Decentralized Decision-Making in Retail Chains: Evidence from Inventory Management^{*}

Victor Aguirregabiria[†] University of Toronto, CEPR Francis Guiton [‡] University of Toronto

July 9, 2023

Abstract

This paper investigates the impact of decentralizing inventory decision-making in multiestablishment firms using data from a large retail chain. Analyzing two years of daily data, we find significant heterogeneity among the inventory decisions made by 634 store managers. By estimating a dynamic structural model, we reveal substantial heterogeneity in managers' perceived costs. Moreover, we observe a correlation between the variance of these perceptions and managers' education and experience. Counterfactual experiments show that centralized inventory management reduces costs by eliminating the impact of managers' skill heterogeneity. However, these benefits are offset by the negative impact of delayed demand information.

Keywords: Inventory management; Dynamic structural model; Decentralization; Information processing in organizations; Retail chains; Managerial skills; Store managers.

JEL codes: D22, D25, D84, L22, L81.

^{*}We are grateful for the valuable comments received from Heski Bar-Isaac, Loren Brandt, Avi Goldfarb, Jiaying Gu, George Hall, Jordi Mondria, Bob Miller, John Rust, Steven Stern, Junichi Suzuki, and Kosuke Uetake, as well as from seminar participants at Cambridge, Stony Brook, Toronto, UBC-Sauder, EARIE conference, IIOC conference, CEA conference, and the Georgetown conference honoring John Rust. The data used in this paper was obtained through a request made under the Canadian Access to Information Act. We sincerely appreciate the generous assistance provided by LCBO personnel.

[†]Department of Economics, University of Toronto. 150 St. George Street, Toronto, ON, M5S 3G7, Canada, victor.aguirregabiria@utoronto.ca.

[‡]Ph.D. Candidate, Department of Economics, University of Toronto. 150 St. George Street, Toronto, ON, M5S 3G7, Canada, francis.guiton@mail.utoronto.ca.

1 Introduction

Multi-establishment firms can adopt various decision-making structures ranging from centralized decisions at the headquarters to a more decentralized approach where decision authority is delegated to individual establishments. Determining the optimal decision-making process for a firm involves weighing different trade-offs. A decentralized decision-making structure empowers local managers to leverage valuable information specific to their respective stores. This information, which may be difficult or time-consuming to communicate to headquarters, can be utilized effectively at the local level. However, decentralization also entails granting autonomy to heterogeneous managers who possess varied skills. This heterogeneity can lead to suboptimal outcomes for the overall firm. When determining the degree of decentralization in their decision-making structure, multi-establishment firms must evaluate this trade-off. They need to consider the benefits of local knowledge and timely decision-making against the potential challenges posed by managerial heterogeneity.

In this study, we investigate the inventory management decisions made by store managers at the *Liquor Control Board of Ontario* (LCBO) and examine the impact on the firm's performance of delegating these decisions to the store level. The LCBO is a provincial government enterprise responsible for alcohol sales throughout Ontario. As a decentralized retail chain, each store has a degree of autonomy in its decision-making process. Specifically, store managers have discretion in two key areas of inventory management: assortment decisions (i.e., determining which products to offer) and replenishment decisions (i.e., when and how much to order for each product). Replenishment decisions involve forming expectations about future demand to determine the optimal order quantity and timing. To conduct our analysis, we utilize a comprehensive dataset obtained from the LCBO, encompassing daily information on inventories, orders, sales, stockouts, and prices for every store and product (SKU) from October 2011 to October 2013 (677 working days).¹ Additionally, we supplement our main dataset with information from LCBO reports and gather data on store managers' education and experience from professional networking platforms such as *LinkedIn*.

The LCBO data and framework provide a unique setting to study inventory management due to the simple pricing mechanism employed, where prices are set as a fixed markup over wholesale cost. This feature of the market allows us to focus specifically on the inventory setting problem without the need to incorporate a model where equilibrium prices are explicitly determined by inventory decisions. By abstracting from the complex relationship between prices and inventory decisions, we can concentrate our analysis on understanding the factors influencing inventory management within the LCBO retail chain.

¹These data were obtained under the Access to Information Act, with assistance from the LCBO personnel.

By employing descriptive evidence and the estimation of reduced-form models of inventory decision rules, we first show substantial heterogeneity across store managers' replenishment decisions. Observable store characteristics – such as demand level, size, category, and geographic location – explain less than half of the differences across store managers in their inventory decision rules.

To gain a deeper understanding of the factors contributing to this heterogeneity, we propose and estimate a dynamic structural model of inventory management. The model allows for differences across stores in demand, storage costs, stockout costs, and ordering costs. Leveraging the high-frequency nature of the daily data, we obtain precise estimates of holding cost, stockout cost, and ordering costs at both the individual product (SKU) and store levels. Our findings reveal significant heterogeneity across stores in all the revealed-preference cost parameters. Remarkably, observable store characteristics only account for less than 50% of this heterogeneity. Furthermore, we uncover a correlation between the remaining heterogeneity and managers' education and experience, suggesting that manager characteristics contribute to this residual variation. We interpret this unexplained heterogeneity as the result of local managers' idiosyncratic perceptions regarding store-level costs.

Using the estimated structural model, we quantify the impact of store manager heterogeneity on inventory outcomes. Overall, eliminating the idiosyncratic heterogeneity in cost parameters produces significant effects on inventory management. Specifically, we observe a 6-day increase in the waiting time between orders, a decrease in the average order amount equivalent to 1.5 days of average sales, and a 21% reduction in the inventory-to-sales ratio. However, the frequency of stockouts remains largely unaffected. These findings indicate that if the idiosyncratic component of costs represents a biased perception by store managers, it has a substantial negative impact on the firm's profitability. This is due to increased storage and ordering costs while having little effect on stockouts and revenue generation.

Finally, we conduct an evaluation of the effects associated with centralizing the decisionmaking of inventory management at the LCBO headquarters. To simulate this counterfactual experiment, we take into account information provided by company reports, which indicate that store-level sales information is processed by the headquarters with a one-week delay. The main trade-off examined in this experiment revolves around the fact that a centralized inventory management system eliminates the influence of store managers' heterogeneous skills and biased perceptions of costs. However, it also relinquishes the advantages derived from store managers' just-in-time information about demand and inventories. Our findings reveal that implementing a centralized inventory management system would result in a substantial reduction in ordering and storage costs, with an average cost decrease of 23% and a 3.7% reduction for the median store. Despite the significant cost reduction, this benefit is nearly completely offset by the negative impact on profits due to the delayed information about demand. Consequently, the net effect on profits is modest, with a mere 2% increase in annual profit for LCBO, equivalent to \$34 million. We further explore the implications of this trade-off for designing a more efficient inventory system that incorporates elements of both centralized and decentralized approaches.

This paper contributes to the growing empirical literature exploring the trade-offs between centralization and decentralization of decision-making in multi-division firms. Notably, DellaVigna and Gentzkow (2019) observe that most retail chains in the US employ uniform pricing across their stores, despite substantial differences in demand elasticities and potential gains from third-degree price discrimination. The authors discuss possible explanations for this phenomenon. In a study of a major international airline company, Hortacsu et al. (2021) analyze a pricing system that combines decision rights across different organizational teams. They find that despite employing advanced techniques, the pricing system fails to internalize consumer substitution effects, exhibits persistent biases in demand forecasting, and does not adapt to changes in opportunity costs. These inefficiencies are primarily attributed to limited coordination between teams. Examining decentralization decisions from a diverse set of firms across 11 OECD countries, Aghion et al. (2021) show that firms that delegate decision power from central headquarters to plant managers exhibited better performance during the Great Recession compared to similar firms with more centralized structures. The empirical evidence presented in their study supports the interpretation that the value of local information increases during turbulent economic times. Our paper contributes to this literature by being, to the best of our knowledge, the first empirical study to examine the trade-offs related to the (de)centralization of inventory management within retail chains, and specifically the role of store managers' heterogeneous skills.

This paper also contributes to the empirical literature on dynamic structural models of inventory behavior. Previous contributions in this area include works by Aguirregabiria (1999), Hall and Rust (2000), Kryvtsov and Midrigan (2013), and Bray et al. (2019) in the context of firm inventories, as well as Eberly (1994), Attanasio (2000), and Adda and Cooper (2000) in the domain of household purchases of durable products. We contribute to this literature by using high-frequency (daily) data at the granular store and product level to estimate cost parameters using a dynamic structural model.

Finally, our paper contributes to the literature on structural models with boundedly rational firms. Most of this literature studies firms' entry/exit decisions (Goldfarb and Xiao, 2011; and Aguirregabiria and Magesan, 2020), pricing decisions (Huang et al., 2022; Ellison et al., 2018), and bidding behavior in auctions (Hortaçsu and Puller, 2008; Doraszelski et al., 2018; Hortaçsu et al., 2019). To the best of our knowledge, our paper represents the first investigation into bounded rationality in firms' inventory decisions. This research also contributes to the existing literature by combining revealed-preference estimates of managers' perceived costs with a decomposition of these costs into the objective component explained by store characteristics and the subjective component associated with managers' education and experience.

The rest of the paper is organized as follows. Section 2 describes the institutional background of the LCBO and presents the dataset and descriptive evidence. Section 3 presents evidence of managers following (S, s) decision rules and illustrates the heterogeneity in these (S, s) thresholds across store managers. Section 4 presents the structural model and its estimation. The counterfactual experiments to evaluate the effects of decentralization are described in section 5. We summarize and conclude in Section 6.

2 Firm and data

2.1 LCBO retail chain

History. LCBO was founded in 1927 as part of the passage of the Ontario Liquor Licence Act.² This act established that LCBO was a crown corporation of the provincial government of Ontario. Today, the wine retail industry in Ontario is a triopoly - consisting of 634 LCBO stores, 164 Wine Rack stores, and 100 Wine Shop stores. Despite its government ownership, LCBO is a profit maximizing company. As described in its governing act, part of its mandate is "generating maximum profits to fund government programs and priorities".³

Store managers. According to the LCBO, store managers are responsible for managing their "store, sales and employees to reflect [their] customers' needs and business goals", with a particular focus on the inventory management of their store. Managers must oversee their store's overall inventory level to ensure that daily demand is met. For incentive purposes, part of the store managers' remuneration depends on the overall sales performance of their store. The managers' pay is therefore closely tied to their stores' profits. In order to satisfy daily demand, managers periodically restock their shelves by ordering products from the nearest distribution center. The order is then delivered by trucks according to a pre-determined

²The information in this section originates from various archived documents from the LCBO. General information about the company and its organization is based on the company's annual reports Liquor Control Board of Ontario (2012) and Liquor Control Board of Ontario (2013), and the collective agreement between the LCBO and OPSEU. Information regarding headquarters' order recommendations (*Suggested Order Quantities*, SOQs) originates from the report Liquor Control Board of Ontario (2016). Additional information regarding the role of store managers originates from an interview we conducted with an LCBO store manager from a downtown Toronto store.

 $^{^{3}}$ See https://www.lcbo.com/content/lcbo/en/corporate-pages/about/aboutourbusiness.html.

route and schedule. In addition to the ordering decisions, managers are also responsible for their store's product assortment, as they must decide which products to offer at their store in order to cater to local demand. Inventory management at LCBO, therefore, entails a dual responsibility for store managers: providing products that are in high demand and keeping these products stocked on the shelves.

Classification of stores. The LCBO classifies its stores into six categories, AAA, AA, A, B, C, and D, ranging from the highest to the lowest. These classifications primarily reflect variations in store size and product assortment. However, there are also differences in the consumer shopping experience across these classifications, with the AAA and AA stores being considered flagship stores.

Headquarters. Headquarters are in charge of the assignment of store managers across the different stores. Assignments are occasionally shuffled due to managers being promoted (demoted) to higher- (lower-) classified stores, with seniority being a main factor in the promotion decision. Another responsibility of headquarters is to assist store managers in their inventory decisions. Headquarters use forecasting techniques to provide recommendations to managers regarding how much to order for each product at their store. In the company's internal jargon, these recommendations are referred to as *Suggested Order Quantities* (SOQs). For each store and product, headquarters generate order recommendations based on the previous week's sales and inventory information.⁴ Importantly, this entails that headquarters process store-level information with a weekly delay. Headquarters do not use *just-in-time* daily information that store managers may be using in their replenishment decisions. This informational friction may play a role in the optimal allocation of decision rights.

Pricing. LCBO and its competitors are subject to substantial pricing restrictions. Prices must be the same across all stores in all markets for a given store-keeping unit (SKU). There is no price variation across the LCBO and its competitors. Retail prices are determined on a fixed markup over the wholesale price set by wine distributors. Furthermore, the percentage markup applies to all the SKUs within broadly defined categories.⁵

No franchising system. The LCBO operates its stores without adopting a franchising system. Instead, all store managers are employees of the LCBO. As a result, store managers are not required to pay any franchise fees, fees per order, or any other types of fees to the firm. The absence of a franchising system ensures that the store managers operate within the organizational structure of the LCBO as employees, without the additional financial obligations associated with a franchising arrangement.

 $^{^{4}}$ More specifically, order recommendations are a function of the Average Rate of Sale (ARS) of the product from the previous week, and of seasonal brand factors.

⁵See Aguirregabiria et al. (2016) for further details about markups at LCBO.

Union. Most employees at the LCBO are unionized under the Ontario Public Service Employees Union (OPSEU). As of 2022, the OPSEU "represents more than 8,000 workers at the Liquor Control Board of Ontario", with their main goal being to "establish and continue harmonious relations between the [LCBO] and the employees". Members of the union include store managers, retail workers, warehouse workers, and corporate workers.

2.2 Data from LCBO

Our analysis combines three data sources: the main dataset provided by the LCBO; data on store managers' experience and education collected from the social media platform *LinkedIn*; and consumers' socioeconomic characteristics from the 2011 Census of Population.

We use a comprehensive and rich dataset obtained from the LCBO, encompassing daily information on inventories, sales, deliveries, and prices of every product sold at each LCBO store. The dataset covers a period of two years, specifically from October 2011 to October 2013, spanning a total of 677 days. With a total of 634 stores operating across Ontario and an extensive product range consisting of over 20,000 different items, our dataset comprises approximately 720 million observations. Moreover, the dataset includes additional valuable information, such as product characteristics, store characteristics (including location, size, and store category), and the store manager's name.

Table 1 presents summary statistics. The average store has an assortment of 2,029 items. Weekly sales per store amount to 12,909 units, translating to an average of 6.36 units sold per item. The average weekly revenue per store is \$162,250, resulting in \$80 in revenue per item per week and \$12.6 in revenue per unit sold. Regarding deliveries, stores receive shipments at approximately 5.37 days per week, with total weekly deliveries containing an average of 12,258 units. Stockout events occur, on average, 415 times per week across all stores.

Notably, these figures exhibit significant variation across different store types. Larger stores, as expected, generate higher weekly revenues. For instance, AAA and AA stores generate average weekly revenues of \$874,367 and \$518,067, respectively, while C and D stores generate average weekly revenues of \$59,328 and \$23,913, respectively. Stockout events appear to be more prevalent in larger stores compared to smaller ones. The average AAA store experiences 1,203 stockout events per week, whereas the average D store encounters 204 stockout events per week. Additionally, larger stores tend to place orders more frequently and in larger quantities. On average, AAA and AA stores receive orders 6.30 and 6.27 days per week, totaling 51,340 and 35,711 units, respectively. In contrast, C and D stores receive orders 5.09 and 3.83 days per week, amounting to 4,395 and 1,660 units, respectively.

The bottom panel in Table 1 presents inventory-to-(daily)sales ratios and ordering fre-

quencies. These statistics are closely related to the (S, s) decision rules that we analyze in Section 3. At the store-product level, the *inventory-to-sales ratio* before and after an order corresponds to the thresholds s and S, respectively. On average, stores maintain enough inventory to meet product demand for approximately 23 days, initiate an order when there is sufficient inventory for about 9 days, and the ordered quantity covers sales for around 18 days. In terms of ordering frequency, the average store and product place an order approximately once every two weeks, equivalent to a frequency of $0.07 \simeq 1/14$.

Average inventory-to-sales ratios tend to decrease with store size/type, although this difference is influenced by the composition effect arising from varying product assortments across store types. For the remainder of the paper, to account for this composition effect and reduce the computational burden of estimating our model across numerous products, we focus on a working sample consisting of a few products carried by all stores. This approach allows us to manage the complexity associated with different product assortments while maintaining the robustness of our analysis.

			Type of	store			
	All	AAA	AA	A	B	C	D
	Mean	Mean	Mean	Mean	Mean	Mean	Mean
	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)
Number of Observations							
Number of Stores	634	5	25	148	157	164	135
Number of Unique Products	22,327	18,200	$18,\!879$	20,860	17,769	$13,\!527$	9,198
Number of Days	677	676	676	676	677	676	675
Sales & Stockouts Per Store							
Revenue per week (\$)	162,249	$874,\!367$	518,067	$320,\!263$	162,588	59,327	23,912
	(158, 268)	(255, 172)	(111, 241)	(75, 387)	(47, 385)	(21, 159)	(10,503)
Units sold per week	12,908	57,160	38,369	$25,\!665$	14,068	4,522	1,730
	(11, 885)	(8,992)	(5,399)	(5,405)	(4,554)	(1, 820)	(819)
Stockout events per week	414	1,202	912	699	371	273	204
	(326)	(429)	(279)	(306)	(152)	(252)	(200)
Inventories Per Store							
Number of products	2,028	6,477	4,852	$3,\!555$	2,102	1,098	742
	(1,416)	(1,018)	(745)	(937)	(706)	(291)	(190)
Delivery days per week	5.37	6.30	6.27	6.23	6.06	5.09	3.83
	(1.08)	(0.01)	(0.05)	(0.11)	(0.31)	(0.66)	(0.88)
Delivery units per week	12,257	51,339	35,710	$24,\!350$	$13,\!476$	$4,\!395$	$1,\!660$
	(11,108)	(6,872)	(4, 848)	(5,063)	(4, 478)	(1,792)	(814)
Inventory Ratios Per Store							
Inventory to (daily) sales ratio	23.31	15.67	15.74	16.89	18.97	25.21	34.23
	(9.86)	(2.16)	(2.74)	(3.32)	(5.24)	(7.20)	(11.85)
Inventory to sales ratio after order	18.48	8.78	10.34	12.18	15.13	21.65	26.87
	(9.16)	(0.62)	(1.55)	(2.76)	(4.40)	(7.38)	(11.72)
Inventory to sales ratio before order	8.62	5.31	6.83	7.82	8.88	9.23	8.92
	(3.59)	(0.60)	(1.40)	(2.39)	(2.97)	(3.66)	(4.99)
Ordering frequency	0.07	0.10	0.10	0.09	0.08	0.06	0.04
	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)

Table 1: Summary Statistics

2.3 Working sample

In our econometric models, estimated in Sections 3 and 4, the parameters are unrestricted at the store-product level. Considering that our dataset comprises nearly 2 million store-product pairs, estimating these models for every store-product combination would be exceedingly time-consuming. To save time while maintaining the integrity of our analysis, we have employed a different approach. Specifically, we estimate the econometric models for every store in our dataset but limited the analysis to a selected subset of 5 products.

We employ two criteria to determine the product basket for our analysis. Firstly, we select products that exhibit high sales across all LCBO stores, ensuring that their inventory decisions significantly impact the firm's overall profitability. Secondly, we include products from each broad category to account for product-level heterogeneity. These categories encompass white wine, red wine, vodka, whisky, and rum. Table 2 provides an overview of the five selected products that satisfy these criteria. By focusing on this subset, we capture a diverse range of products that are representative of the different categories while also being impactful in terms of their sales performance.

	SKU	SKU	SKU	SKU	SKU
	#67	#117	#340380	#550715	#624544
Product Information					
Name	Smirnoff	Bacardi Superior	Two Oceans	Forty Creek	Yellow Tail
	Vodka	White Rum	Sauvignon Blanc	Barrel Select Whisky	Shiraz
Category	Vodka	Rum	White Whine	Whisky	Red Wine
Average Retail Price (\$)	25.28	24.93	9.98	25.86	11.83

Table 2: SKUs - Working Sample

Table 3 provides summary statistics for our working sample. On average, the stores in our working sample sell 91 units per week, equivalent to 18 units per SKU. The average weekly revenue per store is \$1,687, resulting in \$347 in revenue per SKU per week and \$18 in revenue per unit sold. Regarding deliveries, stores in our working sample receive shipments approximately 2 days per week, with each delivery containing an average of 88 units. The average number of stockout events per week per store is 0.37.

Similar to Table 1, these numbers exhibit significant variation across different store types. Larger stores generate higher average weekly revenues compared to smaller ones. For instance, the average AAA and AA stores generate weekly revenues of \$5, 197 and \$4, 511, respectively, while the smaller C and D stores generate average weekly revenues of \$821 and \$302, respectively. Contrary to the full sample, stockout events appear to occur more frequently in

smaller stores than in larger ones within our working sample. The average D store experiences 0.56 stockout events per week, while the average AAA store encounters 0.14 stockout events per week. Delivery patterns in our working sample follow a similar trend to Table 1, with larger stores placing larger and more frequent orders. The average AAA store receives deliveries 4 days per week, totaling 272 units per week, whereas the average D store only receives deliveries 0.6 days per week, amounting to 15 units per week.

			T	ype of stor	re		
	All	AAA	AA	A	B	C	D
	Mean	Mean	Mean	Mean	Mean	Mean	Mean
	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)
Number of Observations							
Number of Stores	634	5	25	148	157	164	135
Number of Unique Products	5	5	5	5	5	5	5
Number of Days	677	676	676	676	677	676	675
Sales & Stockouts Per Store							
Revenue per week (\$)	$1,\!687$	$5,\!197$	4,511	$3,\!192$	1,801	821	302
	(1,402)	(1, 487)	(1,415)	(835)	(660)	(443)	(200)
Units sold per week	91	274	252	173	98	43	15
	(75)	(51)	(65)	(39)	(33)	(22)	(10)
Stockouts events per week	0.37	0.14	0.13	0.21	0.28	0.49	0.56
	(0.39)	(0.07)	(0.11)	(0.15)	(0.23)	(0.35)	(0.60)
Inventories Per Store							
Number of products Offered	4.85	5	5	5	5	4.90	4.42
	(0.51)	(0)	(0)	(0)	(0)	(0.34)	(0.92)
Delivery days per week	1.94	4.04	3.79	3.20	2.28	1.20	0.62
	(1.15)	(0.29)	(0.24)	(0.49)	(0.56)	(0.51)	(0.35)
Delivery units per week	88	272	243	167	95	42	15
	(73)	(55)	(68)	(38)	(33)	(21)	(9)
Inventory Ratios Per Store							
Inventory to sales ratio	20.97	28.27	25.18	24.93	24.07	18.80	14.59
	(7.03)	(8.06)	(7.18)	(6.44)	(6.91)	(4.84)	(3.46)
Inventory to sales ratio after order	20.80	21.02	18.85	20.09	22.48	21.46	19.17
	(4.94)	(10.69)	(4.11)	(4.46)	(6.18)	(4.07)	(3.81)
Inventory to sales ratio before order	13.06	17.48	15.77	16.61	16.25	10.90	7.43
	(5.33)	(7.10)	(3.42)	(3.83)	(4.72)	(3.44)	(2.95)
Ordering frequency	0.15	0.31	0.30	0.25	0.17	0.09	0.05
	(0.09)	(0.06)	(0.06)	(0.06)	(0.05)	(0.04)	(0.03)

 Table 3: Summary Statistics - Working Sample

Figure 1 shows strong heterogeneity across stores in several measures related to inventory management of the five products in our working sample. The figures in panels (a) to (f) are inverse cumulative distributions over stores, together with their 95% confidence bands.⁶

 $^{^6 {\}rm For}$ every store, the 95% confidence interval is based on the construction of store-product-specific rates. The 95% confidence interval is determined by percentiles 2.5% and 97.5% in this distribution.

Panel (a) presents the distribution of the *stockout rate*. For store i, we have:

$$Stockout \ rate_i = \frac{\# \ (product, day) \ observations \ with \ stockout \ in \ store \ i}{\# \ (product, day) \ observations \ for \ store \ i}$$
(1)

The figure shows a large spread of stockout rates: the 10th and 90th percentiles are 0.20% and 2.82%, respectively. Panel (b) shows substantial heterogeneity across stores in the *revenue-loss per-product-day* generated by stockouts. Indexing products by j, and using $\mathcal{J} = 5$ to represent the number of products, the *revenue-loss* for store i is:

Revenue
$$loss_i = \frac{1}{J} \sum_{j=1}^{J} Stockout \ rate_{i,j} \times Average \ daily \ revenue \ without \ stockouts_{i,j}$$
(2)

The 10th and 90th percentiles are \$0.06 and \$1.03 per product-day, respectively. Aggregated at the annual level and over all the products offered in a store, they imply an average annual revenue-loss of approximately \$44,000 at the 10th percentile and \$760,000 at the 90th percentile. Panel (c) presents the ordering frequency of stores in our sample calculated as:

$$Ordering \ frequency_i = \frac{\# \ (product, day) \ observations \ with \ an \ order \ in \ store \ i}{\# \ (product, day) \ observations \ for \ store \ i}$$
(3)

This ordering rate varies significantly across stores, with the 10^{th} percentile being 3.59% and the 90th being 28.41%. Panels (d) to (f) present the empirical distributions for the *inventory-to-sales ratio*, for this ratio just before an order (a measure of the lower threshold s), and for this ratio just after an order (a measure of the upper threshold S). Indexing days by t:

$$Inventory - to - (daily)sales - ratio_i = \frac{\sum_{j=1}^J \sum_{t=1}^T Inventory_{i,j,t}}{\sum_{j=1}^J \sum_{t=1}^T Units \ sold_{i,j,t}}$$
(4)

The distribution of the *inventory-to-sales ratio* shows that stores at the 10^{th} and 90^{th} percentiles hold inventory for 11 days and 33 days of average sales, respectively. For the upper threshold S (in Panel (e)), the values of these percentiles are 9 and 37 days, and for the lower threshold s (in Panel (f)) they are 5 and 19 days.

Given these substantial differences in inventory outcomes across stores, it is interesting to explore how they vary together. We present these correlations in Appendix A.1 (Figure 10). The strongest correlation appears for the positive relationship between our measures of the thresholds S and s. This correlation can be explained by store heterogeneity in storage costs: a higher storage cost implies lower values of both s and S. We confirm this conjecture in the estimation of the structural model in Section 4.



Figure 1: Empirical Distribution (Inverse CDF) of Inventory Outcomes Across Stores

(a) Stockout Frequency

(b) Stockout Cost

To investigate the possibility of stockouts at the warehouse level and their potential impact on store-level stockouts, we also analyze aggregate daily deliveries from the warehouse to all 634 stores. Given that the products in our working sample are popular items, we interpret a zero value in aggregate daily deliveries as a stockout event at the warehouse. The results of our analysis, presented in Section A.4 in the Appendix, reveal that warehouse stockouts are negligible for our working sample. For each product within the working sample, warehouse stockout events occur on no more than 3 out of the 677 days in the sample, which accounts for less than 0.5% of the observed period. These findings suggest that stockouts observed at the store level are primarily driven by factors other than warehouse-level stockouts.

2.4 Data on store managers' education and experience

In addition to the main dataset, we enhance our analysis by incorporating information on store managers' human capital. Leveraging the professional social networking platform *LinkedIn*, we gather data on the education and experience of store managers from their public profiles. Out of the 634 store managers in our dataset, 600 are identifiable by name in the LCBO's records.⁷ Within this subset, we were able to locate public *LinkedIn* profiles for 143 managers, allowing us to retrieve valuable information about their educational background and work experience.

Table 4 presents summary statistics for the variables related to store managers' educational background and work experience. Notably, we observe a pattern in which managers with greater experience and higher educational attainment tend to be assigned to higherclassified stores. This finding suggests a positive correlation between manager characteristics and store classification, indicating that the LCBO may allocate more experienced and highly educated managers to stores of higher importance or larger scale. In Appendix A.2, we provide more detailed information on the positive relationship between manager characteristics and the classification of stores within the LCBO retail chain.

⁷During our sample period, some stores are overseen by interim managers who are not identified by name in our main dataset.

			Ty	pe of sta	ore		
Statistic	All	AAA	AA	A	B	C	D
$\# \ of \ stores$	634	5	25	148	157	163	125
(%)	(100.0)	(0.8)	(4.0)	(23.3)	(24.7)	(25.9)	(21.3)
# managers with observed education	86	1	8	30	32	5	10
(%)	(100.0)	(1.2)	(9.3)	(34.9)	(37.2)	(5.8)	(11.6)
# managers with observed LCBO exper.	136	3	9	48	44	20	12
(%)	(100.0)	(2.2)	(6.6)	(35.3)	(32.4)	(14.7)	(8.8)
# managers with observed industry exper.	72	1	7	30	26	5	3
(%)	(100.0)	(1.4)	(9.7)	(41.7)	(36.1)	(6.9)	(4.2)
$Dummy \ highest \ degree = high \ school \ (Mean)$	0.093	0.000	0.000	0.067	0.094	0.200	0.200
$Dummy \ highest \ degree = college \ (Mean)$	0.395	0.000	0.625	0.333	0.375	0.600	0.400
$Dummy \ highest \ degree = university \ (Mean)$	0.512	1.000	0.375	0.600	0.531	0.200	0.400
Years LCBO exper. (Mean)	13.7	10.0	12.3	12.2	9.7	24.7	17.8
Years industry (non-LCBO) exper. (Mean)	9.8	15.0	8.0	11.2	9.8	5.8	5.3

Table 4: Experience and Education of Store Managers by Store Type

3 (S,s) decision rules

3.1 Model

In this section, we study store managers' inventory behavior through the eyes of (S, s) decision rules. In its simplest form, this decision rule involves time-invariant threshold values. When assuming lump-sum (fixed) ordering costs, quasi-K-concavity of the profit function with respect to orders, and time-invariant expected demand, the profit-maximizing inventory decision rule follows a (S, s) structure (Arrow et al., 1951; Scarf, 1959; Denardo, 1981). The (S, s) rule is characterized by two threshold values: a lower threshold denoted as s, which represents the stock level that triggers a new order (known as the "safety stock level"), and an upper threshold denoted as S, which indicates the stock level to be achieved when an order is placed. Thus, if k_t represents the stock level at the beginning of day t, and y_t represents the orders placed on day t, the (S, s) rule can be defined as follows:

$$y_t = \begin{cases} S - k_t & if \ k_t \le s \\ 0 & otherwise \end{cases}$$
(5)

Hadley and Whitin (1963) and Blinder (1981) provide comparative statics for the thresholds (S, s) as functions of the structural parameters in the firm's profit function. They provide the following results:

$$S = f_S\left(d^e, \ \gamma^h, \ \frac{\gamma^f}{\gamma^c}, \ \frac{\gamma^z}{\gamma^c}\right), \quad s = f_s\left(d^e, \ \gamma^h, \ \frac{\gamma^f}{\gamma^c}, \ \frac{\gamma^z}{\gamma^c}\right), \quad S - s = f_{S-s}\left(d^e, \ \gamma^h, \ \frac{\gamma^f}{\gamma^c}, \ \frac{\gamma^z}{\gamma^c}\right) + - ? + + - - + ?$$
(6)

where d^e represents expected demand, γ^h is the inventory holding cost per period and per unit, γ^f is the fixed (lump-sum) ordering cost, γ^c is the unit ordering cost, and γ^z is the stockout cost per period. These γ 's are the parameters in the structural model that we estimate in Section 4. We investigate the predictions of equation (6) in Section 5.

The optimality of the (S, s) decision rule extends to models with state variables that evolve over time according to exogenous Markov processes. Let \mathbf{z}_t denote the vector of these exogenous state variables, which can include factors influencing demand, unit ordering costs, wholesale prices, and the product's retail price (when taken as given by the store manager), as is the case in our problem. The optimal decision rule in this context follows a (S_t, s_t) structure, where the thresholds S_t and s_t are time-invariant functions of these state variables: $S_t = S(\mathbf{z}_t)$ and $s_t = s(\mathbf{z}_t)$.

In this section, our empirical approach is inspired by the work of Eberly (1994), Attanasio (2000), and Adda and Cooper (2000). These studies utilize household-level data on durable product purchases, specifically automobiles, to estimate (S_t, s_t) decision rules. In these decision rules, the thresholds are functions of household characteristics, prices, and aggregate economic conditions. This approach can be seen as a "semi-structural approach," where the use of (S_t, s_t) rules is motivated by a dynamic programming model of optimal behavior. However, the specification of the thresholds as functions of state variables does not explicitly incorporate the structural parameters of the model.

In Section 4, we present a full structural approach that explicitly incorporates the structural parameters of the model. Furthermore, in Section 5, we utilize the estimated structural model to conduct counterfactual policy experiments, which address the questions that motivated this paper. However, before delving into the full structural analysis, we find it valuable to explore the data using a more flexible empirical framework that remains consistent with the underlying structural model. We investigate heterogeneity between store managers' inventory decisions by estimating (S_t, s_t) rules at the store-product level. These decision rules are consistent with our structural model but they are more flexible. This allows us to gain insights and assess the suitability of the (S_t, s_t) decision rules in capturing the inventory behavior of store managers within the LCBO retail chain. Given that our dataset contains 677 daily observations for every store and product, and that the ordering frequency in the data is high enough to include many orders per storeproduct, we can estimate the parameters in the (S_t, s_t) decision rules at the store-product level. In this section, we omit store and product sub-indexes in variables and parameters, but it should be understood that these sub-indexes are implicitly present.

We consider the following specification for the (S_t, s_t) thresholds:⁸

$$\begin{cases} S_{t} = \exp\{\beta_{0}^{S} + \beta_{d}^{S} \ln d_{t}^{e} + \beta_{p}^{S} \ln p_{t} + u_{t}^{S}\} \\ s_{t} = \exp\{\beta_{0}^{s} + \beta_{d}^{s} \ln d_{t}^{e} + \beta_{p}^{s} \ln p_{t} + u_{t}^{s}\} \end{cases}$$
(7)

where p_t is the product's retail price, d_t^e is the expected demand, and u_t^s and u_t^S represent state variables which are known to the store manager but are unobservable to us as researchers.⁹ The vector of exogenous state variables is $\mathbf{z}_t = (d_t^e, p_t, u_t^s, u_t^S)$. The β 's are reduced form parameters which are constant over time but vary freely across stores and products and are functions of the structural parameters that we present in our structural model in Section 4.

Our measure of expected demand is based on an LCBO report regarding the information that headquarters use to construct order recommendations for each store (Liquor Control Board of Ontario, 2016). Relying on this report, we assume that store managers obtain predictions of demand for each product at their store using information on the product's retail price (p_t) , the average daily sales of the product over the last seven days, (that we represent as $Q_t^{[-7,-1]}$), and seasonal dummies.

Since the observed quantity sold q_t has discrete support $\{0, 1, 2, ...\}$, we consider that demand has a Negative Binomial distribution where the logarithm of expected demand at period t has the following form:

$$\ln d_t^e = \ln \mathbb{E}\left(q_t \mid p_t, Q_t^{[-7, -1]}\right) = \boldsymbol{\alpha}' \ h\left(\ln p_t, \ln Q_t^{[-7, -1]}\right) \tag{8}$$

where q_t is the quantity sold of the store-product at day t, $\boldsymbol{\alpha}$ is a vector of parameters that are constant over time but vary freely across store-products, and $h\left(\ln p_t, \ln Q_t^{[-7,-1]}\right)$ is a vector of monomial basis in variables $\ln p_t$ and $\ln Q_t^{[-7,-1]}$.

⁸We attempted to incorporate demand volatility, represented by $\ln \sigma_t^2$, as an explanatory variable in the decision rule. However, we encountered high collinearity between the time series of $\ln d_t^e$ (expected demand) and $\ln \sigma_t^2$, making it challenging to estimate their separate effects on the thresholds. It is worth noting that according to the Negative Binomial distribution, $\ln \sigma_t^2 = \ln d_t^e + \ln(1 + \alpha d_t^e)$. Consequently, we can interpret the effect of $\ln d_t^e$ on the (S_t, s_t) thresholds as the combined impact of both expected demand and volatility.

⁹For instance, u_t^s and u_t^S may include shocks in fixed and variable ordering costs, or measurement error in our estimate of expected demand.

We denote equation (8) as the sales forecasting function. It deserves some explanation. First, it is important to note that this is not a demand function. For this inventory decision problem, managers do not need to know the demand function but only the best possible predictor of future sales given the information they have. Second, this specification ignores substitution effects between products within the same category or across categories. Ignoring substitution effects in demand is fully consistent with LCBO's report and with the firm's price setting, which completely ignores these substitution effects (see Aguirregabiria et al., 2016).¹⁰ In the Appendix (Section A.3), we present a summary of the estimation results of the sales forecasting function for every store and every product in our working sample.

The (S_t, s_t) model in equations (5) and (7) implies that the decision of placing an order $(y_t > 0)$ or not $(y_t = 0)$ has the structure of a linear-in-parameters binary choice model.

$$\mathbb{1}\{y_t > 0\} = \mathbb{1}\{b_0^s + b_k^s \ln k_t + b_d^s \ln d_t^e + b_p^s \ln p_t + \widetilde{u}_t^s \ge 0\},\tag{9}$$

where $\mathbb{1}\{.\}$ is the indicator function; $\widetilde{u}_t^s \equiv u_t^s/\sigma_{u^s}$ is the standardized version of u_t^s , as σ_{u^s} is the standard deviation of u_t^s ; and there is the following relationship between β^s and b^s parameters: $b_k^s = -1/\sigma_{u^s}$; $b_0^s = \beta_0^s/\sigma_{u^s}$; $b_d^s = \beta_d^s/\sigma_{u^s}$; and $b_p^s = \beta_p^s/\sigma_{u^s}$. These expressions show that, given the parameters b^s , we can identify the parameters β^s and σ_{u^s} . We assume that \widetilde{u}_t^s has a Normal distribution, such that equation (9) is a Probit model, and we estimate the parameters b^s by maximum likelihood.

Our (S_t, s_t) model also implies that in days with positive orders $(y_t > 0)$ the logarithm of the total quantity offered, $\ln(k_t + y_t)$, is equal to the logarithm of the upper-threshold, $\ln(S_t)$, and this implies the following linear-in-parameters (censored) regression model:

$$\ln(k_t + y_t) = \beta_0^S + \beta_d^S \ln d_t^e + \beta_p^S \ln p_t + u_t^S \quad \text{if } y_t > 0.$$
(10)

Equation (10) includes the selection condition $y_t > 0$. That is, the upper-threshold S_t is observed only when an order is placed. This selection issue implies that OLS estimation of equation (10) yields inconsistent estimates of the parameters and the threshold itself. However, the (S_t, s_t) model implies an exclusion restriction that provides identification of the parameters in equation (10). The inventory level k_t affects the binary decision of placing an order or not (as shown in equation (9)), but conditional on placing an order, it does not affect the value of the upper-threshold in the right-hand-side of regression equation (10).

¹⁰Recent papers show that the pricing decisions of important multi-product firms do not internalize substitution or cannibalization effects between the firm's own products. See, for instance, Hortaçsu et al. (2021)'s study of the pricing system of a large international airline company, DellaVigna and Gentzkow (2019) and Hitsch et al. (2021) on uniform pricing at US retail chains, Cho and Rust (2010) on pricing of car rentals, or Miravete et al. (2020) for liquor stores in Pennsylvania.

Therefore, using (9) as the selection equation, we can identify the parameters β^S in (10) using a Heckman two-step approach.¹¹

Part of the variation in parameter estimates across stores and products is attributable to estimation error rather than genuine heterogeneity. For any given parameter $b_{i,j}$, where *i* and *j* represent store and product indices respectively, let $\hat{b}_{i,j}$ denote its consistent and asymptotically normal estimate, with an asymptotic variance of $\sigma_{i,j}^2$. Using this asymptotic distribution, we can establish a relationship between the variances of $\hat{b}_{i,j}$ and $b_{i,j}$ across stores and products: $Var(\hat{b}_{i,j}) = Var(b_{i,j}) + \mathbb{E}(\sigma_{i,j}^2)$, where $\mathbb{E}(\sigma_{i,j}^2)$ represents the mean of asymptotic variances $\sigma_{i,j}^2$ across stores and products. Since $\mathbb{E}(\sigma_{i,j}^2) > 0$, this equation demonstrates that $Var(\hat{b}_{i,j})$ overestimates the true dispersion $Var(b_{i,j})$. To mitigate this excess dispersion or spurious heterogeneity arising from estimation error, we employ a *shrinkage estimator*. The details of this estimator are described in section A.8 in the Appendix.

3.2 Estimation of (S,s) thresholds

Figure 2 presents the average estimates of parameters b_0^s , b_k^s , b_d^s , and b_p^s in the lower threshold for each store, where the average is obtained over the five products in the working sample. We sort stores from the lowest to the largest average estimate such that this curve is the inverse CDF of the average estimate. The red-dashed band around the median of this distribution is the 95% confidence band under the null hypothesis of homogeneity across stores.¹² The signs of the parameter estimates are for the most part consistent with the predictions of the model. These distributions show that the parameter estimates vary significantly across stores. For b_0^s , b_k^s , b_d^s , and b_p^s , we have that 95%, 95%, 98%, and 95% of stores, respectively, lie outside of the Bonferroni confidence interval.

Figure 3 presents the inverse CDF of the store-specific average of the parameters β_0^S , β_d^S , and β_p^S in the upper threshold, as well as the Bonferroni 95% confidence interval under the null hypothesis of homogeneity. As expected, we have strong evidence of heterogeneity in our estimates. For β_0^S , β_d^S , and β_p^S , approximately 96%, 97%, and 97% of stores lie outside of the confidence bands, respectively.

¹¹This exclusion restriction in (S, s) models have been pointed out by Bertola et al. (2005).

¹²The reported 95% confidence interval incorporates the Bonferroni correction for multiple testing. In this context, the implicit null hypothesis is that every store does not differ significantly from the average store. By applying the Bonferroni correction, we account for the increased probability of observing a significant difference by chance when conducting multiple tests.



Figure 2: Empirical Distribution (Inverse CDF) of Estimates b^s for Lower s Threshold

Figure 3: Empirical Distribution (Inverse CDF) of Estimates β^{S} for Upper S Threshold



Given that we have estimates at the store-product level, we can also explore the heterogeneity in these estimates within stores. Table 5 presents a decomposition of the variance of parameter estimates into within-store (between products) and between-store variance. The parameters associated with expected demand and the lower threshold inventory parameter show a between-store variance that is at least as large as the within-store variance. For the constant parameters and the price parameters, the variance is larger across products.

Variance	b_0^s	b_k^s	b_d^s	b_p^s	eta_0^S	eta_d^S	$oldsymbol{eta}_p^S$
Between-store Variance	69.41	0.02	0.28	8.54	110.18	0.27	12.01
Within-store Variance	81.01	0.02	0.07	9.78	183.81	0.11	19.35

Table 5: Variance Decomposition of Parameter Estimates

The parameter estimates imply values for the (S_t, s_t) thresholds. In section A.5 in the Appendix, we investigate heterogeneity across stores in the estimated thresholds. We find strong between-store heterogeneity, especially in the lower threshold s.

4 Structural model

We propose and estimate a dynamic structural model of inventory management. A pricetaking store sells a product and faces uncertain demand. The store manager orders the product from the retail chain's warehouse, and any unsold product rolls over to the next period's inventory. The store-level profit function incorporates four store-specific costs associated with inventory management: per-unit inventory holding cost $(\gamma_{i,j}^h)$, stockout cost $(\gamma_{i,j}^z)$, fixed ordering cost $(\gamma_{i,j}^f)$, and per-unit ordering cost $(\gamma_{i,j}^c)$.

4.1 Non-separability of inventory decisions across products

At LCBO, store managers are responsible for making inventory decisions for thousands of products. These decisions are not independent of each other due to various factors. Firstly, when a product experiences a stockout, consumers may opt to substitute it with a similar available product. As a result, the cost of a stockout for the store is influenced by the availability of substitute products. Secondly, storage costs are affected by the total volume of items and units in the store, which means the inventory management of one product can impact the storage costs of other products. Lastly, the store's ordering cost is dependent on the cost of filling a truck with units of multiple products and transporting them from the warehouse to the store. The decision to order a truck and the cost associated with it are not separable across products.

We can think of a model that fully accounts for the inter-dependence inventory decisions across products. In this model, a store manager maximizes the aggregate profit from all the products taking into account the substitutability of similar products under stockouts and subject to two overall store-level constraints: the total storage capacity of the store and the capacity of the delivery truck. Our store-product level model aligns with this multipleproduct inventory management framework.

By using duality theory, we can show that the marginal conditions for optimality in the multiple-product model are equivalent to those in our single-product model, under appropriate interpretations of our store-product level structural parameters. The inventory holding cost parameter $\gamma_{i,j}^h$ represents the shadow price or Lagrange multiplier associated with the storage capacity constraint at store *i*. Similarly, the fixed and unit ordering costs, $\gamma_{i,j}^f$ and $\gamma_{i,j}^c$ respectively, reflect the shadow prices of the truck's capacity constraint at the extensive and intensive margins. Lastly, the stockout cost parameter $\gamma_{i,j}^z$ accounts for the impact of consumer substitution within the store when a stockout occurs for product *j* at store *i*.

In estimating our model, our approach is valid as long as these Lagrange multipliers do not exhibit significant variation over time. For our counterfactual experiments, we assume that these Lagrange multipliers remain constant in the counterfactual scenario.

To summarize, our approach does not assume the separability of inventory management across products, as this would be an unrealistic assumption. Instead, we employ certain assumptions and shortcuts to address the complexities of the joint inventory management problem, while still maintaining a realistic framework that is consistent with the interdependencies among products.

4.2 Sequence of events and profit

For notational simplicity, we omit store and product indexes. Time is indexed by t. One period is one day. Every day, the sequence of events is the following.

Step (i). The day begins with the store manager observing current stock (k_t) , the retail price set by headquarters (p_t) , and her expectation about the mean and variance of the distribution of log-demand: $\ln d_t^e$ and σ_t^2 , respectively. Given this information, the store manager orders y_t units of inventory from the distribution center. There is *time-to-build* in this ordering decision. More specifically, it takes one day for an order to be delivered to the store and become available to consumers.¹³ The ordered amount y_t is a discrete variable with support set $\mathcal{Y} \equiv \{0, 1, ..., J\}$.

Step (ii). Demand d_t is realized. Demand has a Negative Binomial distribution with log-expected demand and variance:

$$\begin{cases} \ln d_t^e = \boldsymbol{\eta}_0' \operatorname{seas}_t + \eta_p \, \ln p_t + \eta_Q \, \ln Q_t^{[-7,-1]} \\ \sigma_t^2 = d_t^e \, (1 + \alpha \, d_t^e) \end{cases} \tag{11}$$

where η_0 , η_p , η_Q , and α are parameters; \mathbf{seas}_t is a vector of seasonal dummies (i.e., weekend dummy and main holidays dummy); $Q_t^{[-7,-1]}$ is the average daily sales of the product in the store during the last seven days; and α denotes the over-dispersion parameter in the Negative Binomial. We use F_{d_t} to represent the distribution of d_t conditional on $(p_t, Q_t^{[-7,-1]}, \mathbf{seas}_t)$. Importantly, the stochastic demand shock $u_t^d \equiv \ln d_t - \ln d_t^e$ is unknown to the store manager at the beginning of the day when she makes her ordering decision.

Step (iii). The store sells q_t units of inventory, which is the minimum of supply and demand:

$$q_t = \min\{ d_t, k_t \} \tag{12}$$

The store generates flow profits Π_t . The profit function has the following form:

$$\Pi_{t} = (p_{t} - c_{t}) \min\{d_{t}, k_{t}\} + \gamma^{z} \mathbb{1}\{d_{t} > k_{t}\} - \gamma^{h} k_{t} - \gamma^{c} y_{t} - \gamma^{f} \mathbb{1}\{y_{t} > 0\} + \sigma_{\varepsilon} \varepsilon_{t}(y_{t})$$
(13)

where c_t is the wholesale price, and γ^z , γ^h , γ^c , γ^f , and σ_{ε} are store-product-specific structural parameters. When $\gamma^z > 0$, the term $\gamma^z \cdot \mathbb{1}\{d_t > k_t\}$ captures the situation where the cost of a stockout can be smaller than the revenue loss from excess demand because some consumers substitute the product within the store. On the other hand, if $\gamma^z < 0$, this term can represent an additional reputational cost of stockouts that goes beyond the lost revenue (see Anderson et al., 2006). The term $\gamma^h \cdot k_t$ represents the storage cost associated with holding k_t units of inventory at the store. Parameter γ^c denotes the per-unit cost incurred by the store manager when placing an order, and γ^f represents the fixed ordering cost, including the transportation cost from the warehouse to the store. The variable $\varepsilon_t(y_t)$ corresponds to a stochastic shock with a mean of zero that affects ordering costs. More specifically, variables $\varepsilon_t(0)$, $\varepsilon_t(1)$, ..., $\varepsilon_t(J)$ are i.i.d. with a Extreme Value type 1 distribution. Parameter σ_{ε} represents the standard deviation of the shocks in ordering costs.

¹³Based on the interviews we conducted with store managers, the most common delivery lag reported is one day. Delivery lags exceeding three days were described as extremely rare.

These γ parameters are the store manager's perceived costs. For instance, the fixed ordering cost γ^f and the per-unit holding cost γ^h can be interpreted as the manager's perception of the shadow prices (or Lagrange multipliers) associated to the capacity constraints of a delivery truck and of the store, respectively.

Price-cost margins. LCBO's retail prices are a constant markup over their respective wholesale prices. There are different markups for Ontario products (65.5% markup) and non-Ontario products (71.5% markup) (see Aguirregabiria et al., 2016). A constant markup, say τ , implies that the price-cost margin is proportional to the retail price: $p_t - c_t = \mathcal{LI} p_t$, where \mathcal{LI} represents the *Lerner index* that by definition is equal to $\frac{\tau}{1+\tau}$. The *Lerner index* is equal to $\frac{0.655}{1+0.655} = 0.40$ for Ontario products and to $\frac{0.715}{1+0.715} = 0.42$ for non-Ontario products.

Step (iv). Orders placed at the beginning of day t, y_t , arrive to the store at the end of the same day or at the beginning of t + 1. Inventory is updated according to the following transition rule:

$$k_{t+1} = k_t + y_t - q_t \tag{14}$$

Finally, next period price p_{t+1} is realized according to a first order Markov process with transition distribution function $F_p(p_{t+1}|p_t)$.

4.3 Dynamic decision problem

A store manager chooses the order quantity y_t to maximize her store's expected and discounted stream of current and future profits. This is a dynamic programming problem with state variables $\mathbf{x}_t \equiv (k_t, p_t, \ln Q_t^{[-7,-1]}, \mathbf{seas}_t)$ and $\boldsymbol{\varepsilon}_t \equiv (\varepsilon_t(0), \varepsilon_t(1), ..., \varepsilon_t(J))$ and value function $V(\mathbf{x}_t, \boldsymbol{\varepsilon}_t)$. This value function is the unique solution of the following Bellman equation:

$$V(\mathbf{x}_t, \boldsymbol{\varepsilon}_t) = \max_{y_t \in \mathcal{Y}} \{ \pi(y_t, \mathbf{x}_t) + \sigma_{\varepsilon} \varepsilon(y_t) + \beta \mathbb{E} [V(\mathbf{x}_{t+1}, \boldsymbol{\varepsilon}_{t+1}) \mid y_t, \mathbf{x}_t] \},$$
(15)

where $\beta \in (0, 1)$ is the store's one-day discount factor; $\pi(y_t, \mathbf{x}_t)$ is the expected profit function up to the ε_t shock; and $\mathbb{E}[.|y_t, \mathbf{x}_t]$ is the expectation over the i.i.d. distribution of ε_{t+1} , and over the distribution of \mathbf{x}_{t+1} conditional on \mathbf{x}_t . The latter distribution consists of the transition probability $F_p(p_{t+1}|p_t)$ and the distribution of demand d_t conditional on \mathbf{x}_t which together with equations (12) and (14) determines the distribution of $(k_{t+1}, p_{t+1}, \ln Q_{t+1}^{[-7,-1]}, \mathbf{seas}_{t+1})$. The solution of this dynamic programming problem implies a time-invariant optimal decision rule: $y_t = y^*(\mathbf{x}_t, \varepsilon_t)$. This optimal decision rule is defined as the *arg max* of the expression within brackets {} in the right-hand-side of equation (15).

For the solution and estimation of this model, we follow Rust (1987, 1994) and use the integrated value function $V_{\sigma}(\mathbf{x}_t) \equiv \frac{1}{\sigma_{\varepsilon}} \int V(\mathbf{x}_t, \boldsymbol{\varepsilon}_t) d\boldsymbol{\varepsilon}_t$ and the corresponding integrated Bellman

equation. Given the Extreme Value distribution of the ε_t variables, the integrated Bellman equation has the following form:

$$V_{\sigma}(\mathbf{x}_{t}) = \ln \left[\sum_{y \in \mathcal{Y}} \exp \left(\frac{\pi(y, \mathbf{x}_{t})}{\sigma_{\varepsilon}} + \beta \mathbb{E} \left[V_{\sigma}(\mathbf{x}_{t+1}) \mid y, \mathbf{x}_{t} \right] \right) \right].$$
(16)

The expected profit function $\pi(y_t, \mathbf{x}_t)$ is linear in the parameters. That is,

$$\frac{\pi(y_t, \mathbf{x}_t)}{\sigma_{\varepsilon}} = \mathbf{h}(y_t, \mathbf{x}_t)' \gamma, \qquad (17)$$

where γ is the vector of structural parameters $\gamma \equiv (1/\sigma_{\varepsilon}, \gamma^{h}/\sigma_{\varepsilon}, \gamma^{z}/\sigma_{\varepsilon}, \gamma^{f}/\sigma_{\varepsilon}, \gamma^{c}/\sigma_{\varepsilon})'$, and $\mathbf{h}(y_{t}, \mathbf{x}_{t})$ is the following vector of functions of the state variables:

$$\mathbf{h}(y_t, \mathbf{x}_t)' = (\mathcal{LI} \ p_t \ \mathbb{E}[\min\{d_t, k_t\} | \mathbf{x}_t], \ -k_t, \ \mathbb{E}[\mathbbm{1}\{d_t > k_t\} | \mathbf{x}_t], \ -\mathbbm{1}\{y_t > 0\}, \ -y_t), \ (18)$$

where the expectation is taken over the distribution of demand conditional on \mathbf{x}_t .

We consider a discrete space for the state variables \mathbf{x}_t .¹⁴ Let $\mathcal{X} \equiv {\{\mathbf{x}^1, \mathbf{x}^2, ..., \mathbf{x}^L\}}$ be the support set of \mathbf{x}_t . We can represent the value function $V_{\sigma}(.)$ as a vector \mathbf{V}_{σ} in the Euclidean space \mathbb{R}^L , and the transition probability functions of \mathbf{x}_t for a given value of y as an $L \times L$ matrix $\mathbf{F}_{\mathbf{x}}(y)$. Taking this into account, as well as the linear-in-parameters structure of the expected profit $\pi(y_t, \mathbf{x}_t)$, the integrated Bellman equation in vector form is:

$$\mathbf{V}_{\sigma} = \ln \left[\sum_{y \in \mathcal{Y}} \exp \left(\mathbf{H}(y) \ \gamma \ + \ \beta \ \mathbf{F}_{\mathbf{x}}(y) \ \mathbf{V}_{\sigma} \right) \right].$$
(19)

where $\mathbf{H}(y)$ is a $L \times 5$ matrix that in row r contains vector $\mathbf{h}(y, \mathbf{x}^r)'$ for $\mathbf{x}^r \in \mathcal{X}$.

The Conditional Choice Probability (CCP) function, $P(y|\mathbf{x}_t)$, is an integrated version of the decision rule $y^*(\mathbf{x}_t, \boldsymbol{\varepsilon}_t)$. For any $y \in \mathcal{Y}$ and $\mathbf{x}_t \in \mathcal{X}$, the CCP $P(y|\mathbf{x}_t)$ is defined as $\int 1\{y^*(\mathbf{x}_t, \boldsymbol{\varepsilon}_t) = y\} dG(\boldsymbol{\varepsilon}_t)$, where G is the CDF of $\boldsymbol{\varepsilon}_t$. For the Extreme Value type 1 distribution, the CCP function has the Logit form:

$$P(y|\mathbf{x}_t) = \frac{\exp\{\mathbf{h}(y,\mathbf{x}_t)' \gamma + \beta \mathbb{E}[V_{\sigma}(\mathbf{x}_{t+1}) \mid y, \mathbf{x}_t]\}}{\sum_{j=0}^J \exp\{\mathbf{h}(j,\mathbf{x}_t)' \gamma + \beta \mathbb{E}[V_{\sigma}(\mathbf{x}_{t+1}) \mid j, \mathbf{x}_t]\}},$$
(20)

Following Aguirregabiria and Mira (2002), we can represent the vector of CCPs, $\mathbf{P} \equiv \{P(y|\mathbf{x}) : (y, \mathbf{x}) \in \mathcal{Y} \times \mathcal{X}\}$, as the solution of a fixed-point mapping in the probability space: $\mathbf{P} = \psi(\mathbf{P})$. Mapping ψ is denoted the *policy iteration* mapping, and it is the composition

¹⁴In the estimation, we discretize the state space using a k-means algorithm.

of two mappings: $\psi(\mathbf{P}) \equiv \lambda(v(\mathbf{P}))$. Mapping $\lambda(\mathbf{V})$ is the *policy improvement*. It takes as given a vector of values \mathbf{V} and obtains the optimal CCPs as "best responses" to these values. Mapping $v(\mathbf{P})$ is the *valuation* mapping. It takes as given a vector of CCPs \mathbf{P} and obtains the corresponding vector of values if the agent behaves according to these CCPs.¹⁵ In our Logit model, the *policy improvement* mapping has the following vector form, for any $y \in \mathcal{Y}$:

$$\mathbf{P}(y) = \lambda(y, \mathbf{V}) = \frac{\exp\{\mathbf{H}(y) \ \gamma + \beta \ \mathbf{F}_{\mathbf{x}}(y) \ \mathbf{V}\}}{\sum_{j=0}^{J} \exp\{\mathbf{H}(j) \ \gamma + \beta \ \mathbf{F}_{\mathbf{x}}(j) \ \mathbf{V}\}}.$$
(21)

The *valuation* mapping has the following form:

$$\mathbf{V} = \upsilon(\mathbf{P}) = \left[\mathbf{I} - \beta \sum_{y=0}^{J} \mathbf{P}(y) * \mathbf{F}_{x}(y)\right]^{-1} \left[\sum_{y=0}^{J} \mathbf{P}(y) * (\mathbf{H}(y) \ \gamma + euler - \ln \mathbf{P}(y))\right], \quad (22)$$

where euler is Euler's constant, and * is the element-by-element vector product.

4.4 Parameter estimates

For every LCBO store and product in our working sample, we estimate the store-product specific parameters in vector γ using a *Two-Step Pseudo Likelihood (2PML)* estimator (Aguir-regabiria and Mira (2002)). Given a dataset $\{y_t, \mathbf{x}_t : t = 1, 2, ..., T\}$ and arbitrary vectors of CCPs and structural parameters (\mathbf{P}, γ), define the pseudo (log) likelihood function:

$$Q(\mathbf{P},\gamma) = \sum_{t=1}^{T} \ln \psi \left(y_t, \mathbf{x}_t ; \mathbf{P}, \gamma \right), \qquad (23)$$

where $\psi(.)$ is the policy iteration mapping defined by the composition of equations (21) and (22). Note that the likelihood function $Q(\mathbf{P}, \gamma)$ is a function of the store's one-day discount factor β . In the estimation, we fix the value¹⁶ of this discount factor equal to $0.95^{\frac{1}{365}}$. In the first step of the 2PML method, we obtain a reduced form estimation of the vector of CCPs **P** using a Kernel method. In the second step, the 2PML estimator is the vector $\hat{\gamma}$ that maximizes the pseudo-likelihood function when $\mathbf{P} = \hat{\mathbf{P}}$. That is:

$$\widehat{\gamma} = \arg \max_{\gamma} Q(\widehat{\mathbf{P}}, \gamma) \tag{24}$$

 $^{^{15}\}mathrm{See}$ Puterman (2014) for a description of these three mappings in the context of a general dynamic programming problem.

¹⁶We could eventually relax this assumption and treat β as a parameter to be estimated, and allow it to vary across store managers in order to potentially capture different degrees of myopia or impatience.

Aguirregabiria and Mira (2002) show that this estimator is consistent and asymptotically normal with the same asymptotic variance as the full maximum likelihood estimator. In section A.6 in the Appendix, we provide further details on the implementation of this estimator.

In a similar vein to the estimation of (S, s) thresholds in section 3, a portion of the variation in parameter estimates $\hat{\gamma}_{i,j}$ can be attributed to estimation error rather than genuine heterogeneity. To address this issue and mitigate the excessive dispersion or spurious heterogeneity resulting from estimation error, we employ a *shrinkage estimator*. The details of this estimator can be found in section A.8 in the Appendix.

Table 6 presents the medians from the empirical distributions (across stores and products) of our estimates of the four structural parameters, measured in dollar amounts.¹⁷ The median values of the estimates are \$0.0036 for the per-unit inventory holding cost, \$0.0219 for the stockout cost, \$2.9658 for the fixed ordering cost, and \$0.0341 for the per-unit ordering cost. To have an idea of the importance of these dollar amounts, in Section 4.5 below we provide measures of the implied magnitude of each cost relative to revenue. These magnitudes are consistent with other cost estimates in the inventory management literature (see Aguirregabiria [1999], Bray et al. [2019]). Median standard errors and t-statistics in Table 6 show that the inventory holding cost and the fixed ordering cost are very precisely estimated (median t-ratios of 5.32 and 12.65, respectively), while a substantial fraction of the estimates of the stockout cost are quite imprecise (median t-ratio of 0.27).

	Median	Std. Dev.	Median	Median
	Estimate	Estimate	S.e.	t-stat.
γ^h : Per-unit Inventory Holding Cost	0.0036	0.0029	0.0007	5.3259
γ^z : Stockout Cost	0.0219	0.3094	0.1465	0.2672
γ^f : Fixed Ordering Cost	2.9658	1.0847	0.2356	12.6557
γ^c : Per-Unit Ordering Cost	0.0341	0.0607	0.0274	1.4763
# of observed store-product pairs	3,076			
# of store-product pairs with structural estimates	2,589			

Table 6: Structural Estimates of Cost Parameters (in Canadian Dollars)

¹⁷More specifically, we first obtain the two-step PML estimate of the vector $\gamma \equiv (1/\sigma_{\varepsilon}, \gamma^{h}/\sigma_{\varepsilon}, \gamma^{z}/\sigma_{\varepsilon}, \gamma^{f}/\sigma_{\varepsilon}, \gamma^{c}/\sigma_{\varepsilon})'$, and then we divide elements 2 to 5 of this vector by the first element to obtain estimates of costs in dollar amount. We use the delta method to obtain standard errors.



Figure 4: γ Estimates: Shrinkage Estimator

4.5 Relative contribution of the different costs

In this section we assess the magnitude of the different inventory management costs relative to store revenues. The purpose of this exercise is twofold. First, we want to evaluate whether our parameter estimates imply realistic magnitudes for the realization of these costs. And second, it is relevant to measure to what extent the heterogeneity in cost parameters that we have presented above generates heterogeneity in profits across stores. Conditional on their perception of cost parameters, store managers' optimal behavior should compensate – at least partly – for the differences in cost parameters such that heterogeneity in realized costs should be smaller. We want to measure the extent of this compensating effect.

For each component of the inventory management cost, and for every store-product, we calculate the ratio between the realized value of the cost during our sample period and the realized value of revenue during the same period. More specifically, we calculate the following ratios for every store-product: inventory holding cost to revenue; stockout cost to revenue; fixed ordering cost to revenue; variable ordering cost to revenue; and total inventory management cost to revenue. We have an empirical distribution over store-products for each

of these ratios.

Table 7 presents the median and the standard deviation in these distributions. To evaluate the magnitude of these ratios, it is useful taking into account that – according to the LCBO's annual reports – the total expenses to sales ratio of the retail chain is consistently around 16% each year.¹⁸ According to our estimate, the total inventory cost-to-revenue ratio for the median store is approximately 1.37%. This would imply that the retail chain's cost of managing the inventories of their stores would represent around 10% of total costs, which entails that non-inventory related costs would account for approximately 90% of total costs (e.g. labor costs, fixed capital costs, delivery costs). This seems to be of the right order of magnitude. Table 7 shows that the fixed ordering cost is the largest realized cost for store managers at LCBO, followed by storage costs. Realized stockout costs are negligible. This is due to a combination of a small parameter that captures the stockout cost, and infrequent stockouts in our working sample.

	Median	St. Dev.
Inventory Holding Cost to Revenue Ratio (%) Stockout Cost to Revenue Ratio (%) Fixed Ordering Cost to Revenue Ratio (%) Variable Ordering Cost to Revenue Ratio (%)	$\begin{array}{c} 0.2863 \\ 0.0005 \\ 0.8731 \\ 0.2045 \end{array}$	$\begin{array}{c} 0.2182 \\ 0.0197 \\ 0.7728 \\ 0.1895 \end{array}$
Total Inventory Cost to Revenue Ratio (%)	1.3749	0.9235

Table 7: Realized Inventory Management Costs to Revenue Ratios

In section A.9 in the Appendix, we present the empirical distribution across stores and products of each of the four cost-to-revenue ratios. We also show the extent to which managers' inventory decisions compensate for the heterogeneity in the structural parameters.

4.6 Heterogeneity in cost parameters

Below, we investigate two potential sources for the large heterogeneity in our cost parameters: (i) differences across stores, such as store type according to LCBO's classification of stores, physical area, total product assortment, distance to the warehouse, and consumer socioeconomic characteristics; and (ii) differences across local managers. Our goal in this section

¹⁸Of course, these expenses do not include the cost of merchandise.

is to separate the heterogeneity attributable to store characteristics, and the heterogeneity stemming from the managers themselves. We proceed using a sequential approach. First, we regress our parameter estimates on a set of store characteristics. Then, we take the residual components from the first step and regress them on manager characteristics.

	(γ^h)	(γ^z)	(γ^f)	(γ^c)
	Est	Est	Est	Est
	(e e)	(e e)	(e e)	(e e)
	(3.6.)	(3.6.)	(3.6.)	(3.6.)
Store Class				
AA	0.000554	-0.0216	-0.164	0.00302
	(0.000411)	(0.0194)	(0.185)	(0.00587)
Δ	0.000511	-0.0204	-0.203	0.00816
21	(0.000311)	(0.0181)	(0.178)	(0.00510)
	(0.000391)	(0.0101)	(0.178)	(0.00512)
В	-0.000415	-0.0237	0.108	0.00761
	(0.000434)	(0.0204)	(0.190)	(0.00590)
C	-0 00209***	-0.0531**	0 785***	0.00477
C	(0.00200)	(0.0243)	(0.221)	(0.00657)
	(0.000020)	(0.0240)	(0.221)	(0.00001)
D	-0.00334***	-0.0379	1.223^{***}	0.00479
	(0.000564)	(0.0262)	(0.244)	(0.00693)
ln(Product Assortment Size)	0.000336	-0.00932	-0.364***	0.00388^{*}
((0.000205)	(0.00973)	(0.0775)	(0.00201)
	(0.000200)	(0.000.0)	(0.0110)	(0.00201)
In(Population in City)	-0.0000801**	-0.00136	0.00948	0.000413
	(0.0000381)	(0.00220)	(0.0186)	(0.000483)
ln(Median Income in City)	0.000362	0.0493^{*}	-0.0577	0.00795
	(0.000520)	(0.0255)	(0.212)	(0.00646)
	VEC	VEO	VEC	
Location dummies (25 regions, 4 districts)	YES	YES	YES	YES
Product dummies (5 products)	YES	YES	YES	YES
R-squared	0.3954	0.1345	0.5265	0.0598
Observations	2,589	2,589	2,589	2,589

Table 8: Regressions on Store Characteristics

(1) Location dummies based on LCBO's own division of Ontario into 25 regional markets and 4 districts.

(2) Robust standard errors clustered at the store level in parentheses

(3) * means p-value<0.10, ** means p-value<0.05, *** means p-value<0.01

First step: store characteristics. Table 8 presents estimation results from the firststep regressions of each estimated cost parameter against store and location characteristics: LCBO's store type dummies (6 types); LCBO's regional market dummies (25 regions); logarithm of the number of unique products offered by the store; logarithm of population in the store's city; and logarithm of median income level in the store's city. As we have cost estimates at the store-product level, we also include product fixed effects.¹⁹

These store and location characteristics can explain an important part of the variation across stores in inventory holding costs and fixed ordering costs: the R-squared coefficients for these regressions are 0.39, and 0.53, respectively. Fixed ordering costs decline significantly with the number of products in the store, which is consistent with economies of scope in ordering multiple products. Inventory holding costs increase with assortment size and are significantly higher for AAA stores relative to D stores. In contrast, only 6% of the variation in unit ordering costs and 13% of the variation in stockout costs can be explained by these store and location characteristics. These results are robust to other specifications of the regression equation based on transformations of explanatory or/and dependent variables.

Second step: manager characteristics. Table 9 presents the estimation results from the second-step regressions of cost parameters on manager characteristics (educational attainment, years of experience at the LCBO, and other industry experience), after controlling for the variation explained by store characteristics. The overall finding is that managers' education and experience have non-significant effects in these regressions. There are two main reasons that can explain these negligible effects.

First, there is a substantial correlation between store characteristics and managers' skills. More skilled managers tend to be allocated to higher-class stores (positive assortative matching). Therefore, in the first-step regression, where store characteristics are included, these characteristics are also capturing the effect of managers' skills. As a result, the direct effect of managers' skills in the second-step regressions becomes less apparent.

Second, the insignificant effect of managers' skills on the estimated cost parameters aligns with the interpretation that the residual component of these parameters is associated with biased perceptions. More skilled managers may have a better measure of these costs, while less skilled managers may have noisier estimates. However, this does not imply a larger or smaller effect of managers' skills on the mean value of cost parameters. Instead, the effect would appear in the variance of the cost parameters, indicating differences in the precision of their estimates. Indeed, when we regress the variance of the cost parameters on managers' skills, we find evidence supporting this interpretation.

¹⁹Note that the store location dummies capture various factors, including the effect of the distance between the store and the warehouse.

	$(res(\gamma^h))$	$(res(\gamma^z))$	$(res(\gamma^f))$	$(res(\gamma^c))$
	Est.	Est.	Est.	Est.
	(s.e.)	(s.e.)	(s.e.)	(s.e.)
Educational Attainment				
High School	-0.0000165	0.000752	-0.0000456	0.000420
	(0.000237)	(0.0122)	(0.0769)	(0.00271)
University	-0.0000801	0.00474	0.0194	-0.000213
U	(0.000170)	(0.00794)	(0.0577)	(0.00183)
LCBO Experience	0.0000128**	-0.000204	0.00228	-0.0000967
-	(0.00000627)	(0.000270)	(0.00223)	(0.0000669)
Other Experience	0.0000156	0.000576	0.00471	-0.0000587
-	(0.0000125)	(0.000475)	(0.00437)	(0.000115)
R-squared	0.0095	0.0051	0.0058	0.0038
Observations	2,589	2,589	2,589	2,589

Table 9: Regressions on Manager Characteristics

(1) Robust standard errors clustered at the store level in parentheses

(2) * means p-value<0.10, ** means p-value<0.05, *** means p-value<0.01

(3) We use multiple imputation to account for missing values of our explanatory variables

In Section A.10 of the Appendix, we also examine how the variance of the cost parameters depends on store characteristics. We find that the dispersion of the (second-step) manager component of costs is larger on average for lower-class stores. Since managers in these stores generally have lower levels of human capital (i.e. education and experience), we interpret the second-step manager component as a biased perception of the true cost from the point of view of store managers. That is, the (first-step) store component of the costs will be interpreted as the true cost, and the manager component will be interpreted as deviations from this true cost. In order to illustrate the interpretation of the residual component as manager bias, we present in Section A.10 of the Appendix two granular examples in which pairs of stores – located in close proximity to each other – are similar in size, sales, store classification, but have very different levels of manager experience and estimates of the cost parameters.

We explore the interpretations of our cost parameters, and their impact on store-level inventory outcomes, in the subsequent counterfactual experiments of Section 5.

5 Counterfactual experiments

This section presents two sets of counterfactual experiments based on the model that we have estimated in the previous section. First, we study the contribution to inventory management outcomes from the heterogeneity in store managers' perceptions of costs. Second, we evaluate the effects of a counterfactual centralization of inventory management decisions at LCBO headquarters. We present this counterfactual experiment under different scenarios on the information that headquarters has about demand and costs at the store level.

5.1 Removing store managers' idiosyncratic effects

Let $\hat{\gamma}_{i,j}$ be the vector of estimates of cost parameters for product j and store i. Based on the regressions in Tables 8 and 9, we decompose this vector into two additive and orthogonal components: the part explained by store and location characteristics, that we represent as $\hat{\gamma}_{i,j}^{sto}$; and the part explained by local managers, $\hat{\gamma}_{i,j}^{man}$. Below, we construct a counterfactual scenario that removes the idiosyncratic component $\hat{\gamma}_{i,j}^{man}$ from the inventory decision problem for store i and product j. For every store-product (i, j) in our working sample, we implement a separate counterfactual experiment for each of the four cost parameters, and one experiment that shuts down together the manager component of the four cost parameters. This implies a total number of 15, 850 experiments.

We implement each of these experiments by solving the dynamic programming problem and obtaining the corresponding CCPs under the counterfactual values of the structural parameters. We use this vector of CCPs to calculate the corresponding ergodic distribution of the state variables for the store-product.²⁰ Finally, we use the vector of CCPs and the ergodic distribution to calculate mean values of relevant outcome variables related to inventory management. We compare these average outcomes with their corresponding values under the factual values of structural parameters. In terms of outcome variables, we look at the same descriptive statistics as those reported in Table 1 and Figure 1: stockout frequency, ordering frequency, inventory to sales ratio, inventory to sales ratio after an order (i.e. *S* threshold), and inventory to sales ratio before an order (i.e. *s* threshold).

Figures 5 (for stockout frequency), 6 (for ordering frequency), and 7 (for inventory-tosales ratio) summarize the results from these experiments. In each figure, the horizontal axis measures the value of the corresponding parameter $\gamma_{i,j}^{man}$, and the vertical axis measures the difference in the mean value of the outcome variable between the factual and the counterfactual scenario. For instance, in Figure 5(a), the horizontal axis represents $\gamma_{i,j}^{h,man}$,

 $^{^{20}}$ Note that this ergodic distribution incorporates the seasonal effects in the demand part of the model, as seasonal dummies are a component of the vector of state variables of the model.

and the vertical axis measures $\Delta SOF_{i,j} = SOF_{i,j}^{factual} - SOF_{i,j}^{counter}$, where SOF is stockout frequency.

Note that the counterfactual experiment of shutting down $\gamma_{i,j}^{h,man}$ to zero is equivalent to a change in parameter $\gamma_{i,j}^{h}$ from the counterfactual value $\gamma_{i,j}^{h,sto}$ to the factual value $\gamma_{i,j}^{h,sto} + \gamma_{i,j}^{h,man}$. Therefore, we can see the cloud of points in, say, Figure 5(a) as the results of many comparative statics exercises, all of them consisting in changes in the value of parameter γ^{h} . In these figures, there are multiple curves relating a change in γ^{h} with a change in the outcome variable because store-products have different values of the other structural parameters. However, each of these figures shows a monotonic relationship between a change in a cost parameter and the corresponding change in an outcome variable.



Figure 5: Counterfactual Outcome: Stockout Frequency

 Δ Stockout Frequency on y-axis, γ^{man} on x-axis

We compare the relationship between parameters and outcomes implied by these figures with the theoretical predictions from the model as depicted in equation (6). More specifically, note that S - s is negatively related to the ordering frequency (the larger the S - s, the smaller the ordering frequency); s is negatively related to the stockout rate (the larger the s, the smaller the stockout rate); and the levels of both S and s are positively related to the inventory to sales ratio (the larger the S and s, the larger the ratio). The pattern in our figures is fully consistent with Blinder's theoretical predictions for this class of models.

Figure 5 depicts the relationship between cost parameters and the stockout frequency. According to Blinder's formula, the lower threshold s depends negatively on γ^h and γ^f , and positively on γ^z , while the effect of γ^c is ambiguous. Panels (a) to (d) in Figure 5 confirm the signs of these effects on the stockout frequency.



Figure 6: Counterfactual Outcome: Ordering Frequency

 Δ Ordering Frequency on y-axis, γ^{man} on x-axis

In Figure 6, we present the relationship between cost parameters and the ordering frequency. Blinder's formula says that S - s depends negatively on γ^h and positively on γ^f . Panels (a) and (c) confirm the sign of these effects on the ordering frequency. According to Blinder, the sign of the effects of γ^z and γ^c on ordering frequency is ambiguous because they affect the two thresholds S and s in the same direction. In Panel (b), we find a positive relationship between the stockout cost γ^z and ordering frequency. Panel (d) shows that the frequency of placing an order falls when the unit ordering cost increases.





 Δ Inventory to Sales Ratio on y-axis, γ^{man} on x-axis

Figure 7 illustrates the relationship between cost parameters and the inventory-to-sales ratio. Blinder's formula establishes that the two thresholds S and s depend negatively on γ^h and positively γ^z . Panels (a) and (b) confirm these signs for the inventory-to-sales ratio. However, according to Blinder's formula, the sign of the effects of γ^f and γ^c on the inventoryto-sales ratio is ambiguous. Panels (c) and (d) show negative effects of γ^f and γ^c on the inventory-to-sales ratio.

			Parameters	Shut Down	,	
	None	γ^h	γ^z	γ^f	γ^c	All
Store-product level inventory outcomes	Mean	Mean	Mean	Mean	Mean	Mean
	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)	(st.dev)
Stockout Frequency	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003
	(0.0004)	(0.0003)	(0.0004)	(0.0004)	(0.0004)	(0.0003)
Ordering Frequency	0.1776	0.1770	0.1776	0.1714	0.1729	0.1618
	(0.1329)	(0.1353)	(0.1329)	(0.1260)	(0.1313)	(0.1240)
Inventory to Sales Ratio	26.8373	23.3465	26.8376	27.8804	25.7753	22.1358
	(21.1587)	(13.3625)	(21.1589)	(25.1945)	(19.1622)	(11.8119)
Inventory to Sales Ratio After Order	31.6917	26.7060	31.6920	33.7313	29.6416	24.2525
	(29.7575)	(17.9471)	(29.7576)	(36.5440)	(25.5906)	(13.9094)
Inventory to Sales Ratio Before Order	22.6063	18.3920	22.6066	24.6034	20.9836	16.6467
	(22.0222)	(12.2635)	(22.0224)	(28.0980)	(18.7917)	(8.7038)
Change in Total Inventory Cost (%)	-	-4.1913	0.0005	1.5256	-4.5866	-12.0838
		(11.7293)	(0.0048)	(24.1220)	(11.4978)	(17.5204)

Table 10: Removing the Manager Component of Cost Parameters

It is of interest to measure the average effect across stores and products of shutting down the store manager-specific component in costs. Table 10 presents these average effects for each of the four cost parameters and for the combination of the four.²¹ Removing the manager component in all four inventory costs generates a decrease in the mean ordering frequency of 1.6 percentage points, from 17.8% to 16.2%; a decrease in the inventory-to-sales ratio of 4.7 days of average sales, from 26.8 to 22.1 days; a decrease in the lower *s* threshold of 6 days, from 22.6 to 16.6 days; and a decrease in the S - s gap of 1.5 days, from 9.1 to 7.6 days.

Store managers' idiosyncratic perception of costs has a substantial effect on inventory management at the aggregate firm level. It entails a 6-day decrease in waiting time between two orders, an increase in the average order amount of 1.5 days of average sales, and a 21% increase in the inventory-to-sales ratio, but a negligible effect on the frequency of stockouts.

²¹Note that, by construction, the manager component $\gamma_{i,j}^{man}$ has mean zero and is orthogonal to the store component $\gamma_{i,j}^{sto}$. Therefore, if the model implied a linear relationship between outcome variables and structural parameters, then the average effect of shutting down the residual component would be zero. For the same reason, a first-order linear approximation to this average effect is zero. However, the model implies a nonlinear relationship between outcomes and structural parameters such that it is a relevant empirical question to look at these average effects. In fact, we find that the effect is not negligible at all.

Accordingly, if this idiosyncratic component is a biased perception, then it has a substantial negative impact on the firm's profit as it increases storage and ordering costs with almost no effect on stockouts and revenue. The bottom row in Table 10 presents the effect of removing $\gamma_{i,j}^{man}$ on total inventory management cost calculated using $\gamma_{i,j}^{sto}$ but not $\gamma_{i,j}^{man}$. We find that, on average, this cost declines by 12.1%. This substantial effect plays an important role in the counterfactual experiment on centralization that we present in the next section.

5.2 Centralizing inventory decision-making

We now address the main question that motivates this paper: would the LCBO retail chain benefit from managing the stores' inventories at the headquarter level, as opposed to allowing heterogeneous store managers to have autonomy in their inventory decisions? To answer this question, we need to establish some conditions on the headquarters' information about storelevel demand, inventories, and cost parameters. The experiments that we present below are based on the following conditions.

First, based on the institutional details we describe in Section 2.1, we consider that headquarters process store-product level transactions data with a one-week delay. Transmission of information from stores to headquarters occurs in real time, without any substantial delay. However, it takes time to process that information to generate headquarters demand predictions and ordering recommendations. Though a fully automated inventory management system is possible, human supervision can add value by accounting for soft information (Cimini et al., 2019). Accordingly, in the counterfactual centralized system, we replace state variable $Q_t^{[-7,-1]}$ with the one-week lag of this variable, i.e., $Q_t^{[-14,-8]}$.

Second, to compare profits between the centralized and decentralized structures, we must take a stance on what are the "true" cost parameters. We assume that the true cost parameters are $\gamma_{i,j}^{sto}$ which are determined by store and location characteristics. Under the centralized system, the headquarters know these costs and take inventory decisions for every store based on these costs. In contrast, we interpret $\gamma_{i,j}^{man}$ as store managers' behavioral biases and not as "true" costs. Under the decentralized system, store managers make decisions as if the cost parameters were $\gamma_{i,j}^{sto} + \gamma_{i,j}^{man}$, but our measure of their profits is based only on $\gamma_{i,j}^{sto}$. The evaluation of profits under this assumption provides an upper bound for the (profit) gains from centralization. Alternatively, we could assume that $\gamma_{i,j}^{man}$ is a true component of profit that is known to the store manager but unknown to the headquarters. This alternative assumption would provide a lower bound for the gains from centralization.

Based on these assumptions, this counterfactual experiment measures the following tradeoff in the choice between centralized and decentralized inventory management. A negative aspect of decentralization is that store managers have different skills and behavioral biases as captured by the idiosyncratic components $\gamma_{i,j}^{man}$. These biases should have a negative effect on LCBO profits. The positive aspect of decentralization is that store managers have just-in-time information about demand, sales, and inventories, while the firm's headquarters process this information with one week delay. This just-in-time information should have a positive effect on LCBO profits.

Table 11: Decentralized vs. Centralized Profits: Average Daily Profit Per-Store Per-Product

	Mean	Pct. 10%	Pct. 25%	Median	Pct. 75%	Pct. 90%
Centralized Solution (\$)	53.72	10.54	20.29	44.06	80.34	109.21
$Decentralized \ Solution \ (\$)$	52.81	10.16	19.64	43.34	79.66	108.04
$Gains\ in\ Profit\ from\ Decentralization\ (\$)$	-0.91	-3.10	-1.58	-0.66	-0.05	0.80
Gains in Profit from Decentralization $(\%)$	-1.97	-6.09	-3.81	-2.09	-0.19	1.85
Change in Inventory Cost from Decentralization $(\%)$	22.99	-8.06	-2.59	3.72	21.99	62.66

1% change in profit per store-product is approximately \$17 million in total annual profit for LCBO.

Figure 8: Change in Daily Profit From Decentralization (\$)



We depict the results of this experiment in Table 11, and Figures 8 and 9. Similarly as for the counterfactuals in section 5.1, we evaluate the effects using the ergodic distributions of the state variables under the factual and counterfactual scenarios. Table 11 presents means, medians, and several percentiles for the profit per store-product under the centralized and

Figure 9: Change in Daily Inventory Cost From Decentralization (%)



decentralized systems and for the gains from decentralization. To have a better perspective of the implications of these effects, it is useful to take into account that a 1% change in profit per store-product represents approximately \$17 million in total annual profit for LCBO in year 2012.²² At the aggregate level, decentralization has a negative impact on LCBO profits. It implies a 2% decline in profits, which represents \$34 million in annual profits for the retail chain. The effect on the median store is also negative: -2.1%. This relatively modest effect is the result of combining two large effects with opposite signs. The one-week delay in the processing of information in the centralized system has a non-negligible negative impact on profits at every store. However, this negative effect of the centralized system is more than compensated by the large increase in profits due to reducing ordering and storage costs when removing store managers' biased perceptions of costs. This is illustrated in the bottom row of Table 11: on average, decentralization increases total inventory costs by 23%.

The evidence on the mean and median effects in Table 11 is not necessarily sufficient for a retail chain to adopt a centralized inventory management system. A retail chain may need to assess the distributional effects of the gains/losses across its stores before adopting a substantial organizational change, and not simply rely on the average effect. Table 11 and Figure 8 show significant heterogeneity in the impact of decentralization, with a considerable amount of stores benefiting from the decentralized structure: the 90th percentile store has a gain in profit of 1.8%, which is not negligible. Therefore, although centralization would generate positive gains in total profits for the retail chain relative to the existing decentralized

 $^{^{22} \}rm According$ to the 2012-2013 LCBO annual report, the annual profit (net income) of the company was \$1.7 billion. Therefore, 1% of this profit is \$17 million.

structure, the distributional effects of these gains are important to assess.

Figure 9 provides a closer look at the heterogeneous effect from (de)centralization. It presents the empirical distribution of the percentage change in average daily inventory costs. The median of this distribution presents a positive increase in costs (i.e., 3.7%), but most striking is the long right tail of this distribution, which implies a 23% increase in average inventory costs.

6 Conclusion

Retail chains are complex organizations with various divisions and teams, each having its own decision-making authority. Store managers, in particular, play a crucial role within certain retail chains. These managers have the advantage of collecting and processing timely information specific to their individual stores. This store-level information is more consistent and manageable compared to information at the chain-wide level. However, the transfer and processing of this information from stores to headquarters can introduce delays ranging from days to weeks, which can negatively impact decision-making and overall profitability.

On the other hand, store managers exhibit heterogeneous skills, motivations, and levels of effort. A centralized decision-making system that selectively utilizes the most competent managers within the organization can help mitigate the negative effects stemming from the variation in managers' skills.

In this paper, we examine the trade-off between centralized and decentralized decisionmaking in the context of inventory management within a large retail chain. Leveraging a unique dataset containing daily information on inventories, sales, prices, and stockouts at the individual store and product level, we estimate a dynamic structural model to capture store managers' inventory decisions. Using revealed preference as a guiding principle, we obtain separate estimates for each store and product regarding four cost parameters: per unit inventory holding cost, stockout cost, fixed ordering costs, and per unit ordering costs. Our analysis reveals significant heterogeneity across stores in these cost parameters. While observable store and location characteristics account for part of this heterogeneity, a substantial portion can be attributed to idiosyncratic information and perceptions of store managers.

We utilize the estimated model to conduct a counterfactual experiment, evaluating the impact of centralizing inventory management at LCBO. In this experiment, we assume that the idiosyncratic cost component associated with store managers represents a behavioral bias rather than true costs. This assumption allows us to provide an upper-bound estimate for the gains from centralization. Our findings indicate that a centralized inventory management system would lead to a modest 2% increase in LCBO's annual profit. This outcome arises

from the combination of two opposing effects. The negative effect on profits due to the loss of just-in-time information from store managers is outweighed by the significant reduction in ordering and storage costs resulting from the elimination of behavioral biases and skill heterogeneity among store managers (averaging at 23% reduction overall and 3.7% reduction for the median store).

Furthermore, the effects of centralization are highly heterogeneous across stores within the retail chain, with a substantial number of stores experiencing significant losses from adopting centralization. This distributional effect has important implications for the decision-making process when considering organizational changes aimed at maximizing overall company profit (Inderst et al., 2007).

Our empirical findings highlight the advantages of a hybrid inventory management system that combines decentralized decision-making with centralized control. By assigning decision rights to high-skilled store managers and utilizing a centralized system for stores where skill levels are lower, we can eliminate subjective biases while retaining the benefits of just-intime local information for some of the stores. The structural model presented in this paper provides a useful tool for determining the specific allocation of decision rights across stores.

References

- J. Adda and R. Cooper. Balladurette and juppette: A discrete analysis of scrapping subsidies. Journal of Political Economy, 108(4):778–806, 2000.
- P. Aghion, N. Bloom, B. Lucking, R. Sadun, and J. Van-Reenen. Turbulence, firm decentralization, and growth in bad times. *American Economic Journal: Applied Economics*, 13(1):133–69, 2021.
- V. Aguirregabiria. The dynamics of markups and inventories in retailing firms. *The Review* of *Economic Studies*, 66(2):275–308, 1999.
- V. Aguirregabiria and A. Magesan. Identification and estimation of dynamic games when players' beliefs are not in equilibrium. *The Review of Economic Studies*, 87(2):582–625, 2020.
- V. Aguirregabiria and P. Mira. Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica*, 70(4):1519–1543, 2002.
- V. Aguirregabiria, D. Ershov, and J. Suzuki. Estimating the effects of deregulation in the ontario wine retail market. *Working Paper*, 2016.
- E. Anderson, G. Fitzsimons, and D. Simester. Measuring and mitigating the costs of stockouts. *Management Science*, 52(11):1751–1763, 2006.
- J. Arrow, T. Harris, and J. Marschak. Optimal inventory policy. *Econometrica*, 19(3): 250–272, 1951.
- D. Arthur and S. Vassilvitskii. k-means++: The advantages of careful seeding. Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA), page 1027–1035, 2007.
- O. Attanasio. Consumer durables and inertial behaviour: Estimation and aggregation of (s, s) rules for automobile purchases. *The Review of Economic Studies*, 67(4):667–696, 2000.
- G. Bertola, L. Guiso, and L. Pistaferri. Uncertainty and consumer durables adjustment. *Review of Economic Studies*, 72(4):973–1007, 2005.
- A. Blinder. Retail inventory behavior and business fluctuations. Brooking Papers of Economic Activity, 2(1):443–505, 1981.
- R. Bray, Y. Yao, Y. Duan, and J. Huo. Ration gaming and the bullwhip effect. Operations Research, 67(2):453–467, 2019.

- S. Cho and J. Rust. The flat rental puzzle. *The Review of Economic Studies*, 77(2):560–594, 2010.
- C. Cimini, A. Lagorio, F. Pirola, and R. Pinto. Exploring human factors in logistics 4.0: Empirical evidence from a case study. *IFAC-PapersOnLine*, 52(13):2183–2188, 2019.
- S. DellaVigna and M. Gentzkow. Uniform pricing in us retail chains. *NBER Working Paper 23996*, 2019.
- E. Denardo. *Dynamic Programming: Models and Applications*. Prentice Hall PTR, USA, 1981. ISBN 0132215071.
- U. Doraszelski, G. Lewis, and A. Pakes. Just starting out: Learning and equilibrium in a new market. *American Economic Review*, 108(3):565–615, 2018.
- J. Eberly. Adjustment of consumers' durables stocks: Evidence from automobile purchases. Journal of Political Economy, 102(3):403–436, 1994.
- S. Ellison, C. Snyder, and H. Zhang. Costs of managerial attention and activity as a source of sticky prices: Structural estimates from an online market. *NBER Working Paper 24680*, 2018.
- A. Goldfarb and M. Xiao. Who thinks about the competition? managerial ability and strategic entry in us local telephone markets. *American Economic Review*, 101(7):3130– 61, 2011.
- J. Gu and R. Koenker. Empirical bayesball remixed: Empirical bayes methods for longitudinal data. *Journal of Applied Econometrics*, 32(3):575–599, 2017.
- G. Hadley and T. Whitin. Analysis of inventory systems. Prentice-Hall, 1963.
- G. Hall and J. Rust. An empirical model of inventory investment by durable commodity intermediaries. *Carnegie-Rochester Conference Series on Public Policy*, 52:171–214, 2000.
- G. Hitsch, A. Hortaçsu, and X. Lin. Prices and promotions in us retail markets. *Quantitative Marketing and Economics*, pages 1–80, 2021.
- A. Hortaçsu and S. Puller. Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market. *The RAND Journal of Economics*, 39(1):86–114, 2008.
- A. Hortaçsu, F. Luco, S. Puller, and D. Zhu. Does strategic ability affect efficiency? evidence from electricity markets. *American Economic Review*, 109(12):4302–42, 2019.

- A. Hortaçsu, O. Natan, H. Parsley, T. Schwieg, and K. Williams. Organizational structure and pricing: Evidence from a large us airline. *National Bureau of Economic Research*, *Working Paper*, 29508, 2021.
- Y. Huang, P. B. Ellickson, and M. J. Lovett. Learning to set prices. Journal of Marketing Research, 59(2):411–434, 2022.
- R. Inderst, H. M. Müller, and K. Wärneryd. Distributional conflict in organizations. *European Economic Review*, 51(2):385–402, 2007.
- O. Kryvtsov and V. Midrigan. Inventories, markups, and real rigidities in menu cost models. *Review of Economic Studies*, 80(1):249–276, 2013.
- Liquor Control Board of Ontario. Annual report 2011-2012. 2012. https://www.lcbo.com/ content/dam/lcbo/corporate-pages/about/pdf/2011_2012.pdf.
- Liquor Control Board of Ontario. Annual report 2012-2013. 2013. https://www.lcbo.com/ content/dam/lcbo/corporate-pages/about/pdf/LCB0_AR12-13-english.pdf.
- Liquor Control Board of Ontario. An introduction to web based ordering & sos (store order services). 2016. LCBO PSM Trade Symposium.
- E. Miravete, K. Seim, and J. Thurk. One markup to rule them all: Taxation by liquor pricing regulation. *American Economic Journal: Microeconomics*, 12(1):1–41, 2020.
- M. Puterman. Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons, 2014.
- J. Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.
- J. Rust. Structural estimation of markov decision processes. *Handbook of Econometrics*, 4: 3081–3143, 1994.
- H. Scarf. The optimality of (s,s) policies in the dynamic inventory problem. In K. Arrow,S. Karlin, and P. Suppes, editors, *Mathematical Methods in the Social Sciences*, chapter 13.Stanford University Press, Stanford, 1959.

A Appendix

A.1 Correlations between inventory outcomes

Figure 10 below, presents five scatter plots at the store level: (a) stockout frequency against ordering frequency; (b) stockout frequency against inventory-to-sales ratio; (c) stockout frequency against inventory-to-sales ratio after an order is received; (d) stockout frequency against inventory-to- sales ratio before an order is placed; and (e) inventory-to-sales ratio after an order is received against inventory-to-sales ratio before an order is placed. The simple correlations in these figures provide preliminary descriptive evidence on the possible sources of structural heterogeneity, such as heterogeneity across stores in storage cost, stockout cost, ordering cost, or demand uncertainty, which are structural parameters in our model in Section 4.

The strongest correlation appears in Panel (e), for the relationship between our measures of the thresholds S and s.²³ This positive correlation can be explained by store heterogeneity in stockout costs and/or storage costs: a higher stockout cost (storage cost) implies higher (lower) values of both s and S. In contrast, a higher lump-sum ordering cost implies a lower sbut a negligible effect on S. Therefore, the positive correlation between the lower and upper thresholds we observe seems more compatible with store heterogeneity in stockout and/or storage costs rather than with heterogeneity in ordering costs. We confirm this conjecture in the estimation of the structural model in Section 4.

Panel (a) shows a negative relationship between the stockout frequency and the ordering frequency. As one would expect, stores placing orders more frequently tend to have lower stockout rates. Panels (b) and (c) show a small negative relationship between stockout rates and the inventory to sales ratio overall and after an order is received, respectively. These findings are what we would expect: stores that have lower inventory-on-hand on average experience higher stockout rates, and stores that order up to a smaller threshold S also experience higher stockout rates. Relatedly, panel (d) shows a small negative relationship between stockout rates and our measure of the threshold s. Again, this is what we would expect, as stores with a lower safety stock level are more likely to experience higher stockout rates.

²³For the interpretation of this empirical evidence, it is useful to take into the comparative statics of the (S, s) thresholds as functions of the structural parameters in the profit function. We present these comparative statics in equation (6) in Section 3.1, based on results in Hadley and Whitin (1963) and Blinder (1981).









A.2 Correlations between manager and store characteristics

Figure 11 presents correlations between education and experience of managers and store classification. Across all three panels, there seems to be a small positive relationship between store classification and manager characteristics. However the strongest relationship is in Panel (c), where higher educational attainment is associated with higher store classification.

Figure 11: Manager Characteristics and Store Classification



(c) Educational Attainment



A.3 Estimates of Sales Forecasting Equation

Table 12 summarizes our estimation results of the sales forecasting function. For each store and product, we estimate a Negative Binomial regression function using Maximum Likelihood. The set of explanatory variables includes the logarithm of retail price, the logarithm of the store-product sales in the last week, and two seasonal dummies: a weekend dummy, and a holiday dummy for a major holiday. For each product, Table 12 reports the three quartiles in the distribution across stores of parameter estimates and their respective standard errors for the coefficient of log-price, the coefficient of log-lagged-weekly-sales, and the over-dispersion parameter in Negative Binomial model. Given our interest in the forecasting power of this equation, we also report the three quartiles of McFadden's Pseudo R-squared coefficient (i.e., one minus the ratio between the log-likelihoods of the estimated model and a model only with a constant term).

The estimates of for the lagged-sales coefficient show very substantial time persistence for all products and most stores. Standard errors show that this parameter is estimated with enough precision. The estimates for the log-price coefficient are mostly negative and large in absolute value, though they are not precisely estimated as LCBO changes prices quite infrequently. The estimate of the over-dispersion parameter is substantially smaller than one for almost all the stores and products, which implies evidence of over-dispersion and the rejection of the Poisson regression model. The magnitude of the Pseudo R-squared coefficient is on average around 6% for the median store, which seems small. However, it is important to note that the uncertainty about daily sales of a single product and store can be substantially larger than the uncertainty about aggregate sales at monthly level or over products or/and stores.

=	1 st Quartile	2^{nd} Quartile	3 rd Quartile
	Est.	Est.	Est.
	(st.err.)	(st.err.)	(st.err.)
SKU #67 – Smirnoff Vodka			· · · ·
$\ln(p)$	-1.9740	-0.6215	0.3442
(-)	(2.1803)	(1.0068)	(1.1901)
$\ln(Q)$	0.3168	0.5202	0.6352
	(0.0784)	(0.0687)	(0.0567)
lpha	0.2569	0.3344	0.5130
	(0.0336)	(0.1439)	(0.2490)
$Pseudo-R^2$	0.0633	0.0852	0.1034
SKU #117 – Bacardi Superior White Rum			
$\ln(p)$	-4.9425	-2.4639	-0.7968
(-)	(3.5171)	(2.6876)	(1.4862)
$\ln(Q)$	0.2901	0.4855	0.6233
	(0.0925)	(0.0811)	(0.0552)
lpha	0.2246	0.2967	0.4746
	(0.0532)	(0.0401)	(0.0795)
$Pseudo-R^2$	0.0565	0.0840	0.1015
SKU #340380 - Two Oceans Sauvignon Blanc			
$\ln(p)$	-6.1316	-4.7974	-3.3746
(-)	(0.9763)	(0.6034)	(0.6559)
$\ln(Q)$	0.1208	0.2760	0.4100
	(0.0801)	(0.0802)	(0.0882)
lpha	0.4366	0.6963	1.2437
	(0.0421)	(0.1397)	(0.1145)
$Pseudo-R^2$	0.0361	0.0534	0.0710
SKU #550715 – Forty Creek Select Whisky			
$\ln(p)$	-4.4488	-2.2356	-0.5021
(-)	(3.4604)	(3.9153)	(2.3409)
$\ln(Q)$	0.1275	0.2772	0.4409
	(0.0890)	(0.0848)	(0.0723)
lpha	0.1507	0.2112	0.3286
	(0.0182)	(0.0577)	(0.0291)
$Pseudo-R^2$	0.0469	0.0727	0.0880
SKU #624544 – Yellow Tail Shiraz Red			
$\ln(p)$	-4.5672	-3.1739	-2.2505
	(1.1090)	(1.6732)	(1.5370)
$\ln(Q)$	0.2080	0.4003	0.5749
	(0.0762)	(0.0772)	(0.0625)
lpha	0.3849	0.5730	0.9906
	(0.0333)	(0.0425)	(0.1443)
$Pseudo-R^2$	0.0299	0.0491	0.0642

Table 12: Sales Forecasting Equation – Negative Binomial Model

A.4 Stockouts at the Warehouse Level

	Days with a warehouse stockout	% of total days in sample
SKU~#~67	3	0.4~%
SKU # 117	1	0.1~%
SKU ~#~ 340380	1	0.1~%
$SKU \ \# \ 550715$	1	0.1~%
SKU~#~624544	1	0.1~%

Table 13: Warehouse Stockout Events







Figure 13: Daily Number of Stores Receiving Deliveries

A.5 Store heterogeneity in estimated (S,s) thresholds

We investigate the heterogeneity across stores in the estimated thresholds. For each storeproduct pair, we begin by obtaining the *log-lower-threshold* and the *log-upper-threshold* evaluated at the mean value of log-price and retail-specific mean value of log-expected-demand. We denote these log-thresholds as *log-s*₀ and *log-S*₀, respectively. Figure 14 presents the inverse CDF of the store-specific average estimates of *log-s*₀ and *log-S*₀ and the Bonferroni 95% confidence band under the null hypothesis of store homogeneity. These distributions show very significant differences in store level estimates. For the log-lower-threshold, 98% of stores lie outside the confidence bands and therefore have different values of store level logthresholds. For the upper-log-threshold, only 3% of stores lie within the confidence bands, which entails that 97% of stores have different values of store level log-thresholds.

In addition to this between-store heterogeneity, we also observe significant positive correlation between the two log-thresholds. This confirms our previous conjecture from Figure 10, in which we observed a positive correlation between the inventory to sales ratio before and after an order is placed. Again, this correlation can be explained by differences across stores in stockout costs or/and inventory holding costs.

Given that we have estimates at the store-product level, we can also explore within-store heterogeneity. Table 14 below presents a variance decomposition of the log-thresholds s_0

and S_0 . More specifically, we are interested in disentangling how much of the differences we observe in Figure 14 is attributable to variation across stores, and how much is because of differences across products. Table 14 presents an interesting finding: for the lower threshold, between-store variance is significantly larger than within-store variance, while the opposite is true for the upper threshold. That is, the order-up-to quantity seems to be relatively homogeneous across stores, while the safety stock level seems to vary significantly.





Table 14: Variance Decomposition of Log-thresholds

Variance	$Log-s_0$	$Log-S_0$
Between-store Variance	1.03	0.38
Within-store Variance	0.48	0.66

A.6 Details on estimation method of structural parameters

(i) Nonparametric estimation of CCP function. In the first step of the 2PML method, we use the following Kernel method for the estimation of the CCP function. For every $(y, \mathbf{x}) \in \mathcal{Y} \times \mathcal{X}$:

$$\widehat{P}(y|\mathbf{x}) = \frac{\sum_{t=1}^{T} \mathbb{1}\{y_t = y\} K_T(\mathbf{x}_t - \mathbf{x})}{\sum_{t=1}^{T} K_T(\mathbf{x}_t - \mathbf{x})}$$
(25)

where $K_T(\mathbf{u})$ is the Kernel function $1/(1+\sqrt{T}||\mathbf{u}||)$ with ||.|| being the Euclidean distance. (ii) Discretization of state variables. This estimation method applies to models where the vector of state variables \mathbf{x} has discrete support. In principle, our state variables have continuous support. We have applied a K-means clustering method for the discretization of the exogenous state variables. We apply this method separately for each store-product. More specifically, we apply K-means to discretize variables p_t and $\ln Q_t^{[-7,-1]}$. For every store and product pair, we cluster these variables using a k-means algorithm with a squared Euclidean distance metric, and using Arthur and Vassilvitskii (2007)'s k-means++ cluster initialization. For both variables, we impose the number of clusters to be 2. For the endogenous state variable k_t , along with the choice variable y_t , we choose a set of fixed grid points. Specifically, we allow k_t to take values between 0 and 100 with an interval of 2, and y_t to take values between 0 and 48 with an interval of 6. The latter preserves an important aspect of the nature of orders being placed by store managers at LCBO: most orders are placed in multiples of 6, and most order sizes are smaller or equal to 48²⁴. Table 15 below presents the frequency of orders in the choice set. Overall, the grid points in the choice space represent approximately 98% of orders that we observe in the data.

Table 15: Frequency of Orders

	0	6	12	18	24	30	36	42	48
Frequency of orders (in $\%$)	84.29	0.17	10.05	0.04	2.63	0.01	0.91	0.01	0.41

Finally, in Table 16 below, we present a variance decomposition of the state space. Our goal is to assess whether the discretization of the state space is such that it captures most of the variation we observe in the data. The discretization of variables p and $\ln Q$ preserves most of their sample variation, with discretized variance representing approximately 99% and 89% of overall variance, respectively. However, for variable k, the discretization is significantly restrictive. The variance of the discretized variable represents only approximately 20% of the overall variance.

²⁴Note that $\ln d^e$ and σ^2 are indirectly clustered through $\ln Q$ and p, as the sales forecasting equation determines the space of the variables $\ln d^e$ and σ^2

	Overall Variance	Between (Discretized) Variance	Within (Residual) Variance	Between Variance %
k	4,055.56	835.20	2,256.29	20.59
p	51.03	51.00	0.03	99.93
$\ln Q$	1.30	1.16	0.14	89.36

Table 16: State Space Variance Decomposition

(iii) Computing time. Most of the computing time in the implementation of this twostep estimator comes from the calculation of present values, and more specifically from the inversion of matrix $\mathbf{I} - \beta \sum_{y=0}^{J} \mathbf{P}(y) * \mathbf{F}_{x}(y)$ that has dimension $|\mathcal{X}| \times |\mathcal{X}|$. Nevertheless, the computing time to obtain the 2PML for one store-product – using standard computer equipment – was around 20 seconds, and the total computing time for the approximately $634 \times 5 = 3,160$ store-products in our working sample was less than 18 hours.

A.7 Empirical distribution of parameter estimates

Figure 15 plots the empirical density across stores and products of our raw estimates of the four structural parameters, measured in dollar amounts. These empirical densities show substantial heterogeneity across stores and products in the four parameter estimates.



A.8 Shrinkage estimator

To correct for excess dispersion, we consider the following *shrinkage estimator* (see Gu and Koenker [2017]):

$$\widehat{\gamma}_{i,j}^* = \overline{\widehat{\gamma}} + \left(1 - \frac{\widehat{\sigma}_{i,j}^2}{Var(\widehat{\gamma})}\right)^{1/2} \left(\widehat{\gamma}_{i,j} - \overline{\widehat{\gamma}}\right)$$

where $\hat{\gamma}_{i,j}$ is the original parameter estimate; $\hat{\sigma}_{i,j}$ is its standard error, and $\bar{\hat{\gamma}}$ and $Var(\hat{\gamma})$ represent the mean and variance, respectively, in the empirical distribution of $\hat{\gamma}_{i,j}$ across stores and products. This estimator generates a distribution of estimates across stores-products that corrects for the spurious heterogeneity due to estimation error. By construction, we have that $Var(\hat{\gamma}_{ij}^*) = Var(\hat{\gamma}_{ij}) - E(\sigma_{ij}^2)$.

A.9 Heterogeneity in realized inventory management costs

In Figure 16, the blue curves represent the CDFs across stores and products of each of the four cost-to-revenue ratios. These distributions show the following ranges between percentiles

5% and 95%: [0.1%, 0.7%] for the inventory holding cost; [0.0%, 0.03%] for the stockout cost; [0.4%, 2.75%] for the fixed ordering cost; and [0.05%, 0.6%] for the variable ordering cost. We can see that the realized fixed ordering costs are not only the costs with larger contribution to the firms' profit, but also with larger heterogeneity across stores.

The dispersion across stores in these cost-to-revenue ratios is the combination of dispersion in structural parameters and dispersion in decision and state variables affecting these costs. In particular, managers' optimal inventory decisions can partly compensate for the heterogeneity in the structural parameters. For instance, the inventory holding cost to revenue ratio for store i and product j is $\gamma_{i,j}^h \bar{k}_{i,j}/\bar{r}_{i,j}$. A store-product with large per-unit inventory holding cost, $\gamma_{i,j}^h$, will tend to keep smaller levels of inventory than a store-product with a small value of this parameter such that the difference between these stores in the ratio $\gamma_{i,j}^h \bar{k}_{i,j}/\bar{r}_{i,j}$ will be smaller than the difference between their per unit inventory holding cost. To measure the magnitude of this behavioural response by store managers, the red curves in Figure 16 present the CDFs of the cost ratios when we replace the store-product specific structural parameters by their means across products, but we keep the values of decisions and state variables. That is, for the inventory holding cost ratio, the red curve is the CDF of variable $\bar{\gamma}_i^h \bar{k}_{i,j}/\bar{r}_{i,j}$. For each of the four inventory ratios, the counterfactual CDFs in the red curves are steeper than the factual CDFs in the blue curves. Store managers with a perception of higher inventory costs make decisions that entail lower costs of managing their inventories relative to revenue.

Figure 16: Empirical CDFs of Realized Inventory Management Costs to Revenue Ratios

(a) Holding Cost to Revenue



(c) Fixed Ordering Cost to Revenue



(b) Stockout Cost to Revenue



(d) Variable Ordering Cost to Revenue



A.10 Dispersion of cost parameters

In Table 17, we explore how the dispersion of the manager component of costs depends on store and location characteristics. Consistent with the interpretation of managers' biased perception of true costs, we find that managers in high-type stores have a smaller dispersion in this component of costs.

	$(res(\gamma^h))$	$(res(\gamma^z))$	$(res(\gamma^f))$	$(res(\gamma^c))$
	Est.	Est.	Est.	Est.
	(s.e.)	(s.e.)	(s.e.)	(s.e.)
Store Class				
AA	0.000822***	0.0210**	0.0149	0.00154
	(0.000212)	(0.00975)	(0.0589)	(0.00349)
A	0.000793^{***}	0.0343***	-0.0271	0.00295
	(0.000167)	(0.00928)	(0.0583)	(0.00341)
В	0.000690***	0.0359***	-0.00822	0.00179
	(0.000193)	(0.0112)	(0.0656)	(0.00354)
C	0.000529**	0.0270^{*}	0.0814	0.00149
	(0.000241)	(0.0139)	(0.0831)	(0.00393)
D	0.0000842	0.0263^{*}	0.120	0.000249
	(0.000265)	(0.0153)	(0.0989)	(0.00411)
ln(Product Assortment Size)	0.000138	0.00490	-0.0536	0.00204*
	(0.000109)	(0.00616)	(0.0367)	(0.00121)
ln(Population in City)	0.0000201	0.000206	0.0150	0.000304
	(0.0000198)	(0.00118)	(0.00925)	(0.000301)
ln(Median Income in City)	-0.000583**	-0.000222	0.176^{*}	0.00133
、 、	(0.000287)	(0.0157)	(0.106)	(0.00410)
Location dummies (25 regions, 4 districts)	YES	YES	YES	YES
Product dummies (5 products)	YES	YES	YES	YES
R-squared	0.1038	0.0403	0.0712	0.0215
Observations	2,589	2,589	2,589	2,589

Table 17:	Regression	of Cost	Dispersion	on Store	Characteristics
-----------	------------	---------	------------	----------	-----------------

(1) Location dummies based on LCBO's own division of Ontario into 25 regional markets and 4 districts.

(2) Robust standard errors clustered at the store level in parentheses

(3) * means p-value <0.10, ** means p-value <0.05, *** means p-value <0.01

A.11 Manager bias: granular examples

The first pair of LCBO stores we explore in Table 18 is store #452 (1138 Avenue Road) and store #572 (1245 Dupont Street), both located in the Toronto-North area. Although both stores are of type "A" and have similar average weekly sales, their managers have very different years of experience at LCBO and significantly different estimates of cost parameters. Specifically, the manager of store #572 has an additional 16 years of experience at LCBO, a 29% higher average holding cost, a 65% lower average stockout cost, a 13% higher average fixed ordering cost, and a 19% lower average unit ordering cost. The second pair of stores we examine is store #538 (122 Rideau Street) and store #547 (111 Albert Street), both located in the Ottawa-Central area. Again, although the two stores are of type "B" and have similar average weekly sales, the managers have very different experience levels and significantly different cost estimates. Store #547 has 34 fewer years of experience, a 34% lower average holding cost, a 27% higher average stockout cost, a 39% higher fixed ordering cost, and a (mere) 1% lower average stockout cost.

	Store Type	Average Weekly Sales	Manager Years at LCBO	$\hat{\gamma}^h$	$\hat{\gamma}^z$	$\hat{\gamma}^{f}$	$\hat{\gamma}^c$
Store #452 Store #572	A A	92,422 80,376	13 29	$\begin{array}{c} 0.003347 \\ 0.004344 \end{array}$	0.097933 0.033821	2.17968 2.45856	0.061739 0.050037
Store # 538 Store # 547	B B	$\frac{44,254}{37,888}$	35 1	0.006972 0.004608	0.099627 0.126806	2.05306 2.86206	$0.047136 \\ 0.046638$

 Table 18: Manager Biases in Cost Parameters