Too Much Unsold Stock?: Environmental Consequences of Price Commitment

Kecheng Hsu* Kohei Kawaguchi† Kei Kawai‡

October 15, 2026

Abstract

*Department of Economics, School of Business and Management, Hong Kong University of Science and Technology. email: khsu@connect.ust.hk

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

‡Graduate School of Economics, The University of Tokyo. e-mail: kawaixkei@gmail.com

1 Introduction

Every year, new and unused products worth billions of dollars are discarded. One Amazon warehouse in the U.K. alone, for example, was reported to have destroyed 124,000 unused products in a single week, including laptops, electronics, and appliances. According to one estimate, about 4.9% of inventory across a range of industries is discarded in the U.S.A without ever being used by consumers because of overproduction (Avery Dennison, 2022). Not only is the practice of discarding inventory economically inefficient, but it also imposes significant environmental costs.

Recently, public concern over the environmental impact of overproduction has prompted some countries to take legislative action. For example, in 2023, the EU passed legislation banning the destruction of unsold clothing and footwear as part of a broader policy to reduce the environmental footprint of consumer products (European Union, 2024).

Strategic pricing and inventory decisions by firms can exacerbate the problem of unsold stock. For instance, before the cease-and-desist order from the Japan Fair Trade Commission, Seven-Eleven, Japan's largest convenience store chain, banned its franchise stores from discounting unsold food to prevent price competition (Japan Fair Trade Commission, 2010). This resale price maintenance policy significantly reduced competition and increased unsold inventory (Deneckere et al., 1996, 1997), leading to considerable waste.

In this project, we quantitatively study the equilibrium effects of environmental policies targeted at reducing waste in the context of the Japanese publishing industry. We focus particular attention on demand uncertainty and the ability of firms to commit to prices even when there is weak demand. While retail price maintenance is generally illegal in Japan, publishers are exempted from this ban, allowing them direct control over retail prices. This implies that books that turn out to be less popular than expected are not sold and are discarded. Consequently, 35-40% of the books printed in Japan are discarded each year (Tohan Corporation, 2024).

Specifically, we ask how much of the unsold stock are due to price commitment, how large the environmental costs are compared to the pro-competitive effects of price commitment, how much of the unsold stock can be reduced by policy interventions including a deposit-refund (DR) tax system (Fullerton and Wolverton, 2000), and how demand uncertainty is related to the size of the unsold stock.

To address these issues, we first build a dynamic oligopoly pricing model with aggregate demand uncertainty. The model features stockouts and unsold stock on the equilibrium path, departing from much of the previous empirical work on oligopoly pricing that assumes there is no demand uncertainty (Bresnahan, 1989; Nevo, 2001; Ackerberg et al., 2007). We

then estimate the model using data from the Japanese publishing industry. The data include monthly deliveries, sales, and returns of each book from each bookstore, enabling the identification of key structural parameters. Finally, we conduct counterfactual analyses to evaluate the environmental policy.

This project aims to contribute to the environmental economics of waste management (Porter, 2002) by providing a quantitative analysis of the environmental consequences of firms' strategic pricing and inventory decisions on waste. Environmental economics has investigated the optimal tax for reducing waste (Dinan, 1993; Palmer and Walls, 1997; Calcott and Walls, 2000; Fullerton and Wolverton, 2000, 2005; Ino, 2011). This project highlights the tension between non-environmental policies regarding price commitment and environmental policies regarding waste management in terms of consumer and producer surpluses and environmental costs.

This project extends the empirical analysis of oligopoly price competition in multiple dimensions. First, it incorporates demand uncertainty on the supply side and allows for equilibrium stockouts and unsold stock. Second, it highlights the strategic use of price commitment under demand uncertainty. Third, it considers environmental costs as a negative externality in addition to the allocative inefficiency of firms' strategic behavior. The related literature includes Jeon (2022), which studied capacity building under demand uncertainty in container shipping, but did not consider price commitment or any environmental consequences. Kawaguchi et al. (2022) studied the same setting of the Japanese publishing industry but focused on the vertical relationship between bookstores and publishers and the competitive effects of resale price maintenance.

2 Institution

Japan is a suitable setting for studying RPM because it has a large physical-book market, a stable institutional framework for fixed resale prices, and a quantitatively important brick-and-mortar sector during our 2015–2017 sample period. In 2015, book sales totaled 742 billion yen, and average per-capita expenditure on books was 5,838 yen. Traditional bookstores accounted for 64.6 percent of book and magazine sales, compared with 9.6 percent online (Japan Book Publishers Association, 2017). Throughout the paper, a book refers to a new physical copy, sold primarily through traditional bookstores.

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

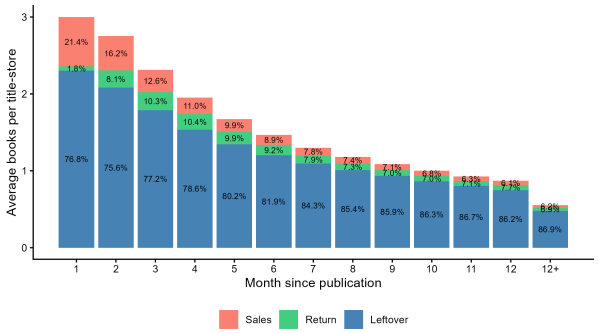
publishers and wholesalers and another between wholesalers and bookstores. RPM does not apply to e-books or used books.

The second distinctive institutional feature is that bookstores have limited control over inventory and assortment. Book wholesalers determine the allocation of new titles and quantities across bookstores with which they have exclusive supply contracts, and they bear the costs of transporting copies from publishers, warehousing them, and delivering them to bookstores. The bookstore's and the wholesaler's interests are aligned by revenue sharing. For each sale, bookstores typically retain 22 percent of the fixed retail price, wholesalers receive 8 percent, and the remaining 70 percent goes to publishers and authors. Unsold books can ordinarily be returned to the wholesaler within a specified period, usually six months, at the original transaction price. This makes price commitment easier to sustain, because retailers can hold the fixed price without bearing the full cost of overstocking. The publisher-wholesaler side effectively controls price, inventory allocation, and the return decision.

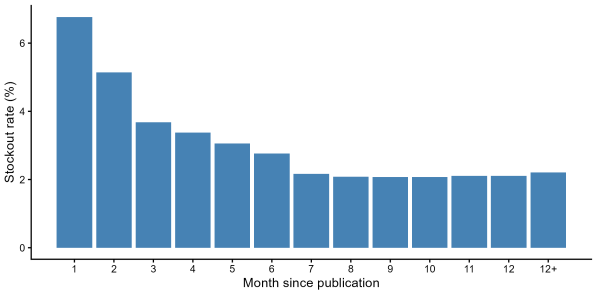
Figure 1a shows how sales, returns, and leftover at the end of the period evolve at an average title-store after publication, based on a 40-municipality national sample drawn with probability proportional to population, covering 2015 to 2017. The bar height equals the start-of-period inventory, the stock on the shelf after delivery, which falls from about 3.0 copies in the first month to about 0.9 by month twelve. Each bar decomposes into three segments. Sales account for about 21 percent of the bar in the first month and gradually decline to about 6 percent by month twelve. Returns are negligible in the first month, rise sharply from the second month, and peak in the fourth month at about 10 percent of the bar. Leftover is the bottom segment and the largest share, falling in absolute terms from 2.3 copies to 0.75 but rising as a share from 77 percent to 87 percent as the title ages. After the first half year, the sales and return shares stabilize at around 7 percent each.

Figure 1b shows the stockout rate by month since publication. A stockout occurs when a title has positive inventory at the start of the month and sells every copy before the month ends. The rate peaks at 6.8 percent in the first month and falls to about 2 percent by month seven, so stockouts are most common when demand is highest relative to the initial delivery.

Together, the two figures reveal the demand uncertainty that the model builds in. Returns arise when demand is lower than expected, and stockouts arise when demand is higher than expected. Because the retail price is fixed under RPM, the publisher cannot adjust the price to resolve the mismatch between demand and inventory. Instead, it adjusts inventory through returns and delivery as sales reveal demand over time, and sets the release price taking this adjustment into account, which makes the setting suitable for studying the effect of price commitment under demand uncertainty.



(a) Sales, returns, and leftover by month since publication



(b) Stockout rate by month since publication

Stockouts also affect the demand side through stockout-based substitution. Reduced-form regressions in Appendix A show that a title that stocked out last month sells more the following month (intertemporal substitution) and that a stockout of one title raises sales of other titles, especially within the same category (intertitle substitution). A stockout in one store raises sales in other stores in the same municipality, but this interstore association disappears once title-month demand shocks are controlled for. The model builds in the first two margins and abstracts away from the third by aggregating bookstores to the municipality and taking the municipality as the market, which the weak interstore substitution justifies. The intertemporal margin operates through old consumers who remain in the market when they do not purchase, so demand spills across periods. The intertitle margin operates through the choice set: once a title stocks out, it drops from the choice sets of consumers who arrive afterwards, who substitute to the remaining titles within the same category.

3 Model

3.1 Setting

Consider a market with publisher competition. Time is discrete, infinite, and indexed by $t = 1, \cdots, \infty$. We consider a stationary environment in which the realization of exogenous variables is the same across periods. We consider this as an approximation of a large market with many consumers and publishers.

In period t, a finite set of new titles $\mathcal{L}^{new}$ is released. Publishers are single-product firms: each title j is owned by a distinct publisher, and the publisher controls only that title over its two-period life cycle. A title is sold for two periods and then discontinued at the end of the second period. The new titles become old titles $\mathcal{L}^{old}$ in the next period. Thus, the collection of titles *listed* in period t is $\mathcal{L} \equiv \mathcal{L}^{new} \cup \mathcal{L}^{old}$. The characteristics of title j, x_j,

is the same across periods, but consumers draw different preference shocks to them because they are different titles.

Publishers face demand uncertainty about title quality ξ_j. At the beginning of period t, before choosing deliveries and (for new titles) release prices, the publisher holds a prior belief

$$\xi_j \sim N(\mu_\xi, \sigma_\xi^2).$$

The publisher observes a private signal

$$s_j = \xi_j + \nu_j, \quad \nu_j \sim N(0, \sigma_s^2),$$

and the implied posterior distribution is $\xi_j \mid s_j \sim N(\mu_{\xi|s}, \sigma_{\xi|s}^2)$, where

$$\mu_{\xi|s} = \frac{\sigma_s^2}{\sigma_\xi^2 + \sigma_s^2}\mu_\xi + \frac{\sigma_\xi^2}{\sigma_\xi^2 + \sigma_s^2}s_j, \quad \sigma_{\xi|s}^2 = \frac{\sigma_\xi^2\sigma_s^2}{\sigma_\xi^2 + \sigma_s^2}.$$

Because of the stationary assumption under the large market, the singal and the posterior distribution are also the same across periods for the same book index j.

Consumers are divided into several segments by their mean indirect utility to the titles. Let $\mathcal{Z}$ denote the set of consumer segments. Every period n_{zm}^{new} consumers of segment z arrive at market m. Consumers leave the market if they purchase a title, whle staying in the market if they do not purchase any title. The number of old consumers of segment z at the beginning of period t in market m is denoted by n_{zmt}^{old}. Let $n_{mt}^{old} \equiv \sum_{z\in\mathcal{Z}} n_{zmt}^{old}$ and $n_{mt}^{new} \equiv \sum_{z\in\mathcal{Z}} n_{zmt}^{new}$. The total number of consumers in market m at the beginning of period t is denoted by $n_{mt} \equiv n_{mt}^{new} + n_{mt}^{old}$.

Publishers set their prices at the time of release and maintain prices throughout the titles' lifecycle. Thus, the price of a title is the same across periods and markets. On the other hand, the publisher can adjust the number of copies delivered to the market. Let p_{jt} be the price of title $j \in \mathcal{L}$. Let d_{jmt} be the number of copies delivered to market m at the beginning of period t and q_{jmt} be the resulting inventory of title j at the beginning of period t in market m.

The indirect utility for consumer i of segment $z(i)$ from purchasing available book $j \in \mathcal{A}_{mt}(i)$ is

$$u_{ij} \equiv \overline{u}_{z(i)j} + \epsilon_{ij} \equiv \boldsymbol{\beta}'_{z(i)}\boldsymbol{x}_j + \boldsymbol{\gamma}_j\boldsymbol{w}_m + \xi_j - \alpha_{z(i)}p_j + \epsilon_{ij}, \tag{1}$$

where $\boldsymbol{\gamma}_j$ is the vector of title-specific coefficients on the demographic covariates $\boldsymbol{w}_m$ of market m. The shock ϵ_{ij} is drawn from an i.i.d. type-I extreme value distribution. Here, we assume that a consumer's preference shock to a title is fixed throughout the periods.

Importantly, as we discuss in the following section, we assume that the preference shocks are observed by the consumer only when they become available for the consumer.

Given the price of titles $\boldsymbol{p}_t$, the number of inventories $\boldsymbol{q}_{mt}$, and the number of consumers $\boldsymbol{n}_{mt}^{old}$, the sales of title j by consumers of segment z in market m in period t is given by s_{jzmt}. The sales function is determined by a consumer arrival process described in the next section. The sales function determines the transition of the state as follows:

$$q_{jm,t+1} = d_{jmt} - \sum_{z \in \mathcal{Z}} s_{jzmt} + d_{jm,t+1}, \tag{2}$$

and

$$n_{zm,t+1}^{old} = n_{zmt}^{old} + n_{zm}^{new} - \sum_{j \in \mathcal{L}} s_{jzmt}. \tag{3}$$

3.2 Consumer Arrival Process with Stockout

Suppose that consumers sequentially arrive at bookstores in random order. The consumer can purchase the title only when the inventory is positive. Let $q_{mt}(i)$ be the inventory of titles available when consumer i arrives. Let $\mathcal{A}_{mt}(i)$ be the set of titles available for purchase when consumer i arrives, that is, $\mathcal{A}_{mt}(i) \equiv \mathcal{A}_{mt}^{new}(i) \cup \mathcal{A}_{mt}^{old}(i)$ for $\mathcal{A}_{mt}^{\tau}(i) \equiv \{j \in \mathcal{L}^{\tau} : q_{mt}(i) > 0\}$ for $\tau = \{new, old\}$. Let $F_{\mathcal{A}}$ be the distribution of the previous period's availability set $\mathcal{A}_{m,t-1}(i)$. Let $s_{zjmt}(\mathcal{A}_t(i), \mathcal{A}_{t-1}^{new}(i))$ be the conditional choice probability of consumer i of segment z for title j given the availability set $\mathcal{A}_{mt}(i)$ and the availability set $\mathcal{A}_{m,t-1}^{new}(i)$ in the previous period, which is characterized in the next section.

The random order defines a joint distribution of consumer characteristics $(z_i, \mathcal{A}_{m,t-1}(i))$ over the order of consumers $[0, n_{mt}]$, which also defines the marginal density of consumer characteristics of i-th visitor denoted by $f_{z\mathcal{A}}(z_i, \mathcal{A}_{imt})$, which is independent of the history up to i-th consumer becuase the ordering is random. Thus, the evolution of the inventory of book j is

$$\frac{dq_{jmt}(i)}{di} = \begin{cases} -\sum_z \sum_{\mathcal{A}_{m,t-1}} f_{z\mathcal{A}}(z, \mathcal{A}_{m,t-1}) s_{zj}(\mathcal{A}_{mt}(i), \mathcal{A}_{m,t-1}^{new}(i)) & \text{if } q_{jmt}(i) > 0, \\ 0 & \text{if } q_{jmt}(i) = 0, \end{cases} \tag{4}$$

while the evolution of sales of book j is given by

$$\frac{ds_{jzmt}(i)}{di} = \sum_{\mathcal{A}_{m,t-1}} f_{z\mathcal{A}}(z, \mathcal{A}_{m,t-1}) s_{zj}(\mathcal{A}_{mt}(i), \mathcal{A}_{m,t-1}^{new}(i)). \tag{5}$$

Consequently, the book j's inventory at the end of the period and the sales during the

period are given by

$$q_{m,t+1} = q_{mt} + \int_0^{n_{mt}} \frac{dq_{imt}(i)}{di} di, \tag{6}$$

and

$$s_{jzmt} = \int_0^{n_{mt}} \frac{ds_{jzmt}(i)}{di} di. \tag{7}$$

Note that this non-stochastically determines the path of inventory and cumulative sales for i-th consumer, $q_{mt}(i)$. Therefore, inventory available for a consumer follows $q_{jmt}(U)$ for $U \in U[0, n_{mt}]$. This induces the distribution of availability set $F_{\mathcal{A}}$ by

$$F_{\mathcal{A}}(\mathcal{A}_{mt}) = \mathbb{P}\{q_{jmt}(U) > 0, \forall j \in \mathcal{A}_{mt}, q_{kmt}(U) = 0, \forall k \neq \mathcal{A}_{mt}\}. \tag{8}$$

We now illustrate how the marginal density of consumer characteristics $f_{z\mathcal{A}}(z, \mathcal{A}_{m,t})$ is determined in $t+1$. The consumer with an empty previous availability set must be a new consumer, representing $\frac{n^{new}_{zm,t+1}}{n_{m,t+1}}$ of the population.[1] If the previous availability set is non-empty, the consumer must be an old consumer and face availability set $\mathcal{A}_{mt}$ in the support of $F_{\mathcal{A}}(\mathcal{A}_{mt})$. The probability that an old consumer has characteristics z and $\mathcal{A}_{mt}$ is

$$p\left(z, \mathcal{A}_{mt}\right) \equiv \frac{\sum_{\mathcal{A}_{m,t-1}} f_{z\mathcal{A}}(z, \mathcal{A}_{m,t-1}) F_{\mathcal{A}}(\mathcal{A}_{mt}) \left(1 - \sum_{j \in \mathcal{A}_{mt}} s_{zj}(\mathcal{A}_{mt}, \mathcal{A}^{new}_{m,t-1})\right)}{\sum_{\mathcal{A}_{mt}} \sum_z \sum_{\mathcal{A}_{m,t-1}} f_{z\mathcal{A}}(z, \mathcal{A}_{m,t-1}) F_{\mathcal{A}}(\mathcal{A}_{mt}) \left(1 - \sum_{j \in \mathcal{A}_{mt}} s_{zj}(\mathcal{A}_{mt}, \mathcal{A}^{new}_{m,t-1})\right)}. \tag{9}$$

The numerator is the probability that a randomly chosen consumer has characteristics z, encounters the availability set $\mathcal{A}_{mt}$, and survives until $t+1$. To see why, note that the numerator can be decomposed into three parts. First, a randomly chosen consumer has characteristics $(z, \mathcal{A}_{m,t-1})$ with probability $f_{z\mathcal{A}}(z, \mathcal{A}_{m,t-1})$ in t. Second, conditional on these characteristics, the consumer faces an availability $\mathcal{A}_{mt}$ with probability $F_{\mathcal{A}}(\mathcal{A}_{mt})$. The probability that a consumer encounters $\mathcal{A}_{mt}$ is independent of her characteristics due to the large market assumption of consumers. Third, conditional on both the consumer characteristics and the availability set, the consumer waits until the next period with probability $1 - \sum_{j \in \mathcal{A}_{mt}} s_{zj}(\mathcal{A}_{mt}, \mathcal{A}^{new}_{m,t-1})$. The denominator aggregates over z and the availability set encountered in t, giving the probability that a randomly chosen consumer survives until $t+1$.

[1] This assumes that no bookstore stocks out entirely.

To summarize, the distribution of consumer characteristics at $t+1$ is

$$f_{z\mathcal{A}}(z,\mathcal{A}_{mt}) = \begin{cases} \frac{n^{new}_{zm,t+1}}{n_{m,t+1}} & \text{, if } \mathcal{A}_{mt} = \emptyset \\ \frac{n^{old}_{m,t+1}}{n_{m,t+1}} p(z,\mathcal{A}_{mt}) & \text{, if } F(\mathcal{A}_{mt}) > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{10}$$

Using (9) and (10), $f_{z\mathcal{A}}(z,\mathcal{A}_{mt})$ can be updated from $f_{z\mathcal{A}}(z,\mathcal{A}_{m,t-1})$, individual conditional choice probabilities $s_{zj}(\mathcal{A}_{mt},\mathcal{A}_{m,t-1})$, $F_{\mathcal{A}}(\mathcal{A}_{mt})$, and the mass of consumer segment in $t+1$ $\left\{\left\{n^{\tau}_{zmt+1}\right\}_{z\in\mathcal{Z}}\right\}_{\tau\in\{new,old\}}$.

3.3 Timing of the Game

In each period, the game unfolds as follows:

1. New consumers arrive in the market and publishers introduce new titles.
2. Publishers observe their information (including the signal s_j for new titles) and simultaneously choose:
   - Release prices $\{p_j\}_{j\in\mathcal{L}^{new}}$ (which are then maintained over the title's life cycle),
   - Deliveries $\{d_{jmt}\}_{j\in\mathcal{L},m\in\mathcal{M}}$ at the beginning of period t.
3. Consumers arrive at bookstores sequentially in random order. For each consumer:
   - Nature draws and reveals the idiosyncratic preference shocks for titles that are available
   - The consumer makes a purchase decision subject to the current inventory
   - If a purchase is made, inventory is updated before the next consumer arrives
   - Consumers who made purchases exit the market.
4. At the end of the period, new titles become old titles and old titles in their second period are discontinued (with remaining copies returned to publishers). Consumers who did not purchase remain as old consumers.

4 Consumer Equilibrium in the Steady State

4.1 Bellman Equations

In the following, we derive the equilibrium of consumers in the market m given the supply-side policy and omit the subscript m for simplicity. Let $V_z(\{u_{ij}\}_{j\in\mathcal{A}^{new}_{i,t-1}}, \boldsymbol{p}_t, \boldsymbol{q}_t, \boldsymbol{n}^{old}_t)$ be the

value function for old consumers of segment z at the beginning of period t. Then, the value function for a new consumer is $V_z(\emptyset, \boldsymbol{p}_t, \boldsymbol{q}_t, \boldsymbol{n}_t^{old})$. The value function of old consumers depends on the availability set $\mathcal{A}_{i,t-1}^{new}$ and their preference shocks to the available titles because the idiosyncratic preference shock to the available titles has been realized.

The old consumer's intra-period problem who has drawn an availability set $\mathcal{A}_{i,t-1}$ last period and $\mathcal{A}_{it}$ this period is

$$W_z(\{u_{ij}\}_{j\in\mathcal{A}_{it}}, \boldsymbol{p}_t, \boldsymbol{q}_t, \boldsymbol{n}_t^{old}) \equiv \max\left\{\max_{j\in\mathcal{A}_{it}} u_{ij}, \delta V_z(\{u_{ij}\}_{j\in\mathcal{A}_{it}^{new}}, \boldsymbol{p}_{t+1}, \boldsymbol{q}_{t+1}, \boldsymbol{n}_{t+1}^{old})\right\} \tag{11}$$

where $\boldsymbol{p}_{t+1}, \boldsymbol{q}_{t+1}, \boldsymbol{n}_{t+1}^{old}$ follow the transition rule described in the previous section.

It is related to $V_z(\{u_{ij}\}_{j\in\mathcal{A}_{i,t-1}^{new}}, \boldsymbol{p}_t, \boldsymbol{q}_t, \boldsymbol{n}_t^{old})$ by

$$\begin{aligned} &V_z(\{u_{ij}\}_{j\in\mathcal{A}_{i,t-1}^{new}}, \boldsymbol{p}_t, \boldsymbol{q}_t, \boldsymbol{n}_t^{old}) \\ &= \int_{\mathcal{A}_{it}} \int_{\{u_{ij}\}_{j\in\mathcal{A}_{it}\setminus\mathcal{A}_{i,t-1}^{new}}} W_z(\{u_{ij}\}_{j\in\mathcal{A}_{it}}) dF_{\{u_{ij}\}_{j\in\mathcal{A}_{it}\setminus\mathcal{A}_{i,t-1}^{new}}} dF_{\mathcal{A}}(\mathcal{A}_{it}, \boldsymbol{p}_t, \boldsymbol{q}_t, \boldsymbol{n}_t^{old}). \end{aligned} \tag{12}$$

4.2 Steady State with a Bounded Perception

We focus on the steady state equilibrium in which the endogenous state variables are time invariant, thereby dropping the dependence of the value functions on $\boldsymbol{p}_t, \boldsymbol{q}_t, \boldsymbol{n}_t^{old}$.

Moreover, we assume that consumers approximate the value function $V_z(\{u_{ij}\}_{j\in\mathcal{A}_{i,t-1}^{new}})$ with $\overline{V}_z(\bar{u}_i\left(\mathcal{A}_{i,t-1}^{new}\right))$ a function of the maximum utility of available new titles in period $t-1$, $\bar{u}_i(\mathcal{A}_{i,t-1}^{new}) \equiv \max_{j\in\mathcal{A}_{i,t-1}^{new}} u_{ij}$, and that consumers anticipate that the next availability titles are the current available titles plus some newly listed titles.

Under this assumption, the perceived intra-period value function of old consumers becomes

$$\begin{aligned} &\max\left\{\max_{j\in\mathcal{A}_{it}} u_{ij}, \delta\overline{V}_z\left(\bar{u}_i(\mathcal{A}_{it}^{new})\right)\right\} \\ &= \max\{\bar{u}_i(\mathcal{A}_{it}^{new}), \bar{u}_i(\mathcal{A}_{i,t-1}^{new}), \delta\overline{V}_z\left(\bar{u}_i(\mathcal{A}_{it}^{new})\right)\} \end{aligned} \tag{13}$$

when $\mathcal{A}_{it}$ is of the form $\mathcal{A}_{it} = \mathcal{A}_{it}^{new} \cup \mathcal{A}_{i,t-1}^{new}$. Otherwise, the perceived intra-period value function is defined as in the first line of (13).

Under this assumption, the perceived value function of old consumers is defined by

$$\begin{aligned} &\overline{V}_z(\bar{u}_i(\mathcal{A}_{i,t-1})) \\ &= \int_{\mathcal{A}_{it}^{new}} \int_{\{u_{ij}\}_{j\in\mathcal{A}_{it}^{new}}} \max\left\{\bar{u}_i(\mathcal{A}_{it}^{new}), \bar{u}_i(\mathcal{A}_{i,t-1}^{new}), \delta\overline{V}_z\left(\bar{u}_i(\mathcal{A}_{it}^{new})\right)\right\} dF_{\{u_{ij}\}_{j\in\mathcal{A}_{it}^{new}}} dF_{\mathcal{A}}(\mathcal{A}_{it}). \end{aligned} \tag{14}$$

Proposition 1 *The following properties hold for* $\delta\overline{V}_z(x)$*:*

1. $\frac{\partial}{\partial x}\left(\delta\overline{V}_z(x)\right) \in [0,1)$ *and* $\delta\overline{V}_z(x)$ *has a unique fixed piont* $x_z^* = \delta\overline{V}_z(x_z^*)$ *such that* $x \geq x_z^*$ *if and only if* $x \geq \delta\overline{V}_z(x)$*.*

2. *The fixed point satisfies*

$$x_z^* = \frac{\delta\mu\rho}{1-\delta(1-\rho)} \tag{15}$$

for

$$\rho \equiv \mathbb{P}\{\bar{u}_z(\mathcal{A}_{it}^{new}) \geq x_z^*\}, \tag{16}$$

$$\mu \equiv \mathbb{E}\{\bar{u}_z(\mathcal{A}_{it}^{new}) | \bar{u}_z(\mathcal{A}_{it}^{new}) \geq x_z^*\}. \tag{17}$$

3. $\delta\overline{V}_z(x) \geq x_z^*$ *for any* x*. Hence, for* $x \leq x_z^*$*,* $\delta\overline{V}_z(x) = x_z^*$*.*

Proof. We show the first property. Note that the derivative of $\overline{V}_z(x)$ is equal to the probability that a title in $\mathcal{A}_{i,t-1}^{new}$ is chosen. To see this, note that

$$\begin{aligned}\overline{V}_z(x) &= \int_{-\infty}^{\infty} \left(x' \vee x \vee \delta\overline{V}_z(x')\right) dF(x') \\ &= \int_{-\infty}^{x} \left(x \vee \delta\overline{V}_z(x')\right) dF(x') + \int_{x}^{\infty} \left(x' \vee \delta\overline{V}_z(x')\right) dF(x')\end{aligned} \tag{18}$$

where $F(x')$ is the distribution of $\bar{u}_i(\mathcal{A}_{it}^{new})$ implied from $dF_{\{u_{ij}\}_{j\in\mathcal{A}_{it}^{new}}} dF_{\mathcal{A}}(\mathcal{A}_{it})$ and the preference shocks. Therefore, writing f for the density of F, the derivative of $\overline{V}_z(x)$ with respect to x is

$$\begin{aligned}\frac{d}{dx}\overline{V}_z(x) &= \left(x \vee \delta\overline{V}_z(x)\right) f(x) + \int_{-\infty}^{x} 1\{\delta\overline{V}_z(x') < x\} dF(x') - \left(x \vee \delta\overline{V}_z(x)\right) f(x) \\ &= \mathbb{P}\{x' < x, \delta\overline{V}_z(x') < x\}.\end{aligned} \tag{19}$$

This implies that, $\delta\overline{V}_z(x)$ has slope at most $\delta < 1$, that is, bounded below 1. Additionally, $\lim_{x\to-\infty} x - \delta\overline{V}_z(x) < 0$ because in the next period, there is a chance to find better titles when the current available titles are extremely bad.

This then implies that $\delta\overline{V}_z(x)$ has a unique fixed point in x, which we denote by x_z^*.

We show the second property. We have

$$\begin{aligned}x_z^* &= \delta\int_{-\infty}^{x_z^*} \left(x_z^* \vee \delta\overline{V}_z(x')\right) dF(x') + \delta\int_{x_z^*}^{\infty} \left(x' \vee \delta\overline{V}_z(x')\right) dF(x') \\ &= \delta\int_{-\infty}^{x_z^*} x_z^* dF(x') + \delta\int_{x_z^*}^{\infty} x' dF(x') \\ &= \delta x_z^* \mathbb{P}\{x' \leq x_z^*\} + \delta\mathbb{E}\{x' | x' \geq x_z^*\}\mathbb{P}\{x' > x_z^*\}.\end{aligned} \tag{20}$$

By arranging the terms, we obtain the second property.

We show the third property. Because $\delta\overline{V}_z(x)$ is the discounted sum of the futurer payoffs under the optimal policy, it is no less than the discounted sum of the future payoffs under an ad hoc policy. Consider a policy such that buying from A_{it}^{new} if $x' \geq x_z^*$ and wait for next period if $x' \leq x_z^*$ and never buys from $A_{i,t-1}^{new}$.

The discounted sum of future payoffs under this policy is

$$\begin{aligned} &\delta\rho\mu + \delta^2(1-\rho)\rho\mu + \delta^3(1-\rho)^2\rho\mu\cdots \\ &= \delta\rho\mu\sum_{k=0}^{\infty}\delta^k(1-\rho)^k \\ &= \frac{\delta\rho\mu}{1-\delta(1-\rho)} \\ &= x_z^*. \end{aligned} \tag{21}$$

Therefore,

$$\delta\overline{V}_z(x) \geq x_z^*. \tag{22}$$

■

With Proposition 1, the old consumers optimal decision is to wait for next period if $\bar{u}_z(\mathcal{A}_{it}^{new}) \leq x_z^*$ and buy $\text{argmax}_{j\in\mathcal{A}_{it}^{new}} u_{ij}$ if $\bar{u}_z(\mathcal{A}_{it}^{new}) > x_z^*$. Therefore, we only need to solve for x_z^* to characterize the consumer's optimal decision.[2]

We can relax the assumption that the consumer perceives that products in $\mathcal{A}_{i,t-1}^{new}$ are all available in next period, that is, some of them can be unavailable in next period.

Let $\widetilde{V}_z(\mathcal{A}_{i,t-1}^{new})$ be the value function under this perception. Because this is the value function under the perception that the choice set is no greater than the previous case, we have:

$$\widetilde{V}_z(\mathcal{A}_{i,t-1}^{new}) \leq \overline{V}_z(\bar{u}_z^{old}(\mathcal{A}_{i,t-1}^{new})). \tag{23}$$

Moreover, because the above strategy is still feasible and the future payoffs of the strategy does not depend on the realization of the utility of products in $\mathcal{A}_{i,t-1}^{new}$, the value of this strategy is evaluated the same under the perception that some of the products in $\mathcal{A}_{i,t-1}^{new}$ are unatailable in next period as under the perception that all products in it are available. This implies that

$$\delta\widetilde{V}_z(\mathcal{A}_{i,t-1}^{new}) \geq x_z^*. \tag{24}$$

[2]This requires that the consumer does not buy old titles published in the last period. That may fail when the second part of the assumption breaks. If a title stocked out last period but is restocked this period, the current availability set contains old titles that were not in the previous availability set. The consumer may then buy those old titles, and the CCP is no longer characterized by this cutoff rule.

Therefore, the consumer's optimal decision is to wait for next period if $\bar{u}_z(\mathcal{A}_{it}^{new}) \leq x_z^*$ because $x_z^* \leq \delta \widetilde{V}_z(\mathcal{A}_{i,t-1}^{new})$ and buy $\text{argmax}_{j \in \mathcal{A}_{it}^{new}} u_{ij}$ if $\bar{u}_z(\mathcal{A}_{it}^{new}) > x_z^*$ because the latter implies that $\bar{u}_z(\mathcal{A}_{it}^{new}) > \overline{V}_z(\bar{u}_z^{old}(\mathcal{A}_{i,t-1}^{new})) \geq \widetilde{V}_z(\mathcal{A}_{i,t-1}^{new})$. This shows that the consumer's optimal decision is still characterized by $\sigma_z(x_z^*)$ even under the slightly modified bounded perception.

4.3 Conditional Choice Probability

The conditional choice probability for a consumer of segment z of title $j \in \mathcal{A}_{it}$ given x_z^* is characterized as follows.

$$\begin{aligned}
s_{zj}(\mathcal{A}_{it}, \mathcal{A}_{i,t-1}^{new}) &= \Pr(u_{ij} \geq x_z^*, u_{ij} \geq u_{ik}, \forall k \in \mathcal{A}_{it}) \\
&= \Pr(\epsilon_{ij} \geq x_z^* - \overline{u}_{zj}, \epsilon_{ik} \leq \epsilon_{ij} + \overline{u}_{zj} - \overline{u}_{zk}, \forall k \neq j \in \mathcal{A}_{it}) \\
&= \Pr(\epsilon_{ij} \geq x_z^* - \overline{u}_{zj}, \epsilon_{ik} \leq \epsilon_{ij} + \overline{u}_{zj} - \overline{u}_{zk}, \forall k \neq j \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}) \\
&= \int_{x_z^* - \overline{u}_{zj}}^{\infty} \prod_{k \neq j \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}} \exp\left(-e^{-(x + \overline{u}_{zj} - \overline{u}_{zk})}\right) e^{-x} \exp(-e^{-x}) dx \\
&= \int_{x_z^* - \overline{u}_{zj}}^{\infty} \exp\left(- \sum_{k \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}} e^{-(x + \overline{u}_{zj} - \overline{u}_{zk})} \right) e^{-x} dx.
\end{aligned}$$

By defining $t = e^{-x}$, the expression becomes

$$\begin{aligned}
s_{zj}(\mathcal{A}_{it}, \mathcal{A}_{i,t-1}^{new}) &= \int_0^{\exp(-(x_z^* - \overline{u}_{zj}))} \exp\left(-t \sum_{k \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}} e^{-(\overline{u}_{zj} - \overline{u}_{zk})} \right) dt \\
&= \frac{1}{-\sum_{k \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}} e^{-(\overline{u}_{zj} - \overline{u}_{zk})}} \exp\left(-t \sum_{k \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}} e^{-(\overline{u}_{zj} - \overline{u}_{zk})} \right) \Bigg|_0^{\exp(-(x_z^* - \overline{u}_{zj}))} \\
&= \left(1 - \exp\left(-e^{-x_z^*} \sum_{k \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}} e^{\overline{u}_{zk}} \right) \right) \frac{e^{\overline{u}_{zj}}}{\sum_{k \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}} e^{\overline{u}_{zk}}}.
\end{aligned}$$

Thus, the individual conditional choice probability for a consumer is

$$s_{zj}(\mathcal{A}_{it}, \mathcal{A}_{i,t-1}^{new}) = \begin{cases} \left(1 - \exp\left(- e^{-x_z^*} H_{it} \right) \right) \frac{e^{\overline{u}_{zj}}}{H_{it}} & \text{, if } j \in \mathcal{A}_{it} \setminus \mathcal{A}_{i,t-1}^{new}, \\ 0 & \text{, if } j \in \mathcal{A}_{i,t-1}^{new}, \end{cases} \tag{25}$$

where $H_{it} \equiv \sum_{j \in \mathcal{A}_{it} \setminus \mathcal{A}^{new}_{i,t-1}} e^{\overline{u}_{zj}}$.

4.4 Definition of the Steady-State Consumer Equilibrium

Given the price $\boldsymbol{p}^*$ and inventory $\boldsymbol{q}^*$, the steady-state consumer equilibrium in a market is a tuple of the mass of consumer segments $\boldsymbol{n}^* \equiv \{\{n^\tau_z\}_{z \in \mathcal{Z}}\}_{\tau \in \{new, old\}}$, aggregate shares $\boldsymbol{s}^* \equiv \{s_{jzm}\}_{z \in \mathcal{Z}, j \in \mathcal{L}, m \in \mathcal{M}}$, the marginal density of consumer characteristics $f^*_{z\mathcal{A}}$, individual conditional choice probabilities $\boldsymbol{s}^*(\mathcal{A}', \mathcal{A}) \equiv \{\{s_{zj}(\mathcal{A}', \mathcal{A})\}_{j \in \mathcal{L}}\}_{z \in \mathcal{Z}}$, cutoff values $\boldsymbol{x}^* \equiv \{x^*_z\}_{z \in \mathcal{Z}}$, the distribution of availability set $F_{\mathcal{A}}$, and intra-period inventory path $\boldsymbol{q}^*(i)$ such that:

Aggregate shares are defined as the steady-state counterpart of (7). That is, for each $j \in \mathcal{L}$, $z \in \mathcal{Z}$, and $m \in \mathcal{M}$,

$$s^*_{jzm} \equiv \int_0^{n^*_m} \frac{ds^*_{jzm}(i)}{di} di,$$

where $n^*_m \equiv \sum_{z \in \mathcal{Z}} (n^{new}_{zm} + n^{old*}_{zm})$ and $\frac{ds^*_{jzm}(i)}{di}$ follows the steady-state counterpart of (5).

1. The individual conditional choice probabilites $\boldsymbol{s}^*(\mathcal{A}', \mathcal{A})$ are determined from the cuttoff values $\boldsymbol{x}^*$ in (25);

2. The cutoff values $\boldsymbol{x}^*$ are the fixed point of the perceived continuation value calculated in (13) that is calculated with the distribution of availability set $F_{\mathcal{A}}$;

3. The distribution of availability set $F_{\mathcal{A}}$ is determined by intra-period inventory path $\boldsymbol{q}^*(i)$ in (8);

4. The intra-period inventory path $\boldsymbol{q}^*(i)$ is determined by individual conditional choice probabilities $\boldsymbol{s}^*(\mathcal{A}', \mathcal{A})$, the distribution of consumer characteristics $f^*_{z\mathcal{A}}$ and $\boldsymbol{n}^*$ in (4) and (6);

5. The marginal density of consumer characteristics $f^*_{z\mathcal{A}}$ is determined by solving the fixed point of $f_{z\mathcal{A}}(z, \mathcal{A})$ in (10) using individual conditional choice probabilities $\boldsymbol{s}^*(\mathcal{A}', \mathcal{A})$, the distribution of availability set $F^*_{\mathcal{A}}$ and the mass of consumer segments n^*.

6. The mass of consumer segments $\boldsymbol{n}^*$ and aggregate shares $\boldsymbol{s}^*$ satisfy (3).

7. Aggregate shares $\boldsymbol{s}^*$ is derived from the individual conditional choice probabilities $\boldsymbol{s}^*(\mathcal{A}', \mathcal{A})$,marginal density of consumer characteristics $f^*_{z\mathcal{A}}$, and mass of consumer segments $\boldsymbol{n}^*$ following (5) and (7).

5 Steady-State Supply-Side Equilibrium

In the steady state, we suppress the time subscript for all time-invariant objects. Prices p_j are committed at release and remain fixed over a title's two-period life cycle. Inventories q_{jm} and deliveries d_{jm} are constant across periods in steady state, where q_{jm} denotes the beginning-of-period inventory of title j in market m and d_{jm} denotes the number of copies of title j delivered to market m at the beginning of the period, for $j \in \mathcal{L}^{new} \cup \mathcal{L}^{old}$. We use superscripts "old" and "new" to separately denote variables for old and new titles whenever it's necessary to make it clear whether we refer to an old or new title.

Publishers are single-product firms. For a title j, if $j \in \mathcal{L}^{new}$, the publisher observes the signal s_j; if $j \in \mathcal{L}^{old}$, the publisher observes the realized quality ξ_j. Taking as given rivals' prices and delivery choices, the publisher chooses deliveries for its title, and if j is new, it also chooses the release price. The publisher discounts next period payoffs by δ.

Given a candidate profile of prices and inventories $(\boldsymbol{p}, \boldsymbol{q})$, let $(\boldsymbol{n}^*, \boldsymbol{s}^*, f^*_{z\mathcal{A}}, \boldsymbol{s}^*(\mathcal{A}', \mathcal{A}), \boldsymbol{x}^*, F^*_{\mathcal{A}}, \boldsymbol{q}^*(i))$ be a steady-state consumer equilibrium in the sense of the previous section, and let $\{s^*_{jzm}\}_{j,z,m}$ denote the corresponding aggregate sales. Write $S_{jm} \equiv \sum_{z \in \mathcal{Z}} s^*_{jzm}$ for total sales of title j in market m, and write S^{new}_{jm} and S^{old}_{jm} for the same object when the title is new or old. All expectations conditional on s_j are taken with respect to the posterior distribution $F_{\xi_j|s_j}$ implied by the prior and signal structure in the setting.

Given rivals' decisions $(\boldsymbol{p}_{-j}, \boldsymbol{d}_{-j})$, a single-product publisher owning title j chooses deliveries for its title, and if j is new, it also chooses the release price, taking the induced steady-state consumer equilibrium as given.

For an old title $j \in \mathcal{L}^{old}$, the publisher observes ξ_j and solves

$$\max_{\{d^{old}_{jm} \geq 0\}_{m \in \mathcal{M}}} \sum_{m \in \mathcal{M}} \left(p_j S^{old}_{jm} - mc_{jm} d^{old}_{jm} \right), \tag{26}$$

taking $(\boldsymbol{p}_{-j}, \boldsymbol{d}_{-j})$ and the induced steady-state consumer equilibrium as given.

For a new title $j \in \mathcal{L}^{new}$, the publisher observes s_j and chooses $(p_j, \{d^{new}_{jm}\}_{m \in \mathcal{M}})$ anticipating the continuation value when the title becomes old:

$$\max_{p_j,\ \{d^{new}_{jm} \geq 0\}_{m \in \mathcal{M}}} \mathbb{E}\Bigg[\sum_{m \in \mathcal{M}} \left(p_j S^{new}_{jm} - mc_{jm} d^{new}_{jm} \right) + \delta \max_{\{d^{old}_{jm} \geq 0\}_{m \in \mathcal{M}}} \sum_{m \in \mathcal{M}} \left(p_j S^{old}_{jm} - mc_{jm} d^{old}_{jm} \right) \ \Big|\ s_j \Bigg], \tag{27}$$

Here, d^{old}_{jm} is the delivery in the second period for a given realization of ξ_j. In steady state,

the inventory when the title is old is $q_{jm}^{old} = \tilde{q}_{jm} + d_{jm}^{old}$, where $\tilde{q}_{jm}$ is the carryover inventory from when the title was new, i.e., $\tilde{q}_{jm} \equiv d_{jm}^{new} - S_{jm}^{new}$.

For an old title $j \in \mathcal{L}^{old}$, given (p_j, ξ_j), the publisher chooses $\{d_{jm}^{old}\}_{m \in \mathcal{M}}$ subject to $d_{jm}^{old} \geq 0$, where d_{jm}^{old} is the delivery in the second period for a given realization of ξ_j. The optimality conditions are, for each realized ξ_j and each $m \in \mathcal{M}$,

$$p_j \frac{\partial S_{jm}^{old}}{\partial d_{jm}^{old}} - mc_{jm} \leq 0, \tag{28}$$

$$d_{jm}^{old} \geq 0, \tag{29}$$

$$\left(p_j \frac{\partial S_{jm}^{old}}{\partial d_{jm}^{old}} - mc_{jm} \right) d_{jm}^{old} = 0, \tag{30}$$

where the dependence of S_{jm}^{old} on d_{jm}^{old} is through the beginning-of-period inventory q_{jm}^{old}.

This Karush-Kuhn-Tucker (KKT) condition implies a "fill-the-gap" structure. Fix (p_j, ξ_j) and market m. By (4) and (7), total sales of title j in market m are weakly increasing in the beginning-of-period inventory q_{jm}^{old} and cannot exceed it. If the title stocks out within the period, then one additional unit of delivery increases total sales one-for-one, i.e., $\partial S_{jm}^{old} / \partial d_{jm}^{old} = 1$. If inventory is high enough that the title remains available throughout the period (no stockout), the availability sets do not change with further deliveries and total sales are locally flat in d_{jm}^{old}, i.e., $\partial S_{jm}^{old} / \partial d_{jm}^{old} = 0$. Therefore, if $p_j > mc_{jm}$, the optimal delivery in market m sets inventory just high enough to eliminate stockout:

$$d_{jm}^{old,*} = \max\{0, D_{jm}^{old}(p_j, \xi_j) - \tilde{q}_{jm}\},$$

where $D_{jm}^{old}(p_j, \xi_j)$ denotes the old title's no-stockout sales in market m, defined as the total sales the title would obtain if it were available throughout the period:

$$D_{jm}^{old}(p_j, \xi_j) \equiv \sum_{z \in \mathcal{Z}} s_{jzm}^{ns,old}(p_j, \xi_j),$$

where $\{s_{jzm}^{ns,old}(p_j, \xi_j)\}_{z \in \mathcal{Z}}$ are the steady-state aggregate sales in the counterfactual consumer equilibrium that is identical except that title j never stocks out during the period. That is, its intra-period inventory path satisfies $q_{jm}^{old}(i) > 0$ for all $i \in [0, n_m^*]$. If $p_j < mc_{jm}$, the optimal choice is $d_{jm}^{old,*} = 0$ and old-period sales equal the minimum of leftover inventory and no-stockout demand, $S_{jm}^{old} = \min\{\tilde{q}_{jm}, D_{jm}^{old}(p_j, \xi_j)\}$.[3] We analogously define $D_{jm}^{new}(p_j, \xi_j) \equiv \sum_{z \in \mathcal{Z}} s_{jzm}^{ns,new}(p_j, \xi_j)$ as the new title's no-stockout sales, where $\{s_{jzm}^{ns,new}(p_j, \xi_j)\}_{z \in \mathcal{Z}}$ are the

[3] On the equilibrium path, $p_j < mc_{jm}$ implies $d_{jm}^{new} = 0$, so $\tilde{q}_{jm} = 0$ and old-period sales are zero.

steady-state aggregate sales in the counterfactual consumer equilibrium that is identical except that the new title j never stocks out during the period.

For a new title $j \in \mathcal{L}^{new}$, given s_j, the publisher chooses $(p_j, \{d_{jm}^{new}\}_{m\in\mathcal{M}})$ anticipating that the title becomes old and deliveries when old are chosen optimally. The optimality conditions are

$$\mathbb{E}\left[\sum_{m\in\mathcal{M}}\left(S_{jm}^{new} + p_j\frac{\partial S_{jm}^{new}}{\partial p_j} + \delta S_{jm}^{old} + \delta p_j\frac{\partial S_{jm}^{old}}{\partial p_j} - \delta mc_{jm}\frac{\partial d_{jm}^{old,*}}{\partial p_j}\right)\middle| s_j\right] = 0. \tag{31}$$

$$\mathbb{E}\left[p_j\frac{\partial S_{jm}^{new}}{\partial d_{jm}^{new}} - \delta mc_{jm}\frac{\partial d_{jm}^{old,*}}{\partial d_{jm}^{new}}\middle| s_j\right] - mc_{jm} \le 0, \quad \forall m \in \mathcal{M}. \tag{32}$$

Holding $\{d_{jm}^{new}\}$ fixed, a higher p_j changes $d_{jm}^{old,*}$ through carryover and through $D_{jm}^{old}(p_j, \xi_j)$. Both enter the last term of (31). The envelope theorem does not apply because old-period profit is not locally flat in d^{old} at the fill-the-gap kink. One fewer old copy cuts sales one-for-one, so the left derivative is $\delta(p_j - mc_{jm}) > 0$. One extra old copy does not raise sales, so the right derivative is $-\delta mc_{jm} < 0$.

Let the stockout probability and the replenish probability be

$$P_{jm}^S = \mathbb{E}\left[\frac{\partial S_{jm}^{new}}{\partial d_{jm}^{new}}\middle| s_j\right], \qquad P_{jm}^R = \mathbb{E}\left[-\frac{\partial d_{jm}^{old,*}}{\partial d_{jm}^{new}}\middle| s_j\right]. \tag{33}$$

Then the delivery FOC can be written as $p_j P_{jm}^S + \delta mc_{jm} P_{jm}^R - mc_{jm} \le 0$. At the margin, extra new delivery moves posterior mass of ξ from stockout to no-stockout just across the threshold ξ^* defined by $q_{jm}^{new} = s_{jm}^{new}(\xi^*)$. Close to ξ^*, sales equal delivery, so that switch contributes zero. On each stockout realization, sales rise one-for-one with new delivery, so the publisher earns p_j whenever a stockout occurs. On the other hand, on each no-stockout realization, the extra unit is unsold and raises carryover. It reduces $d_{jm}^{old,*}$ whenever the publisher still replenishes, and the saved cost is $\delta mc_{jm} P_{jm}^R$. The equality holds only when $d_{jm}^{new} > 0$. When the inequality is strict, $d_{jm}^{new} = 0$.

A steady-state supply-side equilibrium is a collection of release prices and delivery policies such that: (i) for each old title $j \in \mathcal{L}^{old}$, deliveries satisfy (30) given (p_j, ξ_j), (ii) for each new title $j \in \mathcal{L}^{new}$, release price and deliveries satisfy (31) and (32) given s_j and taking rivals' decisions as given, and (iii) the induced consumer-side objects form a steady-state consumer equilibrium under the resulting $(\boldsymbol{p}, \boldsymbol{q})$ as defined in the model. In steady state, all time-invariant objects (prices, inventories, deliveries) are constant across periods.

6 Computation

6.1 Consumer Equilibrium

Equilibrium conditions 5-6 require stationarity of the marginal density of consumer characteristics $f_{z\mathcal{A}}$ and mass of consumer segment $\boldsymbol{n}$ for a given cutoff value $\boldsymbol{x}$, and distribution of availability set $F_{\mathcal{A}}$. To simplify the equilibrium conditions, it is useful to define the mass of consumer characteristics as a new object, denoted by $n_{z\mathcal{A}}(z, \mathcal{A}_{mt}) \equiv n_{m,t+1} f_{z\mathcal{A}}(z, \mathcal{A}_{mt})$. Because there is a one-to-one mapping from $f_{z\mathcal{A}}(z, \mathcal{A}_{mt})$, $\boldsymbol{n}_{m,t+1}$ to $\{\{n_{z\mathcal{A}}(z, \mathcal{A}_{mt})\}_{z\in\mathcal{Z}}\}_{F_{\mathcal{A}}(\mathcal{A})>0}$, we can focus on finding the stationary counterpart of the latter. The evolution of the mass of consumer characteristics is

$$n_{z\mathcal{A}}(z, \mathcal{A}_{mt}) \equiv \begin{cases} \sum_{\mathcal{A}_{m,t-1}} n_{z\mathcal{A}}(z, \mathcal{A}_{m,t-1}) F_A(\mathcal{A}_{mt}) \left(1 - \sum_{j\in\mathcal{A}_{m,t}} s_{zj}(\mathcal{A}_{mt}, \mathcal{A}^{new}_{m,t-1})\right) & \text{if } \mathcal{A}_{mt} \neq \emptyset \\ n^{new}_{zm} & \text{if } \mathcal{A}_{mt} = \emptyset \end{cases} \tag{34}$$

This is derived by multiplying n_{mt} to both the numerator and the denominator of equation (9) and realizing the denominator becomes $n^{old}_{m,t+1}$. Note that (34) is necessary and sufficient for (3) and (10). In the stationary consumer equilibrium, the mass of consumer characteristics $\boldsymbol{n}^*_{z\mathcal{A}}$ is the fixed point of (34) given $\boldsymbol{x}^*$ and $F^*_{\mathcal{A}}$.

In the following, we provide a subroutine to solve for the steady-state consumer equilibrium. Superscript t denotes the number of iterations. The pseudocode is described in $\boldsymbol{n}_{z\mathcal{A}}$, but conversion between $\boldsymbol{n}_{z\mathcal{A}}$ and the pair $(f_{z\mathcal{A}}, \boldsymbol{n})$ is straightforward. See Appendix D (Algorithm 1).

6.2 Supply-side Equilibrium

Because there is a unique mapping between $(\boldsymbol{p}, \boldsymbol{q})$ and $(\boldsymbol{p}, \boldsymbol{d})$, we compute the equilibrium in $(\boldsymbol{p}, \boldsymbol{q})$ space. A profile $(\boldsymbol{p}, \boldsymbol{q})$ is a steady-state supply-side equilibrium if and only if (i) $q^{new}_{jm} = d^{new}_{jm}$ and $q^{old}_{jm} = \tilde{q}_{jm} + d^{old}_{jm}$, (ii) $(\boldsymbol{p}, \boldsymbol{q})$ is a steady-state consumer equilibrium, (iii) each old title satisfies (30) given (p_j, ξ_j), and (iv) each new title satisfies (31) and (32) given s_j.

The fill-the-gap property implies that condition (iii) holds if and only if, for each (j, m),

$$(q^{new}_{jm}, q^{old}_{jm}) = \begin{cases} (0, 0) & \text{if } p_j < mc_{jm}, \\ \left(q^{new}_{jm},\, \max\{\tilde{q}_{jm},\, D^{old}_{jm}(p_j, \xi_j)\}\right) & \text{if } p_j > mc_{jm}. \end{cases} \tag{35}$$

If $p_j < mc_{jm}$, then $d^{new}_{jm} = d^{old}_{jm} = 0$, $S^{new}_{jm} = S^{old}_{jm} = 0$, $P^S_{jm} = 1$, $P^R_{jm} = 0$, so (32) holds automatically and that market contributes zero to (31). Based on this observation, condition

(iv) can be restricted to markets with a positive inventory.

Because q_{jm}^{old} is determined by q_{jm}^{new} through (35), we recover $\boldsymbol{q}^{old}$ from $(\boldsymbol{p}, \boldsymbol{q}^{new})$, and the objects to search over reduce to $(\boldsymbol{p}, \boldsymbol{q}^{new})$. There are $J(M+1)$ such variables, and condition (iv) imposes $J(M+1)$ conditions. We compute the supply equilibrium by evaluating condition (iv), with the delivery condition implemented as complementary slackness (Algorithm 2).

7 Estimation

7.1 Demand parameters

We consider the case with no heterogeneity across consumers within a market, that is, $\mathcal{Z}$ is a singleton. Consumers still differ across markets in demographic attributes. We observe the supply-side variables $\{\boldsymbol{p}, \boldsymbol{q}\}$, and shares $\boldsymbol{s}$. Parameters to be estimated are the inflow in each market n_m^{new} for all m and $(\alpha, \boldsymbol{\beta}, \boldsymbol{\gamma})$. Define

$$\delta_j = \boldsymbol{\beta}' x_j + \xi_j - \alpha p_j \tag{36}$$

to be the intercept of mean indirect utility, which is the same across all markets.

Because inflow matches with outflow in the steady state equilibrium, the inflow estimate $\hat{n}_m^{new}$ is given by

$$\hat{n}_m^{new} = \sum_j^J s_{jm}. \tag{37}$$

We estimate $\boldsymbol{\delta}$ and $\boldsymbol{\gamma}$ by minimizing the relative difference between the predicted shares $\hat{\boldsymbol{s}}$ and the actual shares $\boldsymbol{s}$, rather than the absolute difference. An absolute difference puts most weight on high-sales markets. The relative difference puts weight on low-sales markets and reduces the gap between predicted and actual stockout rates there. The minimum-distance estimator $(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\gamma}})$ satisfies

$$\arg\min_{\boldsymbol{\delta}, \boldsymbol{\gamma}} \sum_m \sum_j \left(\frac{s_{jm} - s_{jm}(n_m^{new}, \boldsymbol{\delta}, \boldsymbol{\gamma})}{s_{jm}} \right)^2, \tag{38}$$

where the predicted share $s_{jm}(n_m^{new}, \boldsymbol{\delta}, \boldsymbol{\gamma})$ is computed from Algorithm 1.[4] The coefficients α and $\boldsymbol{\beta}$ are then derived by regressing $\hat{\boldsymbol{\delta}}$ on price and exogenous title covariates, provided

[4]Estimation proceeds in two reweighted steps. In the first, we fit (38) with equal weight on each market–title cell. In the second, we fit the same criterion with each cell weighted by $1/|s_{jm} - s_{jm}(n_m^{new}, \boldsymbol{\delta}, \boldsymbol{\gamma})|$, evaluated at the step-one parameters. The reported $(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\gamma}})$ are the step-two estimates.

a valid instrument for price exists.

7.2 Supply parameters

From demand estimation we have $\hat{\xi}_j$ and $\hat{\sigma}_\xi$. For a given standard deviation of the signal shock, σ_s, we invert the signal from the price and delivery first-order conditions, with markets restricted to $d_{jm}^{new} > 0$. The interior delivery FOC is affine in mc_{jm}, so $mc_{jm} = p_j P_{jm}^S/(1 - \delta P_{jm}^R)$. Substitute this marginal cost into (31). The signal $\hat{s}_j(\sigma_s)$ solves

$$\sum_{m:\, d_{jm}^{new}>0} \left(\mathbb{E}\big[S_{jm}^{new} + p_j \tfrac{\partial S_{jm}^{new}}{\partial p_j} + \delta S_{jm}^{old} + \delta p_j \tfrac{\partial S_{jm}^{old}}{\partial p_j} \;\big|\; s_j \big] - \tfrac{\delta p_j P_{jm}^S}{1-\delta P_{jm}^R} \mathbb{E}\big[\tfrac{\partial d_{jm}^{old,*}}{\partial p_j} \;\big|\; s_j \big] \right) = 0. \tag{39}$$

With the estimated signal $\hat{s}_j(\sigma_s)$, we subtract $\hat{\xi}_j$ to derive the signal shock $\hat{\nu}_j(\sigma_s) = \hat{s}_j(\sigma_s) - \hat{\xi}_j$ and recompute σ_s using $\hat{\nu}_j$. The estimated σ_s solves the fixed point $g(\sigma_s) = \sigma_s^2$, where

$$g(\sigma_s) = \frac{1}{J} \sum_j \left(\hat{\nu}_j(\sigma_s) - \frac{1}{J} \sum_j \hat{\nu}_j(\sigma_s) \right)^2. \tag{40}$$

At $\hat{\sigma}_s$, $\hat{s}_j = \hat{s}_j(\hat{\sigma}_s)$ and $\widehat{mc}_{jm}$ is the marginal cost implied from the delivery condition above.

7.3 Data

We obtain data on brick-and-mortar bookstores from Tohan Corporation, one of Japan's two largest book wholesalers, covering 2015 to 2017. Each observation is at the bookstore-month-title level and records deliveries, returns, orders, and sales. At the title level, we observe the publisher, author, list price, Nippon Decimal Classification, and publication date, and, where available, the page count from Amazon. We drop titles published before the sample begins, filter problematic records, impute missing title-store-month flows, and adjust first-month sales after publication using the actual release date scraped from Amazon. The detailed rules are described in Appendix B.

Our sample includes the top-selling books in social science (NDC2 = 30). The category is selected by sales volume and by concentration among leading titles. Specifically, ranked by first-year national sales, the five hundred leaders sell about 1.5 million copies and account for 80 percent of the category's 1.9 million copies nationwide. Because sales are concentrated among the leaders, we restrict attention to those titles. For a given period, defined as a calendar half-year, we select the one hundred titles published then with the largest national deliveries over the first three months after publication. Figure 1a justifies this six-month

window. Most sales occur within the first year, which covers both the new period and the old period, so we split that year into two periods of equal length.

A market is a municipality-period in social science. Our sample is restricted to Hokkaido and to the first half of 2015, so the selected titles are those published in that half-year. In each market, the one hundred titles selected for that period compete with later-period versions of themselves. Under the stationarity and large-market assumptions, the distributions of new and old titles are the same across periods. We therefore use the later-period versions of the new titles to approximate the old titles in the period. Pairing each selected title with its own later-period version has the advantage of giving the old and new titles the same characteristics x_j, as the model assumes. Sales, deliveries, and returns of a new title are the sums over the first six months after publication. Those of the same title in the later period are the sums of the same variables over months seven through twelve. Both are counted from the publication date, so a title released late in the period still contributes six months of sales.

We construct returns and inventory at the market-title level. In the model, the publisher has no incentive to return a new title early, because an unsold copy carries over and can still be sold in the later period, and returning it early makes the publisher incur an extra delivery cost when the title later turns out to be popular.[5] To align returns with this implication, we treat any overlap between new-title returns and old-title deliveries in the same market as a within-market transfer rather than a real return. Because a copy returned at one store and an old copy delivered at another are the same book moving between stores, we subtract the overlap from both the new-title return and the old-title delivery. We then move any remaining new-title returns to the end of the later period, so the new-title return rate is zero by construction. Inventory is then constructed after these delivery and return adjustments. A new title's inventory equals its delivery. An old title's inventory equals the new title's delivery minus its sales, plus the old title's delivery after this subtraction.

Table 1 summarizes sales, delivery, inventory, return, the return rate, the stockout rate, and list price at the market–title level for new and old titles separately. There are 42 markets, each with 100 new and 100 old titles. New-title returns are zero by construction, and the old-title return rate averages 78.9 percent. The stockout rate is 6.6 percent for new titles and 1.7 percent for old titles. Sales fall from 2.2 books for a new title to 0.4 books for the same title when old, while inventory remains large relative to sales in both periods.

[5]We do not consider a capacity constraint. In the data, bookstores also return copies at the end of the new period, which can be due to limited shelf space.

Table 1: Descriptive statistics for titles published in the first half of 2015

		New			Old		
Variable	N	Mean	Median	Sd	Mean	Median	Sd
Sales (books)	4,200	2.2	0.0	8.0	0.4	0.0	1.8
Delivery (books)	4,200	4.5	1.0	11.6	0.3	0.0	1.8
Inventory (books)	4,200	4.5	1.0	11.6	2.6	1.0	5.5
Return (books)	4,200	0.0	0.0	0.0	1.5	0.0	3.5
Return rate (%)		0.0	0.0	0.0	78.9	100.0	40.8
Stockout (%)		6.6	0.0	24.9	1.7	0.0	12.9
List price (yen)	4,200	1,078	1,000	346	1,078	1,000	346

The return rate and the stockout rate are computed only on cells with positive inventory (2,410 new and 2,298 old). The return rate is 100 when a cell's return is positive and 0 otherwise.

8 Result

8.1 Demand estimation

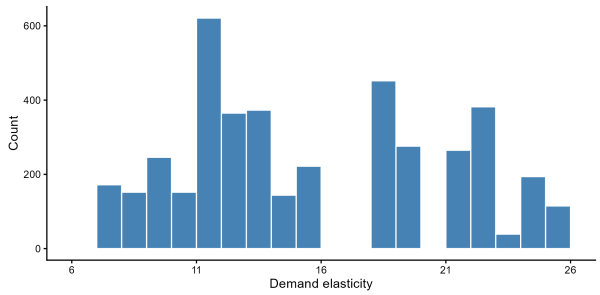
The covariates in x_j are a constant intercept, third-tier NDC indicators, an indicator of whether a page count is observed, and page count when it is observed. The covariates in $\boldsymbol{w}_m$ are selected in two LASSO steps. The first LASSO selects the daytime population ratio. It regresses the log share on municipality demographics after removing book and prefecture fixed effects. Demand is then estimated with only this ratio as demographics, and the model generates a predicted stockout for each market, title, and new or old copy. The second LASSO selects the service-sector employment share, the junior-high share, and the high-school share by regressing the difference between this predicted stockout and the observed stockout on inventory, new-title status, and municipality demographics, after removing book fixed effects. The specification uses the daytime population ratio from the first LASSO and these three shares from the second.

We estimate this regression for social science (NDC2 = 30). Price is instrumented with a Hausman shifter: the average price of the same publisher's titles in other categories during the same half-year. The instrument captures publisher-level cost common across categories, while the restriction to other categories makes the instrument less likely to comove with the focal title's unobserved quality. Table 2 reports this regression, with the constant intercept and the third-tier NDC indicators omitted. The estimated price coefficient is negative, as downward-sloping demand implies, and it is significant at the 10 percent level.[6]

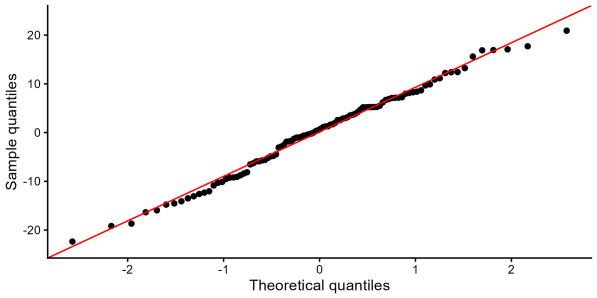
[6]The reported standard error treats the recovered intercept of mean indirect utility, $\boldsymbol{\delta}$, as data, so it does not account for estimation error in $\hat{\boldsymbol{\delta}}$.

Table 2: Regression of the mean indirect utility intercept

	δ
Price	-15.224+
	(7.957)
Number of Pages (If Observed)	0.014
	(0.019)
Num.Obs.	100
R2	0.500
R2 Adj.	0.468
AIC	740.2
BIC	758.4
RMSE	9.13



(a) Distribution of estimated demand elasticities



(b) Normal QQ plot of recovered title quality

We check the scale of the price coefficient through own-price elasticities at the market–title level. Elasticity is undefined when inventory is zero. Figure 2a shows a point mass at zero, which happens when the market stocks out, and the remaining elasticities typically lie between 3 and 26, with a mean of 16. The residual from the regression is unobserved title quality ξ, common across all markets. The model places a normal prior on ξ and a normal signal around it, so the shape of this residual matters for that assumption. Figure 2b plots its quantiles against normal quantiles. The points lie close to the normal line, so the distribution of ξ is approximately normal.

Table 3 reports the moment fit at the market–title level. Each observation is a market–title cell, as in Table 1. New-title returns are zero by construction, because any remaining new-title return is moved to the end of the later period, so the table omits the new-title return and return rate. The fit on mean sales is close. The average sales for a title in a municipality is about two books during the six-month period. The stockout rate is the share of cells with positive inventory that sell out. The main discrepancy is stockouts. Predicted new-title stockouts average 15.9 percent, against 6.6 percent in the data, and predicted old-title

Table 3: Predicted and actual summary moments at the market–title level

			Estimated		Actual	
Group	Moment	Unit	Mean	Sd	Mean	Sd
New	Sales	Books	2.2	8.5	2.2	8.0
	Stockout Rate	Percent	15.9	36.6	6.6	24.9
Old	Sales	Books	0.5	3.2	0.4	1.8
	Return	Books	2.1	4.2	1.5	3.5
	Return Rate	Percent	93.3	25.1	78.9	40.8
	Stockout Rate	Percent	6.7	25.1	1.7	12.9

stockouts average 6.7 percent, against 1.7 percent. One reason is that demand is identified mainly from markets without a stockout, where demand is not truncated by inventory; in markets with a stockout, observed sales are truncated above by inventory and understate true demand. Those non-stockout markets are likely to be high-inventory, high-sales markets with a larger local demand shock, and if that shock is absorbed into the unobserved title quality that is common across all markets, applying it to thinner markets overstates stockout risk.

8.2 Supply estimation

We estimate the supply side with a title-specific signal-noise standard deviation,

$$\sigma_{s,j} = \exp\left(\phi_0 + \phi_1 \mathbf{1}\{\text{new author}_j\}\right).$$

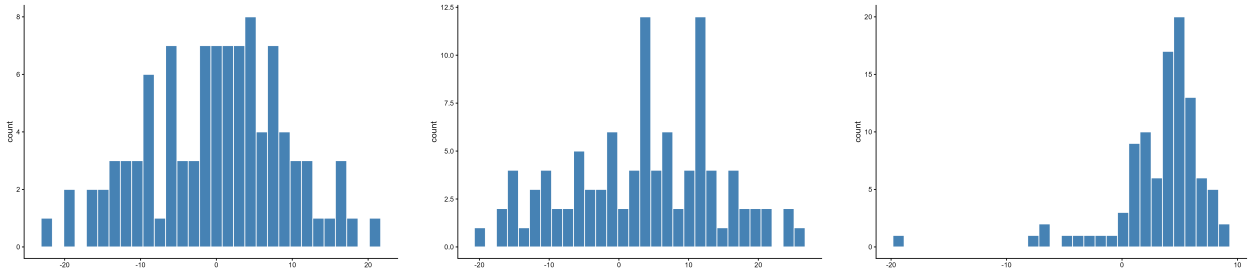
The indicator equals one when the title is the author's first title in the category. The estimated intercept is $\hat{\phi}_0 = 1.200$ and the estimated new-author coefficient is $\hat{\phi}_1 = 0.179$. The implied signal-noise standard deviation is 3.32 for an established author and 3.97 for a new author. A title by a new author is therefore estimated to have more uncertain demand.

Table 4 reports the descriptive statistics for the recovered $\hat{\xi}$, $\hat{s}$, and $\hat{\nu}$ over the one hundred leading titles in the social science category. The unobserved quality $\hat{\xi}$ has mean zero by construction and standard deviation 9.18. The signal shock $\hat{\nu} = \hat{s} - \hat{\xi}$ is less dispersed than $\hat{\xi}$, with standard deviation 3.89. Figure 3 plots the three distributions. The distribution of $\hat{\nu}$ is concentrated on positive values, with a mode near 4.

Given $\hat{s}_j$, the marginal cost is recovered from the interior delivery condition $\widehat{mc}_{jm} = p_j P^S_{jm}/(1 - \delta P^R_{jm})$ only for markets with positive inventory. This expression also highlights an advantage of recovering marginal cost from the delivery condition rather than from the

Table 4: Descriptive statistics for recovered supply-side signals

Series	N	Min	Max	Mean	Median	Sd
$\hat{\xi}$	100.00	-22.36	20.91	-0.00	0.63	9.18
$\hat{s}$	100.00	-19.94	25.92	3.39	3.84	10.63
$\hat{\nu}$	100.00	-19.35	8.85	3.39	4.16	3.89



(a) $\hat{\xi}$ (b) $\hat{s}$ (c) $\hat{\nu}$

Figure 3: Distributions of $\hat{\xi}$, $\hat{s}$, and $\hat{\nu}$ (social science).

Table 5: Descriptive statistics for the implied markup $1 - \widehat{mc}/p$

N	Min	Max	Mean	Median	Sd
2410.00	0.01	1.00	0.27	0.20	0.23

price first-order condition. In the latter approach, marginal cost can turn negative when demand is locally inelastic, which is difficult to interpret economically. In our model, because the seller makes the delivery and price decisions under demand uncertainty, the recovered marginal cost is always positive and proportional to the stockout probability P^S_{jm}. For zero-inventory markets, all we know is the inequality $p_j < mc_{jm}$, not by how much mc_{jm} exceeds p_j. We summarize the implied markup, $1 - \widehat{mc}_{jm}/p_j$, in Table 5 and Figure 4 over the 2,410 title–market observations with positive inventory. The mean markup is 0.27 and the median is 0.20. The distribution has a mode near 0.2, a right tail, and a separate mass of observations near 1.

8.3 Constrained estimation

We compare our baseline marginal cost estimate from the previous section with the classical case, where the seller meets demand without stockout and prices after observing the demand shock.

To isolate the channels that drive the difference, we construct four endpoints in the supply estimation based on whether price and inventory are committed before or after the

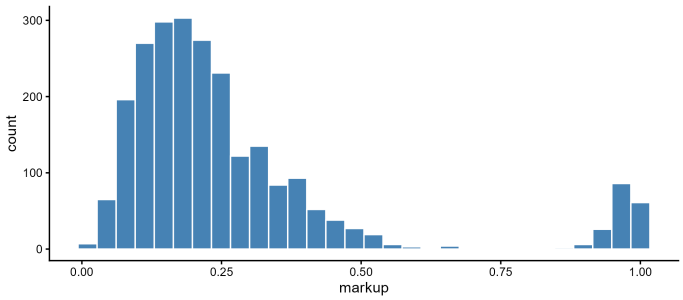


Figure 4: Distribution of the implied markup $1 - \widehat{mc}/p$ (social science).

demand shock. Endpoint A is the classical case, where both price and inventory are set after the shock is observed, while Endpoint D is our baseline, where both are committed before the shock. In between, endpoint B keeps A's timing but evaluates demand at the posterior mean instead of at the realized shock, and Endpoint C sets price before the shock but still lets inventory adjust to the realization. The three gaps are $\Delta_{AB,j} = (mc_j^A - mc_j^B)/p_j^o$, $\Delta_{BC,j} = (mc_j^B - mc_j^C)/p_j^o$, and $\Delta_{CD,j} = (mc_j^C - mc_j^D)/p_j^o$. The first isolates the effect of the demand surprise, the second isolates the effect of setting price after observing the demand realization versus before the shock, and the third isolates the effect of committing inventory before the shock.

For each title j, we hold fixed across endpoints the observed price p_j^o, the estimated shock $\widehat{\xi}_j$, the posterior mean $\mu_{\xi_j|s_j}$, the observed deliveries $\boldsymbol{d}_j^o$, the rivals' policies, and the demand-side parameters. Let $D_{jm}(\xi)$ be the discounted unrestricted sales at the observed price and $\omega_{jm}(\xi)$ its derivative with respect to price, where $D_{jm}(\xi) = D_{jm}^{new}(\xi) + \delta D_{jm}^{old}(\xi)$ with D_{jm}^{new} and D_{jm}^{old} the new and old title's no-stockout sales. Let the corresponding aggregates over markets with $d_{jm}^o > 0$ be $D_j(\xi) = \sum_m D_{jm}(\xi)$ and $\Omega_j(\xi) = \sum_m \omega_{jm}(\xi)$, the aggregate price slope. The magnitude of the aggregate price slope is large when many consumers are near the margin of purchasing. It is small when few consumers purchase, and small when almost all consumers purchase. For A, the seller's interior price FOC at the observed price gives the weighted marginal cost identified at the title level,

$$mc_j^w(\xi) \equiv \frac{\sum_m \omega_{jm}(\xi)\, mc_{jm}}{\Omega_j(\xi)} = p_j^o + \frac{D_j(\xi)}{\Omega_j(\xi)}.$$

Table 6 shows the four endpoints and their implied weighted costs. The marginal cost estimates under A, B, and C are derived from the price FOC, so we only identify a title-level weighted cost. D commits delivery before the shock and uses the baseline costs $\widehat{mc}_{jm}^{\text{base}}$, aggregated to the title level using the same posterior expected derivative weights as in C so that mc_j^C and mc_j^D are comparable.

Table 6: Four endpoints.

Endpoint	Price	Inventory	Implied Weighted Cost
A	After	After	$mc_j^A = p_j^o + \frac{D_j(\widehat{\xi}_j)}{\Omega_j(\widehat{\xi}_j)}$
B	After	After	$mc_j^B = p_j^o + \frac{D_j(\mu_{\xi_j \mid s_j})}{\Omega_j(\mu_{\xi_j \mid s_j})}$
C	Before	After	$mc_j^C = p_j^o + \frac{\mathbb{E}_{\xi_j}[D_j]}{\mathbb{E}_{\xi_j}[\Omega_j]}$
D	Before	Before	$mc_j^D = \frac{\sum_m \mathbb{E}_{\xi_j}[\omega_{jm}]\, \widehat{mc}_{jm}^{\text{base}}}{\mathbb{E}_{\xi_j}[\Omega_j]}$

Table 7: Decomposition of marginal cost estimates.

Variable	Mean	Max	Min	Median	Sd
mc^A/p	0.566	0.763	-0.082	0.589	0.168
mc^B/p	0.582	0.759	0.134	0.595	0.143
mc^C/p	0.330	0.663	0.120	0.324	0.137
E_C/p	0.528	0.753	0.241	0.521	0.110
Cov_C/p	-0.198	0.058	-0.433	-0.219	0.090
mc^D/p	0.319	0.632	0.002	0.334	0.173
$(mc^A - mc^B)/p$	-0.016	0.091	-0.644	0.000	0.085
$(mc^B - mc^C)/p$	0.252	0.429	0.012	0.284	0.103
$(mc^C - mc^D)/p$	0.011	0.637	-0.158	0.006	0.089

Table 7 reports the endpoints and gaps for 97 of 100 titles, removing those outside ± 3 standard deviations. $\Delta_{AB,j}$ has mean -0.016 and median 0.000. The average is small because the demand surprise changes sign across titles and the title-level gaps offset. The average is small because the demand surprise changes sign across titles and the title-level gaps offset. The mean difference between the realized shock and the posterior mean is 1.21, while the standard deviation of realized quality is 8.11. At high ξ_j regions, the price derivative ω_{jm} is virtually zero because in the consumer equilibrium inflow equals outflow and almost all consumers purchase the title, so the price elasticity of demand falls. The weighted cost is $mc_j^w = p_j^o(1 - 1/\text{elasticity})$, so a lower elasticity implies a lower weighted cost. A positive surprise therefore lowers the weighted cost only when the elasticity at $\widehat{\xi}_j$ is lower than the elasticity at $\mu_{\xi_j|s_j}$. Among the 97 kept titles, 78 (80.4%) have this sign pattern, so a positive surprise comes with $\Delta_{AB,j} < 0$ and a negative surprise comes with $\Delta_{AB,j} > 0$. For the other 19 titles the elasticity is higher at the larger of the two shocks. For 14 of them the surprise and $\Delta_{AB,j}$ are both positive, and for 5 both are negative.

$\Delta_{BC,j}$ is the dominant component (mean 0.252, median 0.284) in the change from endpoint A to D. We can further decompose C into the expectation of the estimated marginal

Table 8: $|\Omega_j|$ and the price elasticity at each percentile of ξ, averaged across 100 titles. $|\Omega_j|$ is in books per 1,000 yen. Nodes of ξ start at the first point where the slope is not zero. The slope is zero before that point because sales are zero.

Low-ξ Percentile	$\lvert\Omega_j\rvert$	Mean Elasticity	Median Elasticity	N
10%	0.055	2.950	2.636	100
20%	0.318	2.637	2.457	100
30%	2.10	2.669	2.528	100
40%	17	2.662	2.450	100
50%	110	2.657	2.467	100
60%	360	2.604	2.428	100
70%	1,223	2.435	2.245	100
80%	3,853	1.932	1.698	100
90%	6,873	1.256	1.091	100

cost similar to A and a covariance term,

$$mc_j^C = \underbrace{E_{\xi_j|s_j}\left[mc_j^A(\xi)\right]}_{E_C} + \underbrace{\frac{\mathrm{Cov}_{\xi_j|s_j}\left(\Omega_j(\xi),\, mc_j^A(\xi)\right)}{E_{\xi_j|s_j}[\Omega_j(\xi)]}}_{\mathrm{Cov}_C}, \tag{41}$$

and the decomposition shows that the BC gap is mostly due to the covariance term. Across the 97 kept titles, E_C/p_j^o has mean 0.528 and Cov_C/p_j^o has mean -0.198, so the absolute value of the latter divided by the mean of $\Delta_{BC,j}$ is 78.6%. This covariance term is the extra weight on states with a steeper aggregate price slope when the price condition is averaged before the shock is realized. Those states are less elastic and have a lower implied cost, so the term lowers the cost at C. The term is negative because the expected aggregate price slope in the denominator is negative, while the slope and the implied cost move together. At a low shock, sales are small and the aggregate price slope is close to zero. Demand is elastic, so the markup is low and $mc_j^A(\xi)$ is close to the price. At a higher shock, the aggregate price slope is steeper and demand is less elastic, so the markup is higher and $mc_j^A(\xi)$ is lower. A less negative slope therefore goes with a higher implied cost, and $\mathrm{Cov}(\Omega_j, mc_j^A)$ is positive. Table 8 documents this comparison between a low shock and a higher shock. It reports $|\Omega_j|$ and the price elasticity at each percentile of ξ, averaged across 100 titles.

$\Delta_{CD,j}$ is small on average (mean 0.011, median 0.006) but dispersed (range -0.158 to 0.637). Because the dispersion is large, it is worthwhile to check how the difference varies across titles. Figure 5 plots mc_j^C/p_j^o against mc_j^D/p_j^o, colored by the observed price. Both mc_j^C/p_j^o and mc_j^D/p_j^o are low when the observed price is low. mc_j^C/p_j^o increases with price

Table 9: Averages by price group.

Price Group	Price	Subjective Probability	Sales	Inventory	N
Low ($\leq$ 33rd percentile)	721	0.137	18.9	28.2	33
Middle (33rd to 67th)	1,088	0.346	11.0	19.7	37
High ($>$ 67th percentile)	1,530	0.472	7.4	15.7	27

because the elasticity approximately scales with p_j.[7] A low-price title is therefore less elastic and has a lower mc_j^C/p_j^o. mc_j^D/p_j^o is roughly equal to the subjective stockout probability, because the replenishment probability is close to zero. Low-price titles tend to have a lower subjective stockout probability, with Pearson correlation 0.86 between the subjective stockout probability and price. Table 9 shows that this is mainly because the publisher stocks more of the low-price titles. For each title, sales and inventory are weighted across markets by the same weights as the subjective stockout probability, and the table reports the average of those weighted values across titles. Average inventory is 28.2 books in the low-price group and 15.7 books in the high-price group, while average sales are 18.9 and 7.4 books.[8] Another observation is that mc_j^C/p_j^o tends to lie above mc_j^D/p_j^o in the low-cost region, while mc_j^D/p_j^o lies at or above mc_j^C/p_j^o in the high-cost region. This is because the ratio under D is more dispersed, with standard deviation 0.173 against 0.137.

[7]Holding the continuation value and the distribution of consumers fixed, the elasticity is close to $\alpha p_j(1 - s_{zj})$. This is because one title has a limited effect on the inclusive value H. With many titles, s_{zj} is close to zero, so the elasticity is close to αp_j. In the full equilibrium, the continuation value and the distribution of consumers also move with p_j, but the elasticity still increases with price.

[8]These averages are larger than the roughly two books per title in a municipality reported in the demand estimates. There, each title-market is weighted equally. Here, markets are weighted by the price slope ω_{jm} used for the subjective stockout probability. That slope is larger in magnitude where sales are higher, so the weighted averages are larger.

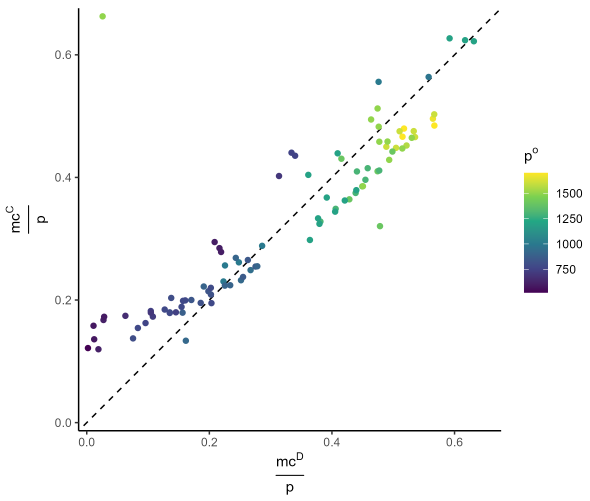


Figure 5: Marginal cost ratios at C and D.

A Stockout-based substitution

This appendix tests stockout-based substitution. Three regressions ask whether a stockout shifts demand to a substitute, either to a later month, to another title, or to another store.

A.1 Sample and variables

The regressions use a national sample of 40 municipalities drawn with probability proportional to population, covering 2015 to 2017. The sample is not restricted to social science or to Hokkaido. Every store-month-title in the 40 municipalities is retained, because the goal is to document how consumers substitute rather than to estimate the structural model. The unit of observation is a store-month-title. For each cell we observe deliveries, sales, and returns. We construct beginning-of-month inventory as the previous month's ending inventory plus this month's delivery. Ending inventory is constructed from the transaction records as described in Appendix B.

The stockout indicator equals one when a title stocks out in a store in a given month and zero otherwise. Each regression controls for the title's own beginning-of-month inventory. It also includes fixed effects, which control for time-invariant differences across titles, stores, calendar months, and months since publication (post_pub). We refer to these fixed effects as the baseline. Standard errors are clustered by title and by store. The dependent variable is sales of a title in a store in a month. The regressors of interest are stockout counts.

A.2 Intertemporal substitution

The first regression asks whether a stockout shifts demand to the following month. The key regressor is a stockout indicator for the same title and store in the previous month. If consumers who could not buy when the title was unavailable return later, sales this month should be higher after a stockout last month, holding current inventory fixed. Table 10 shows that a title that stocked out last month sells 0.066 more copies this month, so demand spills across periods.

A.3 Intertitle substitution

The second regression asks whether a stockout of one title shifts demand to other titles. The key regressors are the number of other titles in the same store that stock out this month, split into titles in the same category (NDC2) and titles in other categories, each scaled per thousand titles so the coefficients are readable. If consumers substitute across titles when one is unavailable, sales of a title should be higher when more other titles stock out, especially

Table 10: Intertemporal substitution

	(1)
Inventory	0.172***
	(0.021)
Inventory last month	-0.005
	(0.015)
Stockout last month	0.066**
	(0.021)
Number of stores	-0.006*
	(0.003)
Num.Obs.	104830169
R2	0.400
R2 Adj.	0.399
R2 Within	0.327
R2 Within Adj.	0.327
FE: book_id	X
FE: store_id	X
FE: year_month	X
FE: post_pub	X

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

within the same category. Columns (1) and (2) use the baseline fixed effects. Table 11 shows that a stockout of another title raises sales, more within the same category (0.071) than across categories (0.036) per thousand titles.

A larger within-category coefficient could be due to demand shocks that are more correlated within a category instead of stockout-based substitution. For example, a widely reported event in a subject, the start of a school term, can sharply raise demand for the whole category in a given month. A positive shock shared by titles in the same category stocks out more of those titles and raises sales of the titles that remain in stock. To control for this category-month level demand shock, columns (3) and (4) replace the baseline year-month fixed effects with NDC2-by-year-month fixed effects. The coefficient on stockouts of other titles in the same store falls from 0.038 in column (1) to 0.035 in column (3) and remains significant, so intertitle substitution within the store is only slightly lower. In column (4), the within-category coefficient rises from 0.071 to 0.129 and the across-category coefficient falls from 0.036 to 0.026, so the within-category concentration strengthens once the category-month demand shock is controlled for.

Table 11: Intertitle substitution

	(1)	(2)	(3)	(4)
Stockouts of other titles (per 1,000)	0.038***		0.035***	
	(0.009)		(0.009)	
Inventory of other titles	-0.000		-0.001	
	(0.000)		(0.000)	
Inventory	0.217***	0.217***	0.217***	0.217***
	(0.020)	(0.020)	(0.020)	(0.020)
Stockouts in the same category (per 1,000)		0.071**		0.129***
		(0.021)		(0.022)
Stockouts in other categories (per 1,000)		0.036***		0.026**
		(0.010)		(0.010)
Inventory of the same category		-0.000		0.000
		(0.001)		(0.001)
Inventory of other categories		-0.000		-0.001
		(0.000)		(0.000)
Num.Obs.	98166609	98166609	98166609	98166609
R2	0.453	0.453	0.453	0.453
R2 Adj.	0.452	0.452	0.452	0.452
R2 Within	0.393	0.393	0.393	0.393
R2 Within Adj.	0.393	0.393	0.393	0.393
FE: book_id	X	X	X	X
FE: store_id	X	X	X	X
FE: post_pub	X	X	X	X
FE: ndc_2_valueˆyear_month			X	X
FE: year_month	X	X		

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

A.4 Interstore substitution

The third regression asks whether sales in one store are affected by a stockout in other stores, either in the same municipality or in a neighboring municipality. If consumers substitute across stores when a title is unavailable locally, sales in a store should be higher when more other stores stock out. Columns (1) and (2) use the baseline fixed effects. Column (1) shows that a stockout in another store in the same municipality is associated with 0.152 higher sales, consistent with interstore substitution. Column (2) adds stockouts in neighboring municipalities, and that coefficient is 0.125.

These positive coefficients can arise from a title-month demand shock that is common across bookstores in the municipality, for example when news or the launch of a related product sharply raises demand for that title. Because the shock raises demand at every

store, more stores stock out and sales rise at the stores that remain in stock. To control for this title-month demand shock, columns (3) and (4) replace the baseline book and year-month fixed effects with title-month fixed effects. The within-municipality coefficient falls to -0.016 and is no longer significant, and the neighbor-municipality coefficient falls to 0.026 and is insignificant. The higher sales when other stores in the same municipality or a neighboring municipality stock out are mostly due to the title-month demand shock.

Table 12: Interstore substitution

	(1)	(2)	(3)	(4)
Stockouts in other stores	0.152***	0.128***	-0.016	-0.033+
	(0.018)	(0.021)	(0.013)	(0.017)
Number of stores	-0.031***	-0.030***	-0.011**	-0.014*
	(0.005)	(0.006)	(0.004)	(0.006)
Inventory	0.211***	0.192***	0.200***	0.183***
	(0.022)	(0.018)	(0.022)	(0.019)
Inventory of other stores	0.008***	0.006**	-0.000	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)
Stockouts in neighbor mun.		0.125***		0.026
		(0.019)		(0.019)
Stores in neighbor mun.		-0.009**		0.001
		(0.003)		(0.002)
Inventory in neighbor mun.		0.006***		0.002
		(0.001)		(0.001)
Num.Obs.	74858140	51680752	74858140	51680752
R2	0.461	0.473	0.538	0.559
R2 Adj.	0.459	0.472	0.516	0.533
R2 Within	0.398	0.401	0.370	0.362
R2 Within Adj.	0.398	0.401	0.370	0.362
FE: store_id	X	X	X	X
FE: post_pub	X	X	X	X
FE: book_idˆyear_month			X	X
FE: year_month	X	X		
FE: book_id	X	X		

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B Data Construction

B.1 Filtering

We exclude four sets of records. We drop titles published on or before the start of the sample, so that each retained title has a complete history of sales, returns, and deliveries after publication. We drop one title with multiple prices, because a uniform retail price is a defining feature of resale price maintenance and a title sold at more than one price violates the arrangement the model assumes. We drop bookstore-months with no sales, which we attribute to inactive stores. Finally, we drop bookstore-title pairs with negative sales or deliveries, which we attribute to recording error.

B.2 Constructing inventory data for brick-and-mortar bookstores

We construct inventory from the transaction records. End-of-month inventory equals cumulative deliveries minus cumulative returns and cumulative sales, each summed through the same month. Beginning-of-month inventory equals the previous month's ending inventory plus this month's delivery, and in the first month a store carries the title it equals that month's delivery. Inventory can be negative for several reasons. Chain stores may transfer books across outlets, so sales are recorded without a matching delivery. Some bookstores transact with multiple wholesalers, so deliveries from other wholesalers are missing from our data even though sales are observed. For books published before 2015, initial inventories may be unobserved because our dataset begins in 2015. If inventory is negative in any month, we add the shortfall to deliveries in the first observed month of the title-store pair.

B.3 Imputing zero sales for brick-and-mortar bookstores

One 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, return, or order 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. For each title-store pair that records at least one such flow within the first six months after publication, we impute zero delivery, zero sales, zero return, and zero order for every missing month in which the bookstore is open. Title-store pairs with no observed flow within the first six months after publication are dropped.

B.4 Adjusting first-month sales for brick-and-mortar bookstores

The POS system records the publication date as the first day of the publication month, so if a title is released later in the month, the month's observed sales cover only part of the month. When the Amazon release date falls in that same month, we scale the month's sales by $31/d$, where d is the number of days from the release date to the end of the month, and add the resulting increment to deliveries.

C Notes

C.1 Update distribution of availability set

Define the l-th stockout book by $j(l)$ and the consumer experiencing the l-th stockout by $i(l) \equiv \inf\{i : q_{j(l)mt}(i) = 0\}$. We can derive $j(l)$ recursively as follows:

$$j(l) = \arg\min_{j \in \mathcal{A}_{mt}(i(l-1))} \left(\frac{q_{jmt}(i(l-1))}{\sum_z \sum_{\mathcal{A}_{m,t-1}} f_{z\mathcal{A}}(z, \mathcal{A}_{m,t-1}) s_{zj}(\mathcal{A}_{mt}(i(l-1)), \mathcal{A}^{new}_{m,t-1})}, n_{mt} - i(l-1) \right), \tag{42}$$

where $j(l) = \emptyset$ if $n_{mt} - i(l-1)$ achieves the minimum. Note that $j(l)$ is defined based on the inventory and availability set of $i(l-1)$. Starting with $i(0) = 0$, $j(l)$ can be iteratively determined until all consumers have made their decisions.

Expressing (8) explicitly, $F_{\mathcal{A}}(\mathcal{A}_{mt})$ can be rewritten as

$$F_{\mathcal{A}}(\mathcal{A}_{mt}) = \begin{cases} \frac{i(l+1)-i(l)}{n_{mt}} & \text{, if } \mathcal{A}_{mt} = \mathcal{A}_{mt}(i(l)) \text{ for } l = 0, \dots l^*_{mt} \\ 0 & \text{, otherwise} \end{cases} \tag{43}$$

C.2 Consumer Equilibrium with no stockout

Consider the case where each title's inventory exceeds the total consumer mass. In this scenario, the only possible steady-state consumer equilibrium is one without stockouts during the period. The distribution of availability sets, $F_{\mathcal{A}}$, degenerates to a singleton, the initial availability set at the beginning of the period, denoted by $\mathcal{A}^*$. The cutoff value $\boldsymbol{x}^*$ can be calculated from (20).

One of the conditions in the definition of a steady-state consumer equilibrium is that the distribution of consumer segments $\boldsymbol{n}^*$ does not change over time. In the steady state, the mass of new consumers must equal the inflow, while the mass of old consumers evolves according to (3). The stationarity of the mass of old consumers implies that the outflow of segment-z consumers equals the inflow n^{new*}_{zm}. The outflow of segment-z consumers comprises

two components: exits from new consumers and exits from old consumers. The mass of exit from the segment-z new consumers is

$$n_{zm}^{new*} \sum_{j \in \mathcal{A}^*} s_{zj}(\mathcal{A}^*, \emptyset).$$

Therefore, for the consumer distribution to remain stationary, the mass of segment-z old consumers exiting the market must satisfy

$$n_{zm}^{new*} \left(1 - \sum_{j \in \mathcal{A}^*} s_{zj}(\mathcal{A}^*, \emptyset)\right).$$

On the other hand, the mass of segment-z consumers exiting the market should equal

$$n_{zm}^{old*} \sum_{j \in \mathcal{A}^{*new}} s_{zj}(\mathcal{A}^*, \mathcal{A}^{*old}),$$

where $\mathcal{A}^{*new}$ is the set of new titles in $\mathcal{A}^*$ and $\mathcal{A}^{*old}$ is the set of old titles in $\mathcal{A}^*$. Equating these two expressions, we can solve for the mass of segment-z old consumers in the equilibrium as

$$n_{zm}^{old*} = \frac{n_{zm}^{new*} \left(1 - \sum_{j \in \mathcal{A}^*} s_{zj}(\mathcal{A}^*, \emptyset)\right)}{\sum_{j \in \mathcal{A}^{*new}} s_{zj}(\mathcal{A}^*, \mathcal{A}^{*old})}.$$

The marginal density of consumer characteristics is

$$f_{z\mathcal{A}}(z, \mathcal{A}) = \begin{cases} \frac{n_{zm}^{new*}}{\sum_{\tau \in \{new,old\}} \sum_{z \in \mathcal{Z}} n_{zm}^{\tau *}} & \text{, if } \mathcal{A} = \phi \\ \frac{n_{zm}^{old*}}{\sum_{\tau \in \{new,old\}} \sum_{z \in \mathcal{Z}} n_{zm}^{\tau *}} & \text{, if } \mathcal{A} = \mathcal{A}^* \end{cases}$$

Sales for title j are also pinned down by

$$\sum_z \left(n_{zm}^{old*} s_{zj}(\mathcal{A}^*, \mathcal{A}^{*old}) + n_{zm}^{new*} s_{zj}(\mathcal{A}^*, \emptyset)\right). \tag{44}$$

The derived $\boldsymbol{n}^*$ and $\boldsymbol{x}^*$ would be part of a consumer equilibrium if the inventory at the beginning of the period is larger than the sales in (44). That is, for all j,

$$q_j^* \geq \sum_z \left(\frac{n_{zm}^{new*} \left(1 - \sum_{j \in \mathcal{A}^*} s_{zj}(\mathcal{A}^*, \emptyset)\right)}{\sum_{j \in \mathcal{A}^{*new}} s_{zj}(\mathcal{A}^*, \mathcal{A}^{*old})} s_{zj}(\mathcal{A}^*, \mathcal{A}^{*old}) + n_{zm}^{new*} s_{zj}(\mathcal{A}^*, \emptyset) \right).$$

A consumer equilibrium with no stockout exists as long as the inventory $\boldsymbol{q}^*$ is sufficiently

large or the inflow of new consumers is sufficiently small.

D Pseudocode

Algorithm 1: Subroutine CONSUMER_EQUILIBRIUM

while *the distance between* $(\boldsymbol{x}^t, \boldsymbol{n}_{z\mathcal{A}}^t)$ *and* $(\boldsymbol{x}^{t-1}, \boldsymbol{n}_{z\mathcal{A}}^{t-1})$ *is at least* ε **do**

Input: $[\boldsymbol{x}^t, \boldsymbol{n}_{z\mathcal{A}}^t]$;

1. Compute the individual choice probabilities $\boldsymbol{s}^t(\mathcal{A}', \mathcal{A})$ from $\boldsymbol{x}^t$ using (25) ;
2. Compute the intra-period inventory path $\boldsymbol{q}^t(i)$ from $\boldsymbol{s}^t(\mathcal{A}', \mathcal{A})$ and $\boldsymbol{n}_{z\mathcal{A}}^t$ using (4) and (6) ;
3. Update the distribution of the availability set $F_{\mathcal{A}}^t$ from $\boldsymbol{q}^t(i)$ using (8) ;
4. Compute aggregate shares $\boldsymbol{s}^t$ from $\boldsymbol{x}^t$ and $\boldsymbol{n}_{z\mathcal{A}}^t$ using (5) and (7) ;
5. Update the continuation values $\boldsymbol{x}^{t+1}$ from $F_{\mathcal{A}}^t$ by solving the fixed point of the perceived continuation value in (13) ;
   (a) For a given $F_{\mathcal{A}}^t$, derive the distribution of $\overline{u}_z(\mathcal{A}^{new})$. The right-hand side of (20) defines a contraction mapping in x_z^*. Solve for $\boldsymbol{x}^{t+1}$ by iterating the mapping from an arbitrary initial value ;
6. Update the mass of consumer characteristics $\boldsymbol{n}_{z\mathcal{A}}^{t+1}$ from $\boldsymbol{n}_{z\mathcal{A}}^t$ given $\boldsymbol{s}^t(\mathcal{A}', \mathcal{A})$ and $F_{\mathcal{A}}$ using (34) ;

Output: $[\boldsymbol{x}^{t+1}, \boldsymbol{n}_{z\mathcal{A}}^{t+1}]$;

if *the while condition is not satisfied* **then**

return $[\boldsymbol{x}^*, \boldsymbol{n}_{z\mathcal{A}}^*]$, *the converged values* ;

Two caveats are in order. First, the initial value of $(\boldsymbol{x}^0, \boldsymbol{n}_{z\mathcal{A}}^0)$ is chosen to be the consumer equilibrium without stockout. Second, the distance between $\boldsymbol{n}_{z\mathcal{A}}^t$ and $\boldsymbol{n}_{z\mathcal{A}}^{t-1}$ is computed by first converting them into sparse matrices. In these matrices, each row corresponds to a previous availability set, each column corresponds to a consumer segment, and the entries represent the mass of consumers. The distance is then measured as the Euclidean distance between these two sparse matrices.

Algorithm 2: Subroutine SUPPLY_EQUILIBRIUM

Input: $[\delta,\ \{mc_{jm}\}_{j\in\mathcal{L},m\in\mathcal{M}},\ \mu_\xi,\ \sigma_\xi,\ \sigma_s,\ \{s_j\}_{j\in\mathcal{L}^{new}},\ \{\alpha_z, \boldsymbol{\beta}_z, n_{zm}^{new}\}_{z\in\mathcal{Z},m\in\mathcal{M}}]$;

For all $j \in \mathcal{L}^{new}$ and $m \in \mathcal{M}$, set $q_{jm}^{new} = 0$ if $p_j < mc_{jm}$;

Recover $\boldsymbol{q}^{old}(\boldsymbol{p}, \boldsymbol{q}^{new})$ by (35) ;

foreach $j \in \mathcal{L}^{new}$ **do**

  Compute quadrature nodes for different realizations of ξ_j using Algorithm 1 ;

  Derive the price FOC (31) by integrating over the posterior $F_{\xi_j|s_j}$ using Newton–Cotes weights ;

foreach $m \in \mathcal{M}$ *and* $j \in \mathcal{L}^{new}$ **do**

  Compute the delivery FOC (32) using a similar procedure to the price FOC ;

Minimize the residual of the price FOC and the delivery FOC over $(\boldsymbol{p}, \boldsymbol{q}^{new})$ until convergence ;

Output:

$[(\boldsymbol{p}^*, \boldsymbol{q}^*)]$;

References

Ackerberg, D., C. Lanier Benkard, S. Berry, and A. Pakes (2007): "Chapter 63 Econometric Tools for Analyzing Market Outcomes," in *Handbook of Econometrics*, ed. by J. J. Heckman and E. E. Leamer, Elsevier, vol. 6, 4171–4276.

Avery Dennison (2022): "Supply Chain Crisis Made Worse As 8% of Stock Ends Up as Waste," https://rfid.averydennison.com/en/home/news-insights/press-releases/supply-chain-crisis-made-worse-as-8-percent-of-stock-ends-up-as-waste.html.

Bresnahan, T. F. (1989): "Chapter 17 Empirical Studies of Industries with Market Power," in *Handbook of Industrial Organization*, Elsevier, vol. 2, 1011–1057.

Calcott, P. and M. Walls (2000): "Can Downstream Waste Disposal Policies Encourage Upstream "Design for Environment"?" *American Economic Review*, 90, 233–237.

Deneckere, R., H. P. Marvel, and J. Peck (1996): "Demand Uncertainty, Inventories, and Resale Price Maintenance," *The Quarterly Journal of Economics*, 111, 885–913.

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

Dinan, T. M. (1993): "Economic Efficiency Effects of Alternative Policies for Reducing Waste Disposal," *Journal of Environmental Economics and Management*, 25, 242–256.

European Union (2024): "Regulation (EU) 2024/1781: Ecodesign Requirements for Sustainable Products," .

Fullerton, D. and A. Wolverton (2000): "Two Generalizations of a Deposit–Refund System," *American Economic Review*, 90, 238–242.

——— (2005): "The Two-Part Instrument in a Second-Best World," *Journal of Public Economics*, 89, 1961–1975.

Ino, H. (2011): "Optimal Environmental Policy for Waste Disposal and Recycling When Firms Are Not Compliant," *Journal of Environmental Economics and Management*, 62, 290–308.

Japan Book Publishers Association (2017): "Shuppan Shihyo (Publication Indicators)," Tokyo.

Japan Fair Trade Commission (2010): "Annual Report 2009," https://www.jftc.go.jp/info/nenpou/h21/index.html.

Jeon, J. (2022): "Learning and Investment under Demand Uncertainty in Container Shipping," *The RAND Journal of Economics*, 53, 226–259.

Kawaguchi, K., J. Qiu, and Y. Zhang (2022): "Competitive Effects of Resale Price Maintenance Through Inventory: Evidence from Publishing Industry," .

Nevo, A. (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 307–342.

Palmer, K. and M. Walls (1997): "Optimal Policies for Solid Waste Disposal Taxes, Subsidies, and Standards," *Journal of Public Economics*, 65, 193–205.

Porter, R. C. (2002): *The Economics of Waste*, New York: Routledge.

Tohan Corporation (2024): "Annual Report 2023," .