_j\},$$

which parsimoniously allows publishers to be less precise about demand for new-author titles.

The parameters are jointly disciplined by moments constructed from observed first-period pric-

ing and delivery decisions. The targeted moments are: (i) publisher-specific average delivery quantities, (ii) the standard deviation of delivery, (iii) the mean realized price for new-author titles, (iv) the mean realized price for existing-author titles, and (v) the coefficients from a regression of price on the number of pages and volume. The publisher-specific delivery moments discipline publisher-level cost heterogeneity, the dispersion of delivery helps pin down residual cost dispersion and belief heterogeneity, and the price moments discipline both the continuous cost coefficients and the amount of uncertainty faced by publishers.

The simulation procedure is as follows. First, for each candidate θ_3, we draw 50 title signals and 50 marginal-cost shocks for each title. The marginal-cost shocks are mean-zero normal draws with standard deviation σ_ε. Second, for each simulated signal, we construct the posterior belief about ξ_j implied by the normal prior and signal structure in Section 5.2.1. Third, to evaluate the publisher's ex ante objective, we approximate the posterior expectation by Monte Carlo integration using 50 draws from that posterior distribution. This yields simulated first-period prices, delivery quantities, potential demand, and realized demand, where realized demand is the minimum of potential demand and delivery evaluated at the recovered title shock. Finally, we compute the simulated moments and choose θ_3 to minimize the weighted distance between the simulated and observed moment vectors. We therefore interpret the supply-side estimates as those that jointly rationalize the pricing and delivery patterns in the data, rather than as the outcome of a separate analytical identification step.

Specifically, let $m_s(\theta_3)$ denote the auxiliary statistics from simulation s, and let m_{obs} denote the observed statistics from the data. The minimum-distance estimator solves

$$\min_{\theta_3} \left(m_{obs} - \bar{m}(\theta_3)\right)' W \left(m_{obs} - \bar{m}(\theta_3)\right), \tag{23}$$

where $\bar{m}(\theta_3) \coloneqq \frac{1}{S}\sum_{s=1}^{S} m_s(\theta_3)$ and W is a weighting matrix.

6.4 Estimation results

Table 2 presents the key estimation results, with fixed effect coefficients in the demand and cost functions suppressed for brevity.

The left panel reports the demand-side estimates. The price coefficient α is estimated to be -0.002 (SE: 0.00004). The implied own-price elasticity of demand is -3.268 and the cross-price elasticity is 0.358, suggesting moderate price sensitivity and limited substitution across bookstores. The estimated nested logit coefficient ω is 0.398 (SE: 0.030), indicating moderate preference correlation within the bookstore nest. The entries $\mu_{0,\xi}$ and $\sigma_{0,\xi}$ summarize the estimated prior distribution of title-level demand shocks across books and serve as prior inputs in the supply-side simulation.

Table 2: Key Parameter Estimates

Demand Parameters			Supply Parameters		
Variable	Estimate	SE	Variable	Estimate	SE
ω	0.398	0.030	$\#Pages$	0.120	0.013
α	-0.002	4.041e-05	Volumes	0.104	0.017
$\mu_{0,\xi}$	-0.862	-	σ_ξ	1.332	0.013
$\sigma_{0,\xi}$	3.899	-	β_ξ	0.391	0.142
			σ_ε	0.623	0.026
Price Elasticity			No. title	100	
Own	-3.268		No. bookstore	45	
Cross	0.358		No. title-bookstore	2624	

Note: The left panel reports demand estimates: the price coefficient (α), the nested logit parameter (ω), and the implied own- and cross-price elasticities. The right panel reports supply estimates: the marginal cost coefficients on the number of pages and volume, the baseline signal-noise parameter (σ_ξ), the additional signal noise for new-author titles (β_ξ), and the standard deviation of the marginal cost shock (σ_ε). Standard errors are reported in the 'SE' column. Fixed effect coefficients are suppressed for brevity. $\mu_{0,\xi}$ and $\sigma_{0,\xi}$ report the mean and standard deviation of the prior distribution of title-level demand shocks carried from the demand step into the supply estimation.

The right panel reports the supply-side estimates. Marginal cost increases with the number of pages, with a coefficient of 0.120 (SE: 0.013), and with volume, with a coefficient of 0.104 (SE: 0.017). The baseline signal-noise parameter σ_ξ is estimated at 1.332 (SE: 0.013), implying that the publisher's signal about title demand is substantially noisy even after conditioning on the prior distribution of ξ_j. The estimate $\beta_\xi = 0.391$ (SE: 0.142) implies that the publisher's signal is noisier for new-author titles than for existing-author titles. The standard deviation of the marginal cost shock is $\sigma_\varepsilon = 0.623$ (SE: 0.026).

Table 3 compares the simulated moments to their data counterparts. The estimator targets publisher-specific average deliveries, but the table reports only their average across publishers for brevity, so it should be read as a summary of fit rather than a complete display of every targeted moment. Among the targeted moments shown, simulated retail prices closely track the data: $1{,}201.815$ yen for new-author titles and $1{,}225.695$ yen for established authors, against $1{,}125.462$ and $1{,}260.115$ yen in the data. The regression of price on physical book characteristics is similarly well matched, with the coefficient on the number of pages at 3.898 versus 4.105 in the data and the coefficient on volume at 0.199 versus 0.169. The standard deviation of delivery also aligns reasonably closely (3.242 versus 3.571). As additional untargeted checks, the model predicts an average delivery of 2.897 versus 3.146 in the data and an average realized demand of 1.057 versus 1.199, indicating that the estimated model also delivers reasonable fit on moments not directly used in estimation.

Table 3: Comparison of Targeted Simulated and Data Moments

Moment	Simulation	Data
Realized Demand	1.057	1.199
Delivery	2.897	3.146
SD of delivery	3.242	3.571
Mean Retail Price of New Author	1201.815	1125.462
Mean Retail Price of Existing Author	1225.695	1260.115
Regression Coefficient of Prices on $\#Page$	3.898	4.105
Regression Coefficient of Prices on Volume	0.199	0.169

Note: This table compares the moments targeted in the simulated minimum-distance estimation between the simulated model and the data. The targeted moments are publisher-specific average delivery quantities, the standard deviation of delivery, mean retail prices for new- and existing-author titles, and regression coefficients of prices on the number of pages and volume. For brevity, the table reports the average of the publisher-specific delivery moments. Average delivery and average realized demand are shown as untargeted fit statistics.

7 The Welfare Effects of RPM

7.1 Counterfactual Settings

The status quo, described in Section 5.2.1, features a publisher choosing bookstore-level deliveries and a uniform retail price before demand is realized. We compare this regime with five counterfactuals. Due to computational constraints, we restrict the counterfactual analysis to a random sample of 10 books and report averages over 10 simulations per book.

The counterfactuals vary three decision rights: who sets retail prices, whether prices can differ across bookstores, and whether prices can adjust after demand is realized. The main text focuses on these institutional differences; Appendix D gives the corresponding optimization problems.

Status quo and two pricing-flexibility benchmarks

- *Status quo*: the publisher sets a uniform RPM price and bookstore-level deliveries before demand is realized.
- *Store-flexible RPM*: the publisher still sets prices before demand is realized, but can choose bookstore-specific prices rather than a single uniform price.
- *Flexible RPM*: the publisher chooses deliveries before demand is realized and then sets bookstore-specific prices after demand is realized.

The comparison from the status quo to store-flexible RPM isolates the cost of uniform pricing across bookstores. The comparison from store-flexible RPM to flexible RPM isolates the cost of

committing to prices before demand is realized. Together, these two margins measure the welfare cost of price rigidity under publisher pricing. When $\rho = 1$, the flexible RPM allocation coincides with vertical integration.

Three RPM-policy counterfactuals

- *No RPM*: the publisher chooses deliveries before demand is realized, while bookstores choose retail prices after demand is realized.
- *Minimum RPM*: the timing is the same as under no RPM, except that the publisher additionally sets a title-specific price floor before demand is realized.
- *Maximum RPM*: the timing is the same as under no RPM, except that the publisher additionally sets a title-specific price ceiling before demand is realized.

The comparison from flexible RPM to no RPM isolates the effect of transferring ex post pricing authority from the publisher to competing bookstores. The comparisons from no RPM to minimum RPM and maximum RPM isolate the additional effects of imposing a floor or a ceiling on those competitive retail prices.

7.2 Decomposition in Empirical Notation

To connect the analytical section to the counterfactual results, define state-contingent consumer surplus in county l for title j as

$$CS_{jl}(\boldsymbol{p}_{jl}, \boldsymbol{n}_{jl}, \xi_j),$$

where $\boldsymbol{p}_{jl} = [p_{jbl}]_{b \in \mathcal{B}_l}$ and $\boldsymbol{n}_{jl} = [n_{jbl}]_{b \in \mathcal{B}_l}$ enter through the consumer-surplus expression above. Let $(\boldsymbol{p}_{jl}^{min}(\xi_j), \boldsymbol{n}_{jl}^{min})$ denote the equilibrium price vector and delivery vector under minimum RPM, and let $(\boldsymbol{p}_{jl}^{0}(\xi_j), \boldsymbol{n}_{jl}^{0})$ denote the corresponding objects under no RPM. Because deliveries are chosen before demand is realized while retail prices are chosen after demand is realized, the delivery vector does not vary with ξ_j whereas the equilibrium price vector does.

We define the minimum-RPM cutoff as

$$\xi_{jl}^{cut} := \sup \left\{ \xi_j : \tilde{Q}_{jbl}(\boldsymbol{p}_{jl}^{min}(\xi_j), n_{jbl}^{min}, \xi_j) < n_{jbl}^{min} \text{ for all } b \in \mathcal{B}_l \right\}.$$

This cutoff partitions the posterior support into the low-demand region $L_{jl} = \{\xi_j : \xi_j \leq \xi_{jl}^{cut}\}$ and the high-demand region $H_{jl} = \{\xi_j : \xi_j > \xi_{jl}^{cut}\}$. The associated posterior probability mass is

$$\pi_{jl}^{L} := \Pr(\xi_j \in L_{jl} \mid \xi_j^*), \qquad \pi_{jl}^{H} := \Pr(\xi_j \in H_{jl} \mid \xi_j^*).$$

Equivalently, the high-demand region is the set of states in which the stockout event

$$S_{jl}^{min}(\xi_j) \coloneqq \mathbb{1}\left\{\exists b \in \mathcal{B}_l : \tilde{Q}_{jbl}(\boldsymbol{p}_{jl}^{min}(\xi_j), n_{jbl}^{min}, \xi_j) = n_{jbl}^{min}\right\}$$

is equal to one. This is the empirical counterpart of the first-stockout state used in the analytical discussion and in the later figures.

Using this partition, the consumer-surplus effect of minimum RPM relative to no RPM is

$$\Delta CS_{jl} \coloneqq \mathbb{E}_{\xi|\xi_j^*}\left[CS_{jl}(\boldsymbol{p}_{jl}^{min}(\xi_j), \boldsymbol{n}_{jl}^{min}, \xi_j) - CS_{jl}(\boldsymbol{p}_{jl}^{0}(\xi_j), \boldsymbol{n}_{jl}^{0}, \xi_j)\right],$$

with low- and high-state components

$$\Delta CS_{jl}^{L} \coloneqq \mathbb{E}_{\xi|\xi_j^*}\left[\left(CS_{jl}(\boldsymbol{p}_{jl}^{min}(\xi_j), \boldsymbol{n}_{jl}^{min}, \xi_j) - CS_{jl}(\boldsymbol{p}_{jl}^{0}(\xi_j), \boldsymbol{n}_{jl}^{0}, \xi_j)\right)\mathbb{1}\{\xi_j \in L_{jl}\}\right],$$

$$\Delta CS_{jl}^{H} \coloneqq \mathbb{E}_{\xi|\xi_j^*}\left[\left(CS_{jl}(\boldsymbol{p}_{jl}^{min}(\xi_j), \boldsymbol{n}_{jl}^{min}, \xi_j) - CS_{jl}(\boldsymbol{p}_{jl}^{0}(\xi_j), \boldsymbol{n}_{jl}^{0}, \xi_j)\right)\mathbb{1}\{\xi_j \in H_{jl}\}\right],$$

so that $\Delta CS_{jl} = \Delta CS_{jl}^{L} + \Delta CS_{jl}^{H}$.

To isolate the two mechanisms within each region, let $\boldsymbol{p}_{jl}^{0}(\xi_j; \boldsymbol{n}_{jl}^{min})$ denote the no-RPM equilibrium price vector evaluated at the delivery vector $\boldsymbol{n}_{jl}^{min}$. Then for $R \in \{L, H\}$, define

$$\Delta CS_{jl}^{R,P} \coloneqq \mathbb{E}_{\xi|\xi_j^*}\left[\left(CS_{jl}(\boldsymbol{p}_{jl}^{min}(\xi_j), \boldsymbol{n}_{jl}^{min}, \xi_j) - CS_{jl}(\boldsymbol{p}_{jl}^{0}(\xi_j; \boldsymbol{n}_{jl}^{min}), \boldsymbol{n}_{jl}^{min}, \xi_j)\right)\mathbb{1}\{\xi_j \in R_{jl}\}\right],$$

$$\Delta CS_{jl}^{R,A} \coloneqq \mathbb{E}_{\xi|\xi_j^*}\left[\left(CS_{jl}(\boldsymbol{p}_{jl}^{0}(\xi_j; \boldsymbol{n}_{jl}^{min}), \boldsymbol{n}_{jl}^{min}, \xi_j) - CS_{jl}(\boldsymbol{p}_{jl}^{0}(\xi_j), \boldsymbol{n}_{jl}^{0}, \xi_j)\right)\mathbb{1}\{\xi_j \in R_{jl}\}\right],$$

where $R_{jl} = L_{jl}$ if $R = L$ and $R_{jl} = H_{jl}$ if $R = H$. By construction, $\Delta CS_{jl}^{R} = \Delta CS_{jl}^{R,P} + \Delta CS_{jl}^{R,A}$ for each region. The analytical section therefore asks how competition and demand uncertainty change the cutoff ξ_{jl}^{cut}, the posterior masses $(\pi_{jl}^{L}, \pi_{jl}^{H})$, and the magnitudes of $(\Delta CS_{jl}^{L,P}$, $\Delta CS_{jl}^{L,A}$, $\Delta CS_{jl}^{H,P}$, $\Delta CS_{jl}^{H,A})$.

Because the high-state channel operates through greater availability, it is also useful to define the delivery effect

$$\Delta N_{jl} \coloneqq \sum_{b \in \mathcal{B}_l}(n_{jbl}^{min} - n_{jbl}^{0})$$

which is the empirical counterpart of the inventory-restoration mechanism emphasized later in the counterfactual analysis.

The later section reports averages of these title-county objects across the counterfactual sample, so for any object X_{jl} we write

$$\bar{X} \coloneqq \frac{1}{|\mathcal{J}^{cf}|}\sum_{j \in \mathcal{J}^{cf}}\frac{1}{|\mathcal{L}_j|}\sum_{l \in \mathcal{L}_j} X_{jl},$$

Table 4: Counterfactual Results Summary

	Delivery (Copies)	Retail Price (JPY)	Welfare (1000 JPY) Consumer	Publisher	Bookstore	Supply	Total
Status Quo	2.26	1299.78	25.73	143.5	94.56	238.06	263.79
	(-25.3%)	(+25.7%)	(-3.6%)	(-50.9%)	(-40.4%)	(-47.2%)	(-44.8%)
	(-22.1%)	(+28.2%)	(-3.6%)	(-49.3%)	(-38.1%)	(-45.4%)	(-43.0%)
Store-flexible	2.33	1378.22	25.75	147.54	97.25	244.79	270.54
	(-23.0%)	(+33.3%)	(-3.5%)	(-49.5%)	(-38.7%)	(-45.7%)	(-43.4%)
	(-19.6%)	(+35.9%)	(-3.6%)	(-47.9%)	(-36.4%)	(-43.8%)	(-41.5%)
Flexible RPM	3.03	1033.92	26.69	292.3	158.59	450.89	477.58
	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)
	(+4.3%)	(+2.0%)	(-0.1%)	(+3.3%)	(+3.7%)	(+3.5%)	(+3.2%)
No RPM	2.9	1013.82	26.71	282.95	152.87	435.83	462.53
	(-4.1%)	(-1.9%)	(+0.1%)	(-3.2%)	(-3.6%)	(-3.3%)	(-3.1%)
	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)	(± 0.0%)
Minimum RPM	2.95	1032.69	26.65	284.57	154.44	439.01	465.67
	(-2.6%)	(-0.1%)	(-0.1%)	(-2.6%)	(-2.6%)	(-2.6%)	(-2.5%)
	(+1.6%)	(+1.9%)	(-0.2%)	(+0.6%)	(+1.0%)	(+0.7%)	(+0.7%)
Maximum RPM	2.9	1013.43	26.71	283.15	152.92	436.08	462.78
	(-4.1%)	(-2.0%)	(+0.1%)	(-3.1%)	(-3.6%)	(-3.3%)	(-3.1%)
	(-0.0%)	(-0.0%)	(-0.0%)	(+0.1%)	(+0.0%)	(+0.1%)	(+0.1%)

Note: The numbers in the first parentheses (blue) are percentage changes relative to the flexible RPM model. The numbers in the second parentheses (bold red) are percentage changes relative to the no RPM model. Reported retail prices and second-period deliveries are expectations over realized demand states. For consumer surplus, only differences across counterfactuals are economically meaningful, because its absolute level depends on the utility normalization with the outside option set to zero. We report levels and percentage changes only to facilitate comparisons across models and outcomes.

where $\mathcal{J}^{cf}$ is the set of titles used in the counterfactual exercise and $\mathcal{L}_j$ is the set of counties in which title j is sold. The central sign comparison in the later welfare discussion is therefore whether

$$\Delta\bar{C}S^H > |\Delta\bar{C}S^L|,$$

that is, whether the average high-state gain outweighs the average low-state loss.

7.3 Results

Table 4 reports equilibrium outcomes, including delivery, expected retail prices, and welfare (expected consumer, publisher, bookstore, and total surplus). We first discuss the cost of pricing rigidity under publisher pricing, then the effect of shifting ex post pricing authority to bookstores, and finally the incremental effects of minimum and maximum RPM. For ease of comparison, the table also reports percentage changes relative to the flexible RPM model (first parentheses, in blue)

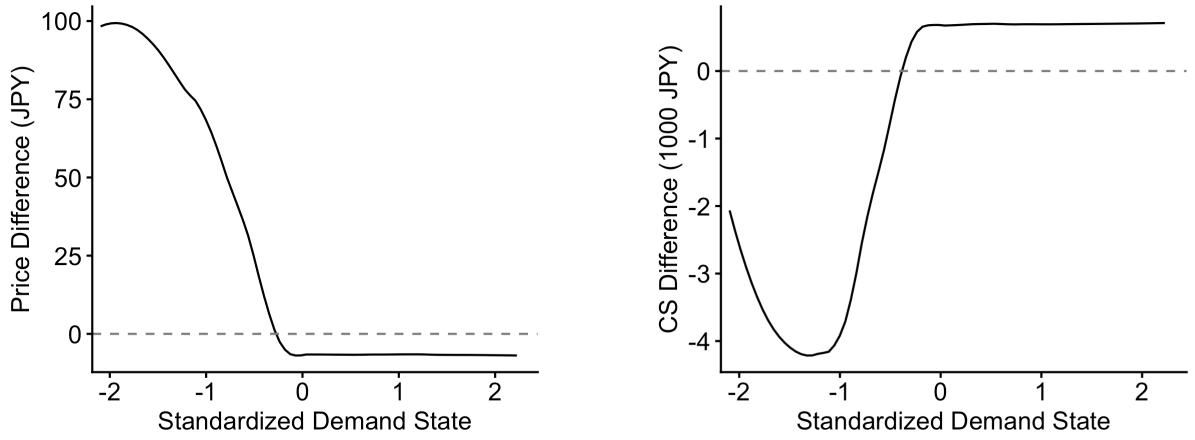
and relative to the no RPM model (second parentheses, in bold red). The discussion follows the sequence of comparisons introduced in the counterfactual-settings subsection and uses the analytical takeaways from Section 4 as an interpretive guide.

Costs of RPM rigidity Comparing the flexible RPM model with the status quo gives the combined cost of two restrictions: uniform pricing across bookstores and the need to commit to prices before demand is realized. Relative to flexible RPM, the status quo leads the publisher to choose substantially lower delivery (-25.3%) and a much higher retail price ($+25.7\%$). This is consistent with the analytical insight that inventory is valuable only insofar as firms can earn a sufficient ex post return on it: when flexible pricing is removed, that return falls and the publisher responds by scaling back delivery. Consumer surplus (-3.6%), publisher surplus (-50.9%), and bookstore surplus (-40.4%) all fall. Overall, these two sources of rigidity account for a 44.8% decrease in total surplus.

Effect of shifting pricing authority to bookstores We next compare flexible RPM with no RPM. This comparison isolates the effect of transferring ex post pricing authority from the publisher to competing bookstores, holding fixed the publisher's control over deliveries. The key analytical insight is that stronger downstream competition weakens ex ante inventory incentives because additional inventory becomes less profitable when retail competition intensifies. The structural estimates are consistent with that pattern: switching to the no RPM model reduces delivery by 4.1%. Two opposing channels then shape retail prices. The conventional channel predicts that competition lowers prices. However, reduced delivery tends to raise prices in delivery-constrained states. On balance, the competition effect dominates, so expected retail prices fall by 1.9% under no RPM and consumer surplus rises slightly ($+0.1\%$). This comes at the cost of lower publisher surplus (-3.2%), bookstore surplus (-3.6%), and total surplus (-3.1%), revealing a trade-off between consumer welfare and industry efficiency.

Effect of minimum RPM We next evaluate whether minimum RPM can restore industry efficiency and how it affects consumer surplus. The central analytical prediction is that minimum RPM can partially restore inventory by softening ex post price competition and raising the marginal return to inventory, while the consumer-surplus effect remains ambiguous. Relative to the no-RPM benchmark, minimum RPM increases average delivery, $\Delta\bar{N}$, by 1.6%, though not to the level achieved under flexible RPM. Averaging across demand states, it also raises retail prices by 1.9%, resulting in a 0.2% decrease in the average consumer-surplus effect, $\Delta\bar{CS}$.

The mechanism is illustrated in Figure 5. The left panel shows the retail price difference (minimum RPM $-$ no RPM): in low-demand states, that is, states in the low-demand region L_{jl}, the floor



(a) Retail Price Difference

(b) Consumer Surplus Difference

Figure 5: Minimum RPM vs. no RPM across Demand States

Note: Each panel plots the difference (Minimum RPM − no RPM) against the standardized demand state. A value above zero indicates that minimum RPM yields a higher level than no RPM.

binds and pushes prices above the competitive level, exactly as in the analytical section. In high-demand states, that is, states in the high-demand region H_{jl}, the stockout event under minimum RPM, $S_{jl}^{min}(\xi_j) = 1$, becomes relevant, and the higher delivery under minimum RPM improves availability, which lowers realized prices. The right panel shows the corresponding consumer-surplus difference: minimum RPM reduces consumer surplus in the low-demand region, captured by the low-state consumer-surplus effect $\Delta CS_{jl}^{L} < 0$, but raises it in the high-demand region, captured by the high-state consumer-surplus effect $\Delta CS_{jl}^{H} > 0$. This is the structural counterpart of the analytical result that the same floor that hurts consumers when it binds can help them in high-demand states through greater availability and lower realized prices. On net, the average low-state loss, $|\Delta \bar{C}S^{L}|$, outweighs the average high-state gain, $\Delta \bar{C}S^{H}$, so the average consumer-surplus effect, $\Delta \bar{C}S$, remains negative. Despite this, the additional instrument raises publisher surplus ($+0.6\%$) and bookstore surplus ($+1.0\%$), with total surplus increasing by 0.7%.

Effect of maximum RPM The maximum-RPM model generates outcomes virtually identical to those under no RPM, up to small numerical errors. The implied price ceiling is therefore non-binding in the estimated environment. Intuitively, competition among bookstores is already strong enough that the publisher does not find it profitable to impose an additional binding ceiling, so the equilibrium coincides with the no-RPM benchmark.

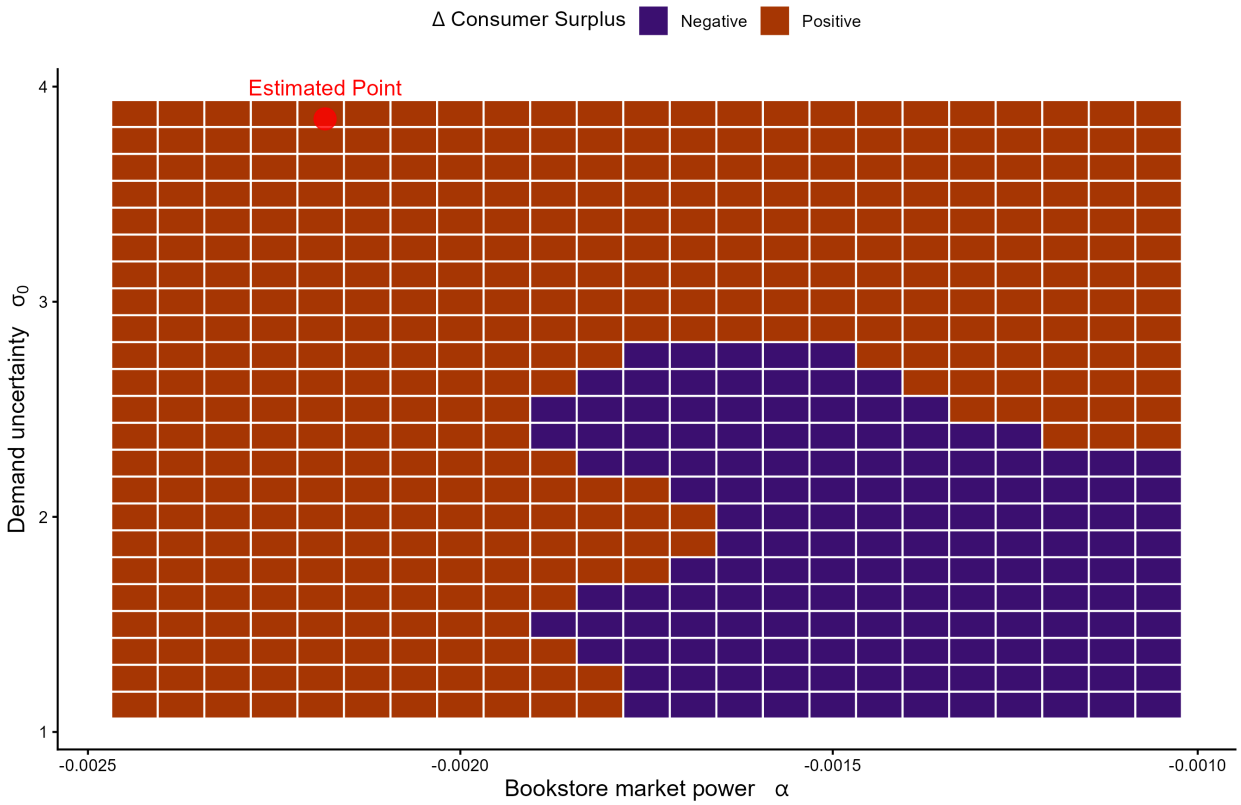
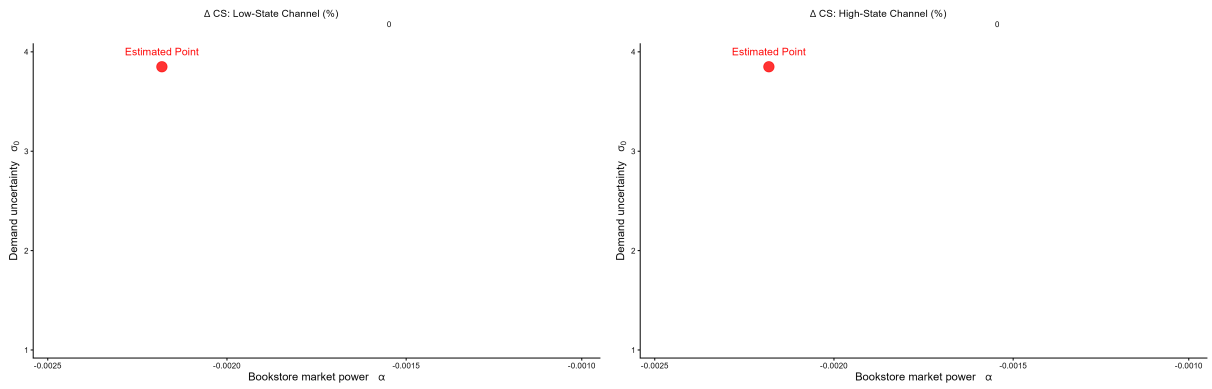


Figure 6: Consumer Surplus under Minimum RPM Relative to No RPM

Note: This figure shows how consumer surplus under minimum RPM changes relative to the no-RPM benchmark as demand uncertainty (σ_0) and price sensitivity (α), and hence the intensity of downstream competition, vary. Orange indicates that minimum RPM raises consumer surplus relative to no RPM, purple indicates that it lowers consumer surplus, and white indicates little change. When σ_0 changes, σ_ξ is adjusted proportionally so that the mean of the posterior distribution does not change.



(a) Low-Demand-State Channel (b) High-Demand-State Channel

Figure 7: Decomposition of the Consumer-Surplus Effect of Minimum RPM

Note: Both panels decompose the consumer-surplus effect of minimum RPM relative to no RPM using the stockout cutoff that mirrors the analytical model. For each title, market, and simulation draw, we define ξ^{cut} as the empirical counterpart of the minimum-RPM cutoff, ξ_{jl}^{cut}, namely the highest demand state at which no bookstore in the market stocks out under minimum RPM. States with $\xi \leq \xi^{cut}$ are classified as low-demand states, and states with $\xi > \xi^{cut}$ are classified as high-demand states. The standardized demand state is used only for comparability across titles. Panel (a) reports the total low-demand-state contribution, that is, the average low-state consumer-surplus effect $\Delta\bar{C}S^{L}$, and panel (b) reports the total high-demand-state contribution, that is, the average high-state consumer-surplus effect $\Delta\bar{C}S^{H}$. The net effect in Figure 6 is their sum.

8 Comparative Statics of Minimum RPM

The baseline counterfactual results show that minimum RPM can raise output and total surplus even though it lowers consumer surplus on average. A natural question is whether this conclusion changes in markets with different levels of demand uncertainty or downstream competition. To address this, we vary two key primitives: the degree of demand uncertainty, σ_0, and the demand price-sensitivity parameter, α, which governs the intensity of downstream competition. We then recompute welfare under no RPM and minimum RPM across all parameter combinations.[30][31] The range for σ_0 runs from zero to the estimated value, so the exercise spans cases from no prior title-level demand uncertainty to the uncertainty level implied by the estimates. The range for α is chosen to span values that imply roughly own-price elasticities from -1 to -10 at observed prices and shares, which is broad enough to cover weak to strong downstream competition in this setting. Figure 6 reports the resulting differences in consumer surplus. This exercise is the structural analogue of the analytical decomposition in Section 4. Consistent with that decomposition, the empirical low-state/high-state split is defined by the minimum-RPM cutoff, ξ_{jl}^{cut}: low states are those in the low-demand region L_{jl}, below the first stockout state, and high states are those in the

[30] When varying σ_0, we adjust σ_ξ proportionally so that the mean of the posterior demand shock remains unchanged.

[31] To limit computational burden, these are local comparative statics around the baseline shock realization: we evaluate all policies at a marginal cost shock and signal equal to zero.

high-demand region H_{jl}, where inventory constraints bind somewhere in the market. The exercise therefore traces how the average low-state consumer-surplus effect, $\Delta\bar{C}S^L$, and the average high-state consumer-surplus effect, $\Delta\bar{C}S^H$, vary with uncertainty and competition through changes in the cutoff, the stockout event, and the associated delivery effects. Because minimum RPM continues to raise quantity and supply-side welfare throughout the parameter range we consider, we focus on consumer surplus, which is the margin on which the sign of the welfare effect changes.

Figure 6 shows that minimum RPM improves consumer surplus only in a limited region of the parameter range we consider. Interpreted through the analytical decomposition, that region is where the average high-state consumer-surplus gain, $\Delta\bar{C}S^H$, exceeds the absolute value of the average low-state consumer-surplus loss, $|\Delta\bar{C}S^L|$. Within the economically plausible range just described, the gains are strongest when downstream competition is sufficiently intense and demand uncertainty is moderate, because minimum RPM then materially relaxes stockouts without generating too large a price distortion before stockouts occur. When uncertainty is low, there is little inventory distortion to fix, so the high-state channel, $\Delta\bar{C}S^H$, is weak. When uncertainty is high, the pre-stockout loss, $\Delta\bar{C}S^L$, often remains large enough to offset the post-stockout gain. At the grid point closest to the estimate, the high-state channel is positive but smaller in magnitude than the low-state loss, so minimum RPM still reduces consumer surplus overall. This pattern reinforces the paper's main policy conclusion: although a procompetitive case for minimum RPM exists in theory, the consumer-benefiting region is too limited within this plausible range to support a presumption in its favor.

Figure 7 makes the decomposition explicit. Panel (a) shows the total low-state contribution, that is, the average low-state consumer-surplus effect $\Delta\bar{C}S^L$, which is most negative where minimum RPM raises prices over a broad range of pre-stockout states. Panel (b) shows the total high-state contribution, that is, the average high-state consumer-surplus effect $\Delta\bar{C}S^H$, which is largest where minimum RPM materially relaxes stockouts in strong-demand states, especially when downstream competition is strong enough to distort inventory incentives under no RPM. The orange region in Figure 6 therefore corresponds to the parameter combinations for which the average high-state consumer-surplus gain exceeds the average low-state consumer-surplus loss in absolute value, that is, $\Delta\bar{C}S^H > |\Delta\bar{C}S^L|$. Elsewhere, the low-state losses remain larger than the high-state gains.

9 Conclusion

This paper studies RPM in an environment in which a procompetitive rationale is theoretically available: retailers must hold inventory before demand is realized, and downstream price competition can distort those inventory choices. First, we connect the legal debate over RPM to a concrete economic mechanism by distinguishing the minimum-RPM and maximum-RPM compo-

nents of fixed book pricing and showing how each operates through a different channel. Second, we develop and estimate a structural model of publishers and bookstores in the Japanese publishing industry that incorporates demand uncertainty, local bookstore competition, and inventory constraints. Third, we use the estimated model to quantify whether this procompetitive justification for RPM is strong enough to improve consumer welfare in practice.

The estimates support the empirical relevance of the mechanism. Demand is price sensitive, with an average own-price elasticity of -3.268, and substitution across bookstores is limited, with an average cross-price elasticity of 0.358. On the supply side, publishers face considerable demand uncertainty, which makes inventory policy an economically important margin and makes the inventory channel a plausible procompetitive justification for minimum RPM in our setting.

The counterfactual results provide two main conclusions. First, minimum RPM is generally difficult to defend on consumer-welfare grounds even in a setting where its procompetitive inventory channel is active. At the estimated parameter values, minimum RPM modestly benefits the supply side, but it lowers consumer surplus by 0.2% relative to the wholesale benchmark. Comparative statics show that a parameter region in which minimum RPM raises consumer surplus does exist, but it occupies only a limited part of the plausible parameter range we consider, where σ_0 runs from no prior uncertainty to the estimated level and α spans values that imply roughly own-price elasticities from -1 to -10. The paper therefore does not reject the theory that minimum RPM may sometimes help consumers, but it shows that such cases are likely to be exceptional even in a setting chosen to favor that mechanism.

Second, distinguishing the upper- and lower-price components of RPM is important both economically and legally. In our counterfactuals, maximum RPM is essentially non-binding in the estimated environment and yields outcomes that are virtually identical to no RPM. Relative to the wholesale benchmark, it changes consumer surplus by approximately 0.0% and publisher surplus by 0.1%. Minimum RPM, by contrast, mainly shifts surplus toward the supply side and does not improve consumer welfare at the estimated parameter values. The key contrast is therefore not between RPM and no RPM in the abstract, but between two distinct restraints that operate through different channels and have different welfare implications.

Taken together, these results imply that the procompetitive case for RPM should be stated narrowly and with attention to mechanism. Our findings support treating maximum RPM and minimum RPM differently. For minimum RPM, the inventory-based procompetitive justification is best understood as a limited defense that may apply in particular market environments, not as a basis for a general presumption of consumer benefit. In the Japanese publishing industry we study, that justification is theoretically relevant but empirically too limited to support a broad claim that minimum RPM is usually beneficial to consumers.

There are several limitations in the analysis. First, the inter-brand competition across book

titles was not considered. If we analyze other categories such as self-help books, the substitution between book titles would be more important. Second, the bargaining powers of the bookstores and publishers were not explicitly estimated because the prevailing revenue sharing rate was observed. If we consider a counterfactual scenario where the wholesale price is determined by bargaining, the estimation of the bargaining power behind the observed revenue-sharing rate is necessary. However, we are skeptical whether the current revenue sharing rate could be regarded as a consequence of static bargaining because there seems to be a historical reason for these numbers. Third, another competitive channel of RPM, such as the enhancement of the in-store advertisement activity, was not considered. In addition, increasing inventory itself might have some advertising effect, because it increases books on the shelves. We did not consider the former because we did not have data on in-store advertising and did not consider the latter because the equilibrium becomes unstable if we consider this channel. By considering either channel, the results will become more in favor of the RPM, and hence our conclusion will not change. Fourth, the data is limited to the bookstores that transact with the data provider and has a point-of-sale system. Thus, the strategic responses of the other bookstores are not considered in the counterfactual simulations. Although the current coverage of the data is wide enough, the analysis based on a more comprehensive data set is desired. Fifth, the analysis does not utilize the variation in the regulatory regime. If similar data is available in countries that experienced a change in the resale price maintenance regime, such as the Norwegian book market, the model prediction could be validated.

References

Asker, John and Heski Bar-Isaac, "Raising Retailers' Profits: On Vertical Practices and the Exclusion of Rivals," *American Economic Review*, February 2014, *104* (2), 672–686.

Bailey, Elizabeth M. and Gregory K. Leonard, "Minimum Resale Price Maintenance: Some Empirical Evidence from Maryland," SSRN Scholarly Paper ID 1532998, Social Science Research Network, Rochester, NY January 2010.

Ball, C., S. Davies, M. Olczak, and Chris M. Wilson, "An Evaluation of the Impact upon Productivity of Ending Resale Price Maintenance on Books," *manuscript*, 2008.

Beck, Jonathan, "Fixed, Focal, Fair? Book Prices Under Optional Resale Price Maintenance," SSRN Scholarly Paper ID 645861, Social Science Research Network, Rochester, NY December 2004.

_ , *Resale Price Maintenance for Books*, Norderstedt: Books On Demand, January 2008.

Bittlingmayer, George, "Resale Price Maintenance in the Book Trade with an Application to Germany," *Journal of Institutional and Theoretical Economics*, 1988, *144* (5), 789–812.

Blair, Roger D. and Francine Lafontaine, "Will Khan Foster or Hinder Franchising? An Economic Analysis of Maximum Resale Price Maintenance," *Journal of Public Policy & Marketing*, March 1999, *18* (1), 25–36.

Bonnet, Céline and Pierre Dubois, "Inference on Vertical Contracts between Manufacturers and Retailers Allowing for Nonlinear Pricing and Resale Price Maintenance," *The RAND Journal of Economics*, 2010, *41* (1), 139–164.

Butz, David A., "Vertical Price Controls with Uncertain Demand," *The Journal of Law and Economics*, October 1997, *40* (2), 433–460.

Cardell, N. Scott, "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity," *Econometric Theory*, 1997, *13* (2), 185–213.

Conlon, Christopher T. and Julie Holland Mortimer, "Demand Estimation under Incomplete Product Availability," *American Economic Journal: Microeconomics*, November 2013, *5* (4), 1–30.

Daljord, Øystein, "Durable Goods Adoption and the Consumer Discount Factor: A Case Study of the Norwegian Book Market," *Management Science*, December 2021.

Deneckere, Raymond, Howard P. Marvel, and James Peck, "Demand Uncertainty, Inventories, and Resale Price Maintenance," *The Quarterly Journal of Economics*, 1996, *111* (3), 885–913.

_ , _ , and _ , "Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition," *The American Economic Review*, 1997, *87* (4), 619–641.

Fishwick, Francis, "Book Prices in the UK Since the End of Resale Price Maintenance," *International Journal of the Economics of Business*, November 2008, *15*, 359–377.

Flath, David and Tatsuhiko Nariu, "Demand Uncertainty and Resale Price Maintenance."

Gilligan, Thomas W., "The Competitive Effects of Resale Price Maintenance," *The RAND Journal of Economics*, 1986, *17* (4), 544–556.

Ho, Katherine, Justin Ho, and Julie Holland Mortimer, "The Use of Full-Line Forcing Contracts in the Video Rental Industry," *American Economic Review*, April 2012, *102* (2), 686–719.

Ippolito, Pauline M. and Thomas R. Overstreet, "Resale Price Maintenance: An Economic Assessment of the Federal Trade Commission's Case against the Corning Glass Works," *The Journal of Law and Economics*, April 1996, *39* (1), 285–328.

Japan Book Publishers Association, "Japan Book Publishers Association - 2017 - An Introduction

to Publishing in Japan.Pdf," 2017.

Japan Fair Trade Commission, "About Retail Price Maintenance of Copyrighted Works," March 2001.

Jullien, Bruno and Patrick Rey, "Resale Price Maintenance and Collusion," *The RAND Journal of Economics*, 2007, *38* (4), 983–1001.

Klein, Benjamin, "Distribution Restrictions Operate by Creating Dealer Profits: Explaining the Use of Maximum Resale Price Maintenance in State Oil v. Khan," *Supreme Court Economic Review*, January 1999, *7*, 1–58.

Krishnan, Harish and Ralph A. Winter, "Vertical Control of Price and Inventory," *American Economic Review*, December 2007, *97* (5), 1840–1857.

Li, Hui, "Intertemporal Price Discrimination with Complementary Products: E-Books and E-Readers," *Management Science*, March 2019, *65* (6), 2665–2694.

_ , "Are E-Books a Different Channel? Multichannel Management of Digital Products," *Quantitative Marketing and Economics*, June 2021, *19* (2), 179–225.

MacKay, Alexander and David A. Smith, "Challenges for Empirical Research on RPM," *Review of Industrial Organization*, March 2017, *50* (2), 209–220.

Marvel, Howard P. and James Peck, "Demand Uncertainty and Returns Policies," *International Economic Review*, August 1995, *36* (3), 691.

_ **and Stephen McCafferty**, "Resale Price Maintenance and Quality Certification," *The RAND Journal of Economics*, 1984, *15* (3), 346–359.

McFadden, Daniel, "Modelling the Choice of Residential Location," *Cowles Foundation Discussion Papers.*, December 1977.

Ornstein, Stanley I. and Dominique M. Hanssens, "Resale Price Maintenance: Output Increasing or Restricting? The Case of Distilled Spirits in the United States," *The Journal of Industrial Economics*, 1987, *36* (1), 1–18.

Spengler, Joseph J., "Vertical Integration and Antitrust Policy," *Journal of Political Economy*, 1950, *58* (4), 347–352.

Telser, Lester G., "Why Should Manufacturers Want Fair Trade?," *The Journal of Law & Economics*, 1960, *3*, 86–105.

Tirole, Jean, *The Theory of Industrial Organization*, Cambridge, MA, USA: MIT Press, August 1988.

Tosdal, H. R., "Price Maintenance in the Book Trade," *The Quarterly Journal of Economics*, November 1915, *30* (1), 86–109.

Williams, Rhys J, "Empirical Effects of Resale Price Maintenance: Evidence from Fixed Book Price Policies in Europe," *Journal of Competition Law & Economics*, January 2024, *20* (1-2), 108–136.

Xia, Tianli, "Welfare Effects of Resale Price Maintenance: Evidence from the Chinese Pharmaceutical Industry," 2025.

Appendices

A Data Sampling and Construction

OpenBD data We merge categorical information from the OpenBD project provided by Calil.jp with our transaction data.[32] OpenBD provides a free API for bibliographic information. Its category code has four digits: the first classifies the target audience, the second the publication form, and the last two the content of the book, such as literature, language, or history.[33] For example, code 8091 denotes a separate edition of Japanese literature for a general audience.

Constructing inventory data for brick-and-mortar bookstores. We construct inventory from the transaction records. Inventory at the beginning of month t equals cumulative deliveries up to month t minus cumulative returns and cumulative sales up to month $t-1$. At the title-store-month level, negative inventory appears in 19% of observations, or 11% in the estimation sample of titles published by the top 1,000 literature authors within six months after publication. These negative values are small in magnitude, with median 1 and mean 2.5.

Several factors may generate this pattern. Chain stores may transfer books across outlets. Some bookstores may transact with multiple wholesalers, so deliveries from other wholesalers are missing from our data even though sales are observed directly from the bookstore. For books published before 2015, some initial inventories may also be unobserved because our dataset begins in 2015. To address this issue, we add the smallest title-store-specific initial inventory needed to keep inventory nonnegative in every period.

Imputing zero sales data for brick-and-mortar bookstores. A second data issue is that we observe a title at a bookstore in a given month only when there is at least one unit of delivery, sales, or return in that month. If no such activity occurs, the title-store-month observation is missing. A missing observation can therefore mean either no activity or unavailability.

We classify a missing title-store-month observation as no activity, with zero delivery, zero sales, and zero return, if the same title has at least one observed delivery, sale, or return at that bookstore at some point within six months after publication. If a title has no observed delivery, sale, or return at that bookstore at any time within six months after publication, we classify all missing observations for that title-store pair as unavailability. We then impute zeros only for the observations classified as no activity.

B Evidence on Stockout-based Substitution

This appendix provides reduced-form evidence on whether consumers substitute after stockouts across time, bookstores, or titles. The purpose is not to identify a structural substitution model, but to assess whether the assumption of limited stockout-based substitution used in the main text is reasonable. We report results for both the all-category sample and the literature sample.

[32] See `https://openbd.jp/`, accessed on November 19, 2021.

[33] The classification of the third digit includes 1. General. 2. Philosophy. 3. History. 4. Social Science. 5. Natural Science. 6. Engineering. 7. Industry. 8. Arts. 9. Language. 10. Literature.

Table 5: Substitution patterns

	Inter-temporal	Across-store	Across-title
All category			
Last-period stockout indicator	0.0277		
	(0.0263)		
Rival-store stockout indicator		0.0975**	
		(0.0290)	
Within-category stockout indicator			2e-04***
			(0.0000)
Cross-category stockout indicator			1e-04***
			(0.0000)
Num.Obs.	38,026,060	38,026,060	38,026,060
R2	0.500	0.407	0.597
Literature			
Last-period stockout indicator	0.0085		
	(0.0402)		
Rival-store stockout indicator		0.0908**	
		(0.0302)	
Cross-category stockout indicator			-3e-04
			(3e-04)
Cross-category stockout indicator			3e-04***
			(1e-04)
Num.Obs.	9,596,490	9,596,490	9,596,490
R2	0.673	0.575	0.779

Note: As the controls, we include the inventory level of the title at the bookstores and other bookstores, if across-bookstore substitution is considered. We also include title, bookstore, post-publication months, and year-month fixed effects.

For title j at bookstore b in county l in month t, we estimate

$$Sales_{jbt} = \alpha \mathbb{1}\{stockout_{jbt}\} + \beta X_{jbt} + \epsilon_{jbt}, \tag{24}$$

$Sales_{jbt}$ is the sales volume for title j at bookstore b at time t. The stockout measure is defined differently depending on the substitution margin being studied. The control vector X_{jbt} includes inventory controls for the focal title and, when relevant, for competing bookstores. We also include title, bookstore, post-publication-month, and year-month fixed effects. Table 5 reports the results.

Inter-temporal substitution To test inter-temporal substitution, we define the stockout indicator as whether title j stocked out at bookstore b in month $t-1$:

$$\mathbb{1}\{stockout_{jbt}\} \coloneqq \mathbb{1}\{Sales_{jb,t-1} = Inventory_{jb,t-1}\}.$$

Columns 1 and 5 of Table 5 report the results for the all-category and literature samples, respectively. In both cases, the coefficient on the lagged stockout indicator is statistically insignificant, providing no evidence of inter-temporal substitution.

Across-bookstore substitution To test across-bookstore substitution, we define the stockout measure as the number of stockouts for title j at other bookstores in the same county in month t:

$$\mathbb{1}\{stockout_{jbt}\} \coloneqq \sum_{d \in \mathbb{B}_l, d \neq b} \mathbb{1}\{Sales_{jdt} = Inventory_{jdt}\}.$$

Columns 2 and 6 of Table 5 show positive and statistically significant coefficients in both samples. This pattern is consistent with some across-bookstore substitution, but it may also reflect common title-level demand shocks that raise sales and stockouts simultaneously across stores. We therefore do not interpret this coefficient as clean causal evidence of substitution. Its magnitude is nevertheless small: one additional stockout at another bookstore is associated with only about 0.09 additional units sold at bookstore b.

Across-title substitution To test across-title substitution, we define two stockout counts at bookstore b: one for other titles in the same category as title j and one for titles in different categories,

$$\mathbb{1}\{stockout^1_{jbt}\} \coloneqq \sum_{k \neq j, k \in \mathbb{C}_j} \mathbb{1}\{Sales_{kbt} = Inventory_{kbt}\},$$

$$\mathbb{1}\{stockout^2_{jbt}\} \coloneqq \sum_{k \neq j, k \notin \mathbb{C}_j} \mathbb{1}\{Sales_{kbt} = Inventory_{kbt}\},$$

where $\mathbb{C}_j$ denotes the category of title j. The corresponding coefficients capture within-category and between-category substitution, respectively.

Columns 3–4 and 7–8 of Table 5 report the results for the all-category and literature samples. In the all-category sample, both coefficients are positive and statistically significant, so we cannot rule out across-title substitution. However, the magnitudes are economically very small: raising

sales of the focal title by one unit would require more than 5,000 same-category stockouts or more than 10,000 different-category stockouts.

In the literature sample, the coefficient on same-category stockouts is insignificant, while the coefficient on other-category stockouts is positive and significant. This suggests little evidence that consumers substitute across literature titles within the same bookstore after a stockout. Although stockouts in non-literature categories may shift some demand toward literature titles, the effect is again small: increasing sales of the focal literature title by one unit would require roughly 3,300 stockouts in other categories.

Taken together, these results support the modeling choice in the main text. We find no evidence of inter-temporal substitution, only weak and non-causal evidence of across-bookstore substitution, and economically small across-title effects. For that reason, we abstract from stockout-based substitution in the demand model.

C Consumer Surplus Calculation

This appendix derives the stockout-adjusted consumer-surplus index used in the main text. The derivation relies on two maintained assumptions from the demand model: consumers do not continue searching after a stockout, and when a bookstore receives more visitors than available copies, those copies are allocated by random rationing. Under those assumptions, the relevant welfare object for a given book j in county l is

$$CS_{jl} = -\alpha^{-1} \sum_{b \in \{0\} \cup \mathcal{B}_l} \mathbb{P}(i \text{ chooses bookstore } b) \mathbb{E}_\epsilon(u_{ijbl} \mid i \text{ chooses bookstore } b)\, \mathbb{P}(i \text{ can buy from } b), \tag{25}$$

where, under random rationing,

$$\mathbb{P}(i \text{ can buy from } b) = \min \left\{ 1, \frac{n_{jbl}}{M q_{jbl}(\boldsymbol{p}_{jl}, \xi_j)} \right\}$$

for bookstore options and equals one for the outside option. For notational convenience, let option $b = 0$ denote the outside option and define $\delta_{j0l} = 0$. The only nonstandard term relative to the usual nested-logit formula is the service probability, which captures stockouts. Under random rationing, the service event is independent of the individual taste shock conditional on the chosen option, so it is sufficient to derive the conditional expected utility term for the nested-logit model.

Conditional on consumer i choosing option b, the expected utility is

$$\begin{aligned}
&\mathbb{E}_\epsilon\left(u_{ijbl} \mid i \text{ chooses option } b\right) \\
&= \delta_{jbl} + \mathbb{E}\left(\epsilon_{ijbl} \mid i \text{ chooses option } b\right) \\
&= \delta_{jbl} + \gamma + \log \sum_{g=0}^{1} \left[\sum_{d \in g} \exp\left(\frac{\delta_{jdl}}{1-\omega}\right)\right]^{(1-\omega)} - \delta_{jbl} \qquad (26) \\
&= \gamma + \log \sum_{g=0}^{1} \left[\sum_{d \in g} \exp\left(\frac{\delta_{jdl}}{1-\omega}\right)\right]^{(1-\omega)}.
\end{aligned}$$

where γ is Euler's constant. This expression is common across all bookstore options because it depends on the inclusive value of the market rather than on the identity of the chosen bookstore.

To derive Equation (26), it is sufficient to derive the conditional expectation of the idiosyncratic shock under the nested-logit model. We do so using the standard GEV representation. Let alternatives be indexed by $r \in \mathcal{R}$, let V_r denote deterministic utility, and let $G(e^{V_1}, \ldots, e^{V_R})$ be the homogeneous degree-one generating function for a GEV model. Under regularity conditions, McFadden (1977) shows that the choice probability of alternative r is

$$s_r = \frac{e^{V_r} G_r(e^{V_1}, \ldots, e^{V_R})}{G(e^{V_1}, \ldots, e^{V_R})}, \qquad (27)$$

where G_r is the derivative of G with respect to the rth argument.

Using this representation, the conditional expectation of the idiosyncratic shock for alternative r, given that r is chosen, is

$$\begin{aligned}
&\mathbb{E}[\epsilon_r \mid r \text{ chosen}] \\
&= \int_{-\infty}^{\infty} \epsilon e^{-\epsilon} G_r(e^{V_1 - V_r - \epsilon}, \ldots, e^{-\epsilon}, \ldots, e^{V_R - V_r - \epsilon}) \exp\left\{-G(e^{V_1 - V_r - \epsilon}, \ldots, e^{-\epsilon}, \ldots, e^{V_R - V_r - \epsilon})\right\} \, d\epsilon \\
&= \int_{-\infty}^{\infty} \epsilon e^{-\epsilon} G_r(e^{V_1}, \ldots, e^{V_R}) \exp\left\{-e^{-\epsilon} e^{-V_r} G(e^{V_1}, \ldots, e^{V_R})\right\} \, d\epsilon.
\end{aligned} \qquad (28)$$

The second line uses the fact that G_r is homogeneous of degree zero. Letting $t = e^{-\epsilon}$ yields the standard result

$$\mathbb{E}[\epsilon_r \mid r \text{ chosen}] = \gamma + \log(G_r) - \log(s_r), \qquad (29)$$

where γ is Euler's constant. In the standard logit case, $G_r = 1$, so the expectation reduces to the familiar logit formula. In the nested-logit model used in this paper, the generating function implies

$$G_r = e^{V_r \frac{\omega}{1-\omega}} \left[\sum_{k \in g(r)} \exp\left(\frac{V_k}{1-\omega}\right)\right]^{-\omega}, \qquad (30)$$

where $g(r)$ denotes the nest containing alternative r. Substituting this expression and the nested-logit share formula into the previous equation yields

$$\mathbb{E}\left(\epsilon_{ijbl} \mid i \text{ chooses bookstore } b\right) = \gamma + \log \sum_{g=0}^{1} \left[\sum_{d \in g} \exp\left(\frac{\delta_{jdl}}{1-\omega}\right)\right]^{(1-\omega)} - \delta_{jbl},$$

which implies Equation (26). Combining this conditional expected utility with the service probability above gives the stockout-adjusted consumer-surplus index used in the main text. If consumers could continue searching after a stockout, a different surplus object would be required.

D Counterfactual Institutions Used in the Main Text

Let $\mathbf{p}_j \coloneqq [p_{jbl}]'_{b \in \mathcal{B}_l,\, l \in \mathcal{L}}$ and $\mathbf{d}_j \coloneqq [d_{jbl}]'_{b \in \mathcal{B}_l,\, l \in \mathcal{L}}$ collect prices and deliveries across bookstores and across regions.

Overview This appendix formalizes the counterfactual regimes reported in the welfare section. In all of them, the publisher retains the inventory decision. The models differ only in who sets retail prices, whether prices can vary across bookstores, and whether minimum or maximum RPM constrains those prices. As in the main text, the publisher observes the title-specific signal before choosing any ex ante decision, and all ex ante expectations are taken with respect to the posterior distribution of ξ_j conditional on that signal. We suppress the conditioning on ξ_j^* in some notation below when it is not needed for readability.

D.1 Store-flexible RPM

After observing the signal but before the realization of ξ_j, the publisher chooses bookstore-specific prices p_{jbl} and deliveries d_{jbl} to maximize

$$\max_{\mathbf{p}_j, \mathbf{d}_j} \quad \mathbb{E}_{\xi \mid \xi_j^*} \left\{ \sum_{l \in \mathcal{L}} \sum_{b \in \mathcal{B}_l} \left[\rho\, p_{jbl} \tilde{Q}_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j) - MC_{jbl} d_{jbl} \right] \right\}, \tag{31}$$

subject to

$$\begin{aligned} & p_{jbl} \geq 0, \\ & d_{jbl} \geq 0 \qquad \forall\, b \in \mathcal{B}_l,\ l \in \mathcal{L} \end{aligned}$$

D.2 Flexible RPM

In this counterfactual, the publisher chooses deliveries before demand is realized, but retains pricing authority and sets retail prices after the demand shock is realized.

Given realized demand ξ_j and pre-committed inventories $\mathbf{d}_{jl}$, the publisher chooses store-specific

prices p_{jbl} for each bookstore to maximize its realized profit:

$$\max_{\mathbf{p}_j} \quad \sum_{l\in\mathcal{L}}\sum_{b\in\mathcal{B}_l}\Big[\rho\, p_{jbl}\,\tilde{Q}_{jbl}\big(\mathbf{p}_{jl}, d_{jbl}, \xi_j\big)\Big], \tag{32}$$

subject to

$$p_{jbl} \geq 0 \qquad \forall\, b \in \mathcal{B}_l,\ l \in \mathcal{L}$$

Anticipating the pricing policy in (32), in the first stage the publisher chooses inventories to maximize expected profit:

$$\max_{\mathbf{d}_j} \quad \mathbb{E}_{\xi|\xi_j^*}\left\{\sum_{l\in\mathcal{L}}\sum_{b\in\mathcal{B}_l}\Big[\rho\, p^*_{jbl}(\xi_j; d_{jl})\,\tilde{Q}_{jbl}\big(\mathbf{p}^*_j(\xi_j; \mathbf{d}_{jl}), d_{jbl}, \xi_j\big) - MC_{jbl}\, d_{jbl}\Big]\right\}, \tag{33}$$

subject to

$$d_{jbl} \geq 0 \qquad \forall\, b \in \mathcal{B}_l,\ l \in \mathcal{L},$$

where $p^*_{jbl}(\xi_j; \mathbf{d}_{jl})$ denotes the optimal ex post price that solves (32) for a given realization of ξ_j and inventories $\mathbf{d}_{jl}$.

D.3 No RPM

Bookstore's Problem Fix a book j, bookstore b, and county l. In the counterfactual, the publisher chooses delivery d_{jbl} before demand is realized. After the demand shock ξ_j is realized, bookstore b sets its retail price p_{jbl}, taking its inventory d_{jbl} as given and facing the revenue-sharing rule. Given ξ_j, the bookstore's problem is

$$\max_{p_{jbl}} \quad (1-\rho)\, p_{jbl}\,\tilde{Q}_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j), \tag{34}$$

subject to

$$p_{jbl} \geq 0.$$

Since $(1-\rho)$ is a constant multiplicative factor, the bookstore chooses p_{jbl} to maximize revenue, taking as given competitors' prices $\{p_{jb'l}\}_{b'\neq b}$ and the delivery vector $\mathbf{d}_{jl}$. Let $\mathbf{p}^*_{jl}(\xi_j; \mathbf{d}_{jl})$ denote the Nash equilibrium pricing outcome in region l.

Publisher's Problem After observing the signal but before the realization of ξ_j, the publisher chooses deliveries $\mathbf{d}_j$ to maximize expected profits, anticipating the bookstores' pricing decisions after demand is realized. Under revenue sharing, the publisher receives share ρ of realized revenue

and pays the marginal cost for each delivered unit. The publisher's problem is

$$\max_{\mathbf{d}_j} \quad \mathbb{E}_{\xi|\xi_j^*}\left\{\sum_{l\in\mathcal{L}}\sum_{b\in\mathcal{B}_l}\left[\rho\, p^*_{jbl}(\xi_j;\mathbf{d}_{jl})\,\tilde{Q}_{jbl}\big(\mathbf{p}^*_{jl}(\xi_j;\mathbf{d}_{jl}), d_{jbl}, \xi_j\big) - MC_{jbl}\, d_{jbl}\right]\right\}, \tag{35}$$

subject to

$$d_{jbl} \geq 0 \qquad \forall\, b \in \mathcal{B}_l,\ l \in \mathcal{L}.$$

D.4 Minimum RPM

This counterfactual is identical to the no-RPM model except that, after observing the signal, the publisher also chooses a title-specific floor $\underline{p}_j$ before demand is realized. The bookstore's ex post pricing problem becomes

$$\max_{p_{jbl}} \quad (1-\rho)\, p_{jbl}\, \tilde{Q}_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j)\,, \tag{36}$$

subject to

$$p_{jbl} \geq \underline{p}_j.$$

The publisher's ex ante problem is the same as in the no-RPM model, except that it jointly chooses deliveries and the floor and anticipates this constrained pricing game.

D.5 Maximum RPM

This counterfactual is identical to the no-RPM model except that, after observing the signal, the publisher also chooses a title-specific ceiling $\bar{p}_j$ before demand is realized. The bookstore's ex post pricing problem becomes

$$\max_{p_{jbl}} \quad (1-\rho)\, p_{jbl}\, \tilde{Q}_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j)\,, \tag{37}$$

subject to

$$p_{jbl} \leq \bar{p}_j,$$

$$p_{jbl} \geq 0.$$

The publisher's ex ante problem is again the same as in the no-RPM model, except that it jointly chooses deliveries and the ceiling and anticipates this constrained pricing game.

E Extension: Decentralized Inventory Decision

This appendix considers an extension in which the inventory decision is returned to bookstores and the publisher instead chooses a wholesale price. These models are not part of the main welfare

comparison in Section 7; they are included as an extension and for comparison with the centralized-inventory benchmark above. We keep the revenue-sharing setting at rate ρ. Let the indicator $RP = 1$ denote that the free-return policy is allowed, with $RP = 0$ otherwise. The setup is general enough to incorporate institutions with or without revenue sharing and with or without the free-return policy. For pricing, we consider three scenarios: (i) no RPM, (ii) maximum RPM, and (iii) minimum RPM.

E.1 No RPM

In this model, the bookstores can freely adjust prices after demand uncertainty is resolved. We illustrate this model as follows.

Let unsold inventory be denoted by

$$s_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j) \;\equiv\; d_{jbl} - \tilde{Q}_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j) \;=\; \max\{0,\, d_{jbl} - Q_{jbl}(\mathbf{p}_{jl}, \xi_j)\}.$$

Bookstore's Pricing Problem After the demand is realized, given $(\mathbf{d}_{jl}, \xi_j)$, the bookstore chooses p_{jbl} to maximize

$$\max_{p_{jbl}} \quad (1-\rho)\, p_{jbl}\, \tilde{Q}_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j) \;+\; \text{RP} \cdot w_j\, s_{jbl}(\mathbf{p}_{jl}, d_{jbl}, \xi_j), \tag{38}$$

subject to

$$p_{jbl} \geq 0.$$

Let $\mathbf{p}^*_{jl}(\xi_j, \mathbf{d}_{jl})$ be the Nash equilibrium outcome of the pricing game across bookstores in region l.

Bookstore's Delivery Problem Before ξ_j is realized, bookstore b chooses d_{jbl} anticipating the ex post pricing rule $\mathbf{p}^*_{jl}(\xi_j, \mathbf{d}_{jl})$. The bookstore's expected profit is

$$\max_{d_{jbl}} \quad \mathbb{E}_\xi\Big[(1-\rho)\, p^*_{jbl}(\xi_j, \mathbf{d}_{jl})\, \tilde{Q}_{jbl}\big(\mathbf{p}^*_{jl}(\xi_j, \mathbf{d}_{jl}), d_{jbl}, \xi_j\big) \;-\; w_j d_{jbl} + \text{RP}\cdot w_j\, s_{jbl}\big(\mathbf{p}^*_{jl}(\xi_j, \mathbf{d}_{jl}), d_{jbl}, \xi_j\big)\Big], \tag{39}$$

subject to

$$d_{jbl} \geq 0.$$

Let $d^*_{jbl}(w_j)$ denote the bookstore's optimal delivery choice.

Publisher's Problem The publisher chooses w_j anticipating the bookstores' delivery choices and ex-post pricing.

Let realized sales under induced decisions be

$$\tilde{Q}_{jbl}(\xi_j; w_j) \;\equiv\; \tilde{Q}_{jbl}\Big(\mathbf{p}^*_{jl}(\xi_j, \mathbf{d}_{jl}),\, d^*_{jbl}(w_j),\, \xi_j\Big), \tag{40}$$

and

$$s_{jbl}(\xi_j; w_j) \;\equiv\; s_{jbl}\Big(\mathbf{p}^*_{jl}(\xi_j, \mathbf{d}_{jl}),\, d^*_{jbl}(w_j),\, \xi_j\Big). \tag{41}$$

The publisher's expected profit can be written compactly as

$$\max_{w_j} \quad \mathbb{E}_\xi \left\{ \sum_{l\in\mathcal{L}} \sum_{b\in\mathcal{B}_l} \Big[\rho\, p^*_{jbl}(\xi_j, \mathbf{d}_{jl}(w_j))\, \tilde{Q}_{jbl}(\xi_j; w_j) + (w_j - MC_{jbl}) d^*_{jbl}(w_j) - RP \cdot w_j\, s_{jbl}(\xi_j; w_j) \Big] \right\}, \tag{42}$$

subject to

$$w_j \geq 0.$$

E.2 Minimum RPM

Next, we drop fixed RPM but allow minimum RPM in the extension. The timeline and the bookstore's objective function remain identical to the no-RPM extension. The key distinction is that the publisher sets a **price floor** $\underline{p}_j$, in addition to the wholesale price, before demand is realized. Consequently, bookstores must set their retail prices p_{jbl} to satisfy the constraint $p_{jbl} \geq \underline{p}_j$ in the price-setting problem.

E.3 Maximum RPM

Finally, we drop fixed RPM but allow maximum RPM in the extension. The timeline and the bookstore's objective function remain identical to the no-RPM extension. The key distinction is that the publisher sets a **price ceiling** $\bar{p}_j$, in addition to the wholesale price, before demand is realized. Consequently, bookstores must set their retail prices p_{jbl} to satisfy the constraint $p_{jbl} \leq \bar{p}_j$ in the price-setting problem.