University of Toronto Department of Economics



Working Paper 722

Decentralized Decision-Making in Retail Chains: Evidence from Inventory Management

By Victor Aguirregabiria and Francis Guiton

March 27, 2022

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

Victor Aguirregabiria †Francis Guiton ‡University of Toronto, CEPRUniversity of Toronto

March 28, 2022

Abstract

This paper examines the effects of decentralizing decision-making in multi-establishment firms. Using a unique dataset from the Liquor Control Board of Ontario (LCBO), we assess the impact of allowing store managers to control the inventory replenishment decisions of their stores. We first present evidence of strong heterogeneity across the inventory decisions of the 634 store managers in the retail chain. We then study the sources of this heterogeneity by focusing on differences across stores in the structure of the profit function. We estimate a separate dynamic structural model of inventory management for each store manager in the retail chain using two years of daily data. We find very substantial heterogeneity across stores in unit costs of holding inventory, placing orders, and stockouts, and in fixed ordering costs. Using counterfactual experiments based on the estimated model, we find that a centralized inventory management system would yield a 0.4% increase in annual profit for LCBO, equivalent to \$6.8 million. This modest effect is the result of combining two substantial effects with opposite signs: the negative effect of delaying the processing of information is more than compensated by the large reduction in ordering and storage costs from eliminating store managers' behavioral biases and heterogeneous skills (i.e., 25.5%) on average, 6.3% for the median store). Furthermore, the gains from centralization are very heterogeneous across stores in the retail chain, with both very substantial losses and gains. These distributional effects within multi-divisional companies can be relevant when deciding whether to adopt an organizational change.

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 helpful comments from Loren Brandt, Avi Goldfarb, Jiaying Gu, Jordi Mondria, Alvaro Pinzon, Sixia Xie, and seminar participants at the University of Toronto, and at the 2021 EARIE conference.

[†]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 within their organization. Decisions can be centralized at the headquarter level or delegated to individual establishments in a more decentralized structure. Finding a decision-making process that is optimal for a firm requires the evaluation of different trade-offs. A more decentralized decision-making structure allows local managers to make use of valuable information at the store level that can be inaccessible or difficult to communicate *just-in-time* to headquarters. However, decentralized decision-making gives autonomy to local managers who are heterogeneous in their skills to acquire and process information, which can translate into sub-optimal outcomes for the firm. A multi-establishment firm must assess this trade-off when choosing the degree of decentralization in its decision-making structure.

In this paper, we study inventory management decisions of store managers at the *Liquor Control Board of Ontario* (LCBO), and the impact on the firm's performance of allowing these decisions to be autonomous at the store manager level. The LCBO is a provincial government enterprise that sells alcohol across the province of Ontario. As a decentralized retail chain, each store has some autonomy in its decision-making process. More specifically, managers have discretion regarding two aspects of their store's inventory: which products to offer (*assortment decisions*); and when and how much to order of each product (*replenishment decisions*). The latter involves forming expectations about future demand in order to determine both the optimal order quantity and when to place the order. We use a very unique dataset from LCBO with daily information on inventories, orders, sales, stockouts, and prices at every store and for every product (SKU) from October 2011 to October 2013 (677 days).

First, we find strong evidence in support of store managers following (S, s) rules in their replenishment decisions.¹ Most importantly, after controlling for the different demand level at each store, we find evidence of substantial heterogeneity across stores in their (S, s) thresholds. Observable store characteristics – such as size and geographic location – explain only part of the differences across store managers in their inventory decision rules.

Second, we propose and estimate a dynamic structural model of inventory management to study the determinants of the heterogeneity in store managers' (S, s) rules. In this model, we focus on the differences across stores in the structure of the profit function, and more specifically, in the costs of holding inventory, stockouts, and ordering. Given the daily frequency of the data, we obtain precise estimates of holding cost, stockout cost, and ordering costs at the level of individual product (SKU) and store. We find substantial heterogeneity across stores in all the

¹The lower threshold s defines the safety stock level, and the upper threshold S is the order-up-to stock level.

(revealed-preference) cost parameters. Similarly as for the (S, s) thresholds, observable store characteristics explain less than 50% of this heterogeneity. The remaining heterogeneity can be attributed to local managers' idiosyncratic perceptions about store-level costs, because of either superior private information or misconceptions about the actual costs.

Third, we use the estimated store-product-specific structural models to measure the contribution of store manager heterogeneity to different inventory outcomes. Overall, removing store managers' idiosyncratic heterogeneity in the four cost parameters generates, on average, a substantial effect on inventory management, with a 6-day increase in waiting time between two orders, a decrease in the average order amount equivalent to 1.6 days of average sales, a 22% decrease in the inventory to sales ratio, but a negligible effect on the frequency of stockouts. Accordingly, if the idiosyncratic component of costs is a biased perception by store managers, then it has a substantial negative effect on the firm's profit as it increases storage and ordering costs with almost no effect on stockouts and revenue.

Finally, we evaluate the effects of centralizing the decision-making of inventory management at the LCBO headquarters. This counterfactual experiment takes into account that, according to company reports, store-level information regarding sales is processed by headquarters with a weekly delay (see Section 2.1). Accordingly, the main trade-off that this experiment measures is that a centralized inventory management system eliminates (or reduces) store managers' heterogeneous skills and behavioral biases, but it cannot benefit from store managers' just-in-time information about demand, sales, and inventories. We find that a centralized inventory management system would yield a 0.4% increase in annual profit for LCBO, which corresponds to approximately \$6.8 million. This 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 negative impact on profits. However, this negative effect is more than compensated by the large increase in profits due to reducing ordering and storage costs when removing store managers' behavioral biases and heterogeneous skills (i.e., 25.5% on average, 6.3% for the median store). Furthermore, the gains from centralization are very heterogeneous across stores, with both very substantial losses and gains: percentiles 10 and 90 are -3.5% and 3.3%, respectively. Distributional effects within multi-divisional companies can be relevant when deciding whether to adopt an organizational change (see Inderst, Müller, and Wärneryd (2007)).

This paper contributes to several literatures. First, we contribute to a growing empirical literature on centralization vs. decentralization of decision making in multi-division firms. DellaVigna and Gentzkow (2019) show that most retail chains in the US charge uniform prices across stores, despite there being large differences in demand elasticities and gains from thirddegree price discrimination. The authors discuss possible explanations for this finding. Hortacsu, Natan, Parsley, Schwieg, and Williams (2021) study the pricing system of a large international airline company that combines different decision rights across multiple organizational teams. The authors find that this pricing system — despite using advanced optimization techniques and automation — does not internalize consumer substitution to other products, uses persistently biased demand forecasts, and does not respond to changes in opportunity costs. These inefficiencies are mostly related to the limited coordination between teams.²

Second, our paper relates to a large literature on *Team Theory*, pioneered by Marschak and Radner (1972). Team theory studies information processing and decision making in firms characterized by dispersed information and by physical constraints that make it costly to communicate and process information. In our model, store managers have their own perception of the firm's profit, which can be the result of private incentives, different skills, or behavioral biases. In this aspect, our framework relates to the literature on coordination in organizations and the trade-off between coordination and acting independently.³

Third, the paper is related to the empirical literature of (S, s) decision rules, with contributions by Blinder (1981), Mosser (1991), Aguirregabiria (1999), Hall and Rust (2000), Kryvtsov and Midrigan (2013), Bray, Yao, Duan, and Huo (2019), and Bray and Stamatopoulos (2021) for firm inventories, and by Eberly (1994), Attanasio (2000), Adda and Cooper (2000), and Foote, Hurst, and Leahy (2000) for household purchases of durable products. We contribute to this literature by using high-frequency (daily) data at the store and product level to estimate cost parameters using a dynamic structural model for each store and product.

Finally, our paper contributes to the literature in empirical industrial organization on structural models with boundedly rational firms. Most of this literature studies firms' entry/exit decisions (Goldfarb and Xiao, 2011, 2019; and Aguirregabiria and Magesan, 2020), pricing decisions (Huang, Ellickson, and Lovett, 2020; Ellison, Snyder, and Zhang, 2018), and bidding behavior in auctions (Hortaçsu and Puller, 2008; Doraszelski, Lewis, and Pakes, 2018; Hortaçsu, Luco, Puller, and Zhu, 2019). As far as we know, our paper is the first one to investigate bounded rationality in firms' inventory decisions.

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.

²In our application, LCBO combines centralized uniform pricing with decentralized inventory management.

³See Alonso, Dessein, and Matouschek (2008), Hart and Holmstrom (2010), and Dessein, Garicano, and Gertner (2010).

2 Firm and data

2.1 LCBO retail chain

LCBO was founded in 1927 as part of the passage of the Ontario Liquor Licence Act.⁴ 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 to "[generate] maximum profits to fund government programs and priorities".

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

As a retail chain, LCBO operates stores across the province of Ontario. Stores report to headquarters, who oversee the overall chain operations. In terms of inventories, stores order product from the closest distribution centre. The order is then periodically delivered by trucks, according to a pre-determined route and schedule.

Store managers at the LCBO are responsible for the overall operation of their store. As an incentive, a fixed percentage of the managers' salaries is based on the performance of their store. In terms of their role, the main responsibility of store managers is to oversee their store's inventory level for each product they offer, and to ensure that daily demand is met. In order to do so, managers have autonomy to decide product assortment and replenishment of inventory in their store, and how much to restock those products. Managers therefore have a dual responsibility when it comes to the inventory management of their store: providing products that are in high demand, and keeping these products stocked. In this paper, we focus on the inventory replenishment decisions.

In order to assist store managers in their inventory decisions, headquarters provide daily recommendations to managers regarding how much to order of each product at their store. In the company's internal jargon, these recommendations are referred to as *Suggested Order Quantities*

⁴The 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). 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.

⁵See Aguirregabiria, Ershov, and Suzuki (2016) for further details about markups at LCBO.

(SOQs). For each store in the retail chain, headquarters generate daily order recommendations based on the previous week's sales and inventory information, and therefore provide a schedule for the retail chain's overall inventory replenishment. The store-level information used to generate the order recommendations include the *Average Rate of Sale (ARS)* of the product from the previous week, as well as seasonal brand factors.

Note that the headquarters's recommendation is based on sales information with a one week lag. This informational friction plays an important role in the decentralization decision. Store managers have more updated sales information than the headquarters. However, despite having lagged information, headquarters may use more sophisticated statistical models than some store managers to predict future demand. We take into account this trade-off in our evaluation of the benefits/costs of decentralizing inventory decision-making.

2.2 Data

We use a unique dataset from the LCBO with daily information on inventories, sales, deliveries, and prices of every product sold at each LCBO store between October 2011 and October 2013 (677 days). With 634 stores across Ontario and more than 20,000 different products sold, the dataset includes approximately 700 million observations in total. The dataset also includes information about product characteristics, store characteristics such as location, size, and store type, and about the managers themselves.⁶

Table 1 presents basic summary statistics of the dataset. The average store has an assortment of 2,029 items (SKUs); sells 12,909 units per week (6.36 units per SKU); generates a weekly revenue of \$162,250 (\$80 per SKU and week, and \$12.6 per unit sold); receives deliveries 5.37 days per week, with these deliveries containing 12,258 units; and has 415 stockout events per week. According to the table, these numbers vary significantly across the different types of stores. As expected, larger stores generate higher weekly revenues than smaller ones: the average AAAand AA stores generate weekly revenues of \$874, 367 and \$518,067 respectively, while the average C and D stores generate weekly revenues of \$59,328 and \$23,913, respectively. Stockout events seem to occur more frequently at larger stores than smaller ones on average, with 1,203 stockout events occurring at the average AAA store per week against 204 stockout events per week at the average D store. Finally, larger stores seem to place orders more frequently and in greater quantity. The average AAA and AA stores receive orders 6.30 and 6.27 days per week for a total of 51,340 and 35,711 units respectively, while the average C and D stores only receive orders 5.09 and 3.83 days per week for a total of 4,395 and 1,660 units, respectively.

⁶LCBO classifies their stores into six categories, from largest (i.e. highest-selling) to smallest (i.e. lowest-selling): AAA, AA, A, B, C, and D.

			Type of	store			
	All	AAA	AA	A	B	C	D
	Mean						
	(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

The bottom panel in Table 1 presents inventory-to-(daily)sales ratios and ordering frequencies. These statistics are 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 correspond to thresholds s and S, respectively. On average, a store keeps enough inventory to cover demand of a product for 23 days, places an order when there is inventory for 9 days, and the ordered amount covers sales for 18 days. For the average store and product, orders are placed once every two weeks (i.e., $0.07 \simeq 1/14$). Average inventory-to-sales ratios decrease with store size/type, though this difference captures also a composition effect due to different product assortments across store types. In sections 3 to 5, we control for these composition effects by focusing on a set of products carried by all the stores.

Figure 1 shows strong heterogeneity across stores in several measures related to inventory management. The figures in panels (a) to (f) are inverse cumulative distributions over stores,

together with their 95% confidence bands.⁷

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

$$Stockout \ rate_i = \frac{\# \ product - day \ observations \ with \ stockout_i}{\# \ total \ product - day \ observations_i}$$
(1)

The figure shows a significant spread of stockout rates: the 10^{th} percentile is 2% and the 90^{th} percentile is 10%.

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}_i and J_i to represent the set and the number of products in store i, respectively, the *revenue-loss* for store i is:

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

The 10th and 90th percentiles are \$0.20 and \$1.00 per product-day, respectively. Aggregated at the annual level and over all products offered in a store, they imply an average annual revenue-loss of approximately \$150,000 at the 10th percentile and \$750,000 at the 90th percentile.

Panel (c) presents the ordering frequency of stores in our sample calculated as:

$$Ordering \ frequency_i = \frac{\# \ product - day \ observations \ with \ order_i}{\# \ total \ product - day \ observations_i}$$
(3)

This ordering rate varies significantly across stores, with the 10^{th} percentile being 0.05 and the 90^{th} being 0.13.

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:

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

The distribution of the *inventory-to-sales ratio* shows that stores at the 10^{th} and 90^{th} percentiles hold inventory for 14 days and 35 days of average sales, respectively. For the upper threshold *S* (in Panel (e)), the values of these percentiles are 10 and 30 days, and for the lower threshold *s* (in Panel (f)) they are 5 and 13 days.

⁷For every store, the 95% confidence interval is based on the construction of store-product specific stockout rates. For a given store, we have thousands of these product specific rates. The 95% confidence interval is determined by percentiles 2.5% and 97.5% in this distribution.



(c) Ordering Frequency









(d) Inventory to Sales Ratio







Given these substantial differences across stores inventory outcomes, it is interesting to explore how they vary together. In Figure 2 below, we present five scatter plots at the store level: (a) stockout frequency against ordering frequency; (b) stockout frequency against inventory-tosales 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) inventoryto-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.⁸ 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 s but 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 small positive relationship between the stockout frequency and the ordering frequency. That is, stores placing orders more frequently tend to have higher stockout rates. This correlation could be explained by heterogeneity in inventory holding costs, as a higher storage cost implies both a larger ordering frequency (with a smaller amount per order) and a larger stockout rate due to a lower safety s threshold.

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.

⁸For 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).











(d) Stockouts and Inventory, Before Order





(b) Stockouts and Inventory

2.3 Working sample

In the econometric models that we estimate in Sections 3 and 4, the parameters are unrestricted at the store-product level. There are almost 2 million store-product pairs in our data. For the sake of time saving, we have not estimated these models for every store-product in the data. This would have been especially time-consuming for the dynamic structural model in Section 4. Instead, we have estimated these models for every store, but only for a selection of 5 products. In this section, we explain our criteria in the selection of these products and present summary statistics for this working sample.

The product basket we select for estimation in Sections 3 and 4 relies on two criteria. First, we choose products that are high-selling across all LCBO stores. That is, we include products that are in high demand across the province, which entails that inventory decisions for these products have a substantial contribution to the firm's profit. Second, since our parameters are unrestricted at the store-product level, we include products from each broad category in order to account for product-level heterogeneity. These product categories include white wine, red wine, vodka, whisky, and rum. Based on these criteria, Table 2 presents the selected 5 products.

	SKU	SKU	SKU	SKU	SKU
	#67	#117	#340380	#550715	#624544
Product Information					
Name	Smirnoff Vodka	Bacardi Superior White Rum	Two Oceans Sauvignon Blanc	Forty Creek Barrel Select Whisky	Yellow Tail 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 below presents summary statistics for our working sample. The average store in our working sample sells 91 units per week (18 units per SKU); generates a weekly revenue of \$1,687 (\$347 per SKU and week, and \$18 per unit sold); receives deliveries 2 days per week, with these deliveries containing 88 units; and has 0.37 stockout events per week. Similarly to Table 1, these numbers vary significantly across the different types of stores. While the average AAA and AA store generates weekly revenues of \$5,198 and \$4,512 respectively in our working sample, the smaller C and D stores generate only \$821 and \$302 on average, respectively. Contrary to our full sample, stockout events seem to occur more frequently in smaller stores than larger ones in our working sample (0.56 stockouts per week for the average D store, and 0.14 stockouts per week for the average AAA store). Finally, deliveries seem to follow the same trend as in Table

1, with bigger stores placing larger and more frequent orders. The average AAA store receives deliveries 4 days per week for a total of 272 units per week, while the average D store only receives deliveries 0.6 days per week for a total of 15 units per week.

Given our selection criteria of high-selling products for our working sample, the mean values for the inventory-to-sales ratios – overall, after an order is received, and before an order is placed – are all above the mean values for full sample: 20.40, 19.30, and 13.38 in the working sample versus 14.35, 17.63, and 10.29 in the full sample. That is, our working sample presents higher inventory on hand across stores, and higher values of S and s on average.

			T	ype of stor	re		
	All	AAA	AA	Â	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

3 (S,s) decision rules

3.1 Model

In this section, we study store managers' inventory behaviour through the eyes of (S, s) decision rules. In its simpler version, the threshold values that define this decision rule are time-invariant. Under lump-sum (fixed) ordering costs, time-invariant expected demand and ordering costs, and general conditions on the profit function, the optimal (profit-maximizing) inventory decision rule has a (S, s) structure (Arrow, Harris, and Marschak, 1951; Scarf, 1959; Denardo, 1981). This rule is described by a lower threshold s that determines the stock level that triggers a new order (the so called *safety stock level*), and an upper threshold S that represents the stock level to reach when an order is placed. That is, if k_t represents the stock level at the beginning of day t, and y_t is orders at day t, then a (S, s) rule implies:

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

$$\begin{cases} S = f_S\left(d^e, \gamma^h, \frac{\gamma^f}{\gamma^c}, \frac{\gamma^z}{\gamma^c}\right) \\ + & - & ? & + \\ s = f_s\left(d^e, \gamma^h, \frac{\gamma^f}{\gamma^c}, \frac{\gamma^z}{\gamma^c}\right) \\ + & - & - & + \\ S - s = f_{S-s}\left(d^e, \gamma^h, \frac{\gamma^f}{\gamma^c}, \frac{\gamma^z}{\gamma^c}\right) \\ + & - & + & ? \end{cases}$$
(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 and per unit. In fact, these γ 's are the parameters in the structural model that we estimate in Section 4.

The optimality of the (S, s) decision rule extends to models with state variables – other than

inventory – that evolve over time according to an exogenous Markov process.⁹ Let \mathbf{z}_t be the vector of these exogenous state variables. For instance, \mathbf{z}_t may include variables affecting the state of demand, unit ordering costs, wholesale prices, and the product's retail price, when this price is taken as given by the store manager, as it is in our problem. The optimal decision rule has a (S_t, s_t) structure, where the thresholds S_t and s_t are time-invariant functions of these state variables: $S(\mathbf{z}_t)$ and $s(\mathbf{z}_t)$.

Our empirical approach in this section is in the same spirit as Eberly (1994), Attanasio (2000), and Adda and Cooper (2000). These papers use household-level data on purchases of a durable product (automobiles) to estimate (S_t, s_t) decision rules where the thresholds are functions of household characteristics, prices, and aggregate economic conditions. This can be interpreted as a *semi-structural approach* in the sense that using (S_t, s_t) rules is motivated by a dynamic programming model of optimal behavior, but the specification of the thresholds as functions of state variables does not incorporate explicitly the structural parameters of the model. Section 4 presents a full structural approach, and in Section 5, we use the estimated structural model to make counterfactual policy experiments that provide answers to the questions motivating this paper. Here, we find useful to explore the data using a less restrictive empirical framework that is still consistent with the structural model.

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 store-product, we can estimate the parameters in the (S_t, s_t) decision rules at the store-product level. This is important in order to distinguish product and store differences in these parameters, such that we do not spuriously interpret as store heterogeneity what is actually the result of different product assortments across stores. 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:

$$\begin{cases} S_{t} = \exp\{\beta_{0}^{S} + \beta_{d}^{S} \ln d_{t}^{e} + \beta_{p}^{S} \ln p_{t} + u_{t}^{S}\} \\ s_{t} = \exp\{\beta_{0}^{s} + \beta_{d}^{s} \ln d_{t}^{e} + \beta_{p}^{s} \ln p_{t} + u_{t}^{s}\} \end{cases}$$
(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.¹⁰

 $^{^{9}}$ By exogenous process, here we mean that the transition probability function of these state variables does not depend on the firm's orders or inventories.

¹⁰For 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.

That is, 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 to each store (Liquor Control Board of Ontario, 2016). According to this report, LCBO central managers obtain predictions of demand for each store-product using information on the product's retail price (p_t) , the average daily sales of the store-product at the previous week (that we represent as Q_{t-1}), and seasonal dummies. We consider that demand has a Negative Binomial distribution where the logarithm of expected demand at period t has the following form:

$$\ln d_t^e = \ln \mathbb{E} \left(q_t | p_t, Q_{t-1} \right) = \boldsymbol{\alpha}' \ h \left(\ln p_t, \ln Q_{t-1} \right) \tag{8}$$

where q_t is the quantity sold of the store-product at day t, α is a vector of parameters which are constant over time but vary freely across store-products, and $h(\ln p_t, \ln Q_{t-1})$ is a vector of monomial basis in variables $\ln p_t$ and $\ln Q_{t-1}$.

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 and its parameters are not causal parameters. For this inventory decision problem, managers do not need to know the demand function and in particular the causal effect of price on sales. They only need to obtain the best possible predictor of future sales given the information they have. Obtaining this predictor does not require estimating the causal effect of price on demand. 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, Ershov, and Suzuki, 2016). There is substantial and growing empirical evidence showing 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, Hortacsu, Natan, Parsley, Schwieg, and Williams (2021)'s study of the pricing system of a large international airline company, DellaVigna and Gentzkow (2019) and Hitsch, Hortacsu, and Lin (2021) on uniform pricing at US retail chains, Cho and Rust (2010) on pricing of car rentals, or Miravete, Seim, and Thurk (2020) for liquor stores in Pennsylvania.

In the Appendix (Section 7.1), 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; $\tilde{u}_t^s \equiv u_t^s / \sigma_{u^s}$ 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:

$$b_k^s = \frac{-1}{\sigma_{u^s}}; \quad b_0^s = \frac{\beta_0^s}{\sigma_{u^s}}; \quad b_d^s = \frac{\beta_d^s}{\sigma_{u^s}}; \quad b_p^s = \frac{\beta_p^s}{\sigma_{u^s}}$$
(10)

These equations imply that, given the parameters b^s , we can identify the parameters β^s and σ_{u^s} . We assume that \tilde{u}_t^s has a Logistic distribution, such that equation (9) is a Logit 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 regression model:

$$\ln(k_t + y_t) = \beta_0^S + \beta_d^S \ln d_t^e + \beta_p^S \ln p_t + u_t^S \quad \text{if } y_t > 0.$$
(11)

Equation (11) 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 (11) 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 (11). 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 upperthreshold in the right-hand-side of regression equation (11). This exclusion restriction in (S, s)models have been pointed out before by Bertola, Guiso, and Pistaferri (2005). Therefore, using (9) as the selection equation, we can identify the parameters β^S in (11) using a Heckman two-step approach.

3.2 Estimation of (S, s) thresholds

Figure 3 presents the cloud of point estimates of parameters b_0^s , b_k^s , b_d^s , and b_p^s in the lower threshold for every store and product. For each parameter, the figure presents also a curve with the average estimate for each store, where the average is obtained over the five products. 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.¹¹ 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 90%, 95%, 97%, and 90% of stores lie outside of Bonferroni confidence interval, respectively.



¹¹This 95% confidence interval applies Bonferroni correction for multiple testing.

Figure 4 presents the cloud of point estimates of the parameters β^{S} in the upper threshold, as well as the inverse CDF of the store specific average, and 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 96% of stores lie outside of the confidence bands, respectively.









Table 4 presents a decomposition of the variance of parameter estimates into within-store (between products) and between-stores variance. The parameters associated to 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	$oldsymbol{eta}_0^S$	eta_d^S	eta_p^S
Between-store Variance	42.36	0.03	0.61	8.01	47.14	0.31	9.31
Within-store Variance	80.56	0.03	0.05	15.12	91.99	0.08	17.86

 Table 4: Variance Decomposition of Parameter Estimates

The estimated parameters imply estimates for the (S_t, s_t) thresholds. We now investigate heterogeneity across stores in these thresholds. For each store-product 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 5 presents the cloud of point estimates of *log-s*₀ and *log-S*₀ for every store-product, the inverse CDF of the store-specific average, and the Bonferroni 95% confidence band under the null hypothesis of store homogeneity. These distributions show very significant differences in both store-product level estimates and store level estimates. For the log-lowerthreshold, 96% of stores lie outside the confidence bands and therefore have different values of store level log-thresholds. For the upper-log-threshold, only 1% of stores lie within the confidence bands, which entails that 99% 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 2, 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.

In order to explore the sources of this heterogeneity, Table 5 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 5 is attributable to variation across stores, and how much is because of differences across products. Table 5 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.

Variance	$Log-s_0$	$Log-S_0$
Between-store Variance	1.54	0.39
Within-store Variance	0.99	0.68

Table 5: Variance Decomposition of Log-thresholds

4 Structural model

We propose and estimate a dynamic structural model of inventory management. A price-taking 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 next period's inventory. The store-level profit function incorporates four store-specific costs associated with inventory management: a per-unit inventory holding cost, a stockout cost, a fixed ordering cost, and a per-unit ordering cost. For notational simplicity, we omit store and product indexes.

4.1 Sequence of events and profit

Time is indexed by t and 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.

Step (ii). The store manager orders y_t units of inventory from the distribution centre. 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. The ordered amount y_t is a discrete variable with support set $\mathcal{Y} \equiv \{0, 1, ..., J\}$.

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

$$\begin{cases} \ln d_t^e = \eta_{0,t} + \eta_p \, \ln p_t + \eta_Q \, \ln Q_{t-1} \\ \sigma_t^2 = d_t^e (1 + d_t^e \alpha) \end{cases}$$
(12)

where η_p , and η_Q are parameters; $\eta_{0,t}$ represents seasonal effects (i.e., weekend dummy and main holidays dummy, denoted as *seas*); α denotes the overdispersion parameter; and Q_{t-1} is the average daily sales of the product in the store during the previous week. We use F_{d_t} to represent the distribution of d_t at day t, i.e., conditional on p_t , Q_{t-1} , and seasonal effects. 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 (iv). The store sells q_t units of inventory, which is the minimum of supply and demand:

$$q_t = \min\{ d_t, k_t \} \tag{13}$$

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

$$\Pi_t = (p_t - c_t) \min\{d_t, k_t\} - \gamma^z \mathbb{1}\{d_t > k_t\} - \gamma^h k_t - \gamma^c y_t - \gamma^f \mathbb{1}\{y_t > 0\} + \sigma_\varepsilon \varepsilon_t(y_t)$$
(14)

where c_t is the wholesale price, γ^z , γ^h , γ^c , γ^f , and σ_{ε} are store-product-specific structural parameters: $\gamma^z \ \mathbb{1}\{d_t > k_t\}$ is the cost of a stockout in addition to the lost revenue; $\gamma^h k_t$ is the storage cost associated with holding k_t units of inventory at the store; γ^c is a per-unit cost incurred by the store manager when placing an order; γ^f is a fixed ordering cost, and $\varepsilon_t(y_t)$ is a stochastic and mean zero shock in 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.

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, Ershov, and Suzuki, 2016). A constant markup implies that the price-cost margin is proportional to the retail price: $p_t - c_t = L p_t$, where $L = \frac{\tau}{1+\tau}$, τ is the markup, and L stands for Lerner index. This index is equal to L = 0.40 for Ontario products and L = 0.42 for non-Ontario products.

Step (v). Orders y_t placed at the beginning of day 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:

$$k_{t+1} = k_t + y_t - q_t \tag{15}$$

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

4.2 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 $\mathbf{x}_t \equiv (k_t, p_t, \ln Q_{t-1}, seas_t)$ and $\varepsilon_t \equiv (\varepsilon_t(0), \varepsilon_t(1), ..., \varepsilon_t(J))$ and value function $V(\mathbf{x}_t, \varepsilon_t)$. This value function is the unique solution of the following Bellman equation:

$$V(\mathbf{x}_t, \varepsilon_t) = \max_{y_t \in \mathcal{Y}} \{ \pi(y_t, \mathbf{x}_t) + \sigma_{\varepsilon} \varepsilon(y_t) + \beta \mathbb{E} [V(\mathbf{x}_{t+1}, \varepsilon_{t+1}) \mid y_t, \mathbf{x}_t] \},$$
(16)

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 (13) and (15) determines the distribution of $(k_{t+1}, p_{t+1}, \ln Q_t, seas_{t+1})$. 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 (16).

For the solution and estimation of this model, we follow Rust (1987, 1994) and use the integrated value function $V_{\sigma}(\mathbf{x}_t) \equiv \frac{1}{\sigma_{\varepsilon}} \int V(\mathbf{x}_t, \varepsilon_t) d\varepsilon_t$ and the corresponding integrated Bellman equation. Given the Extreme Value distribution of the ε_t variables, the integrated Bellman equation has the following form:

$$V_{\sigma}(\mathbf{x}_{t}) = \ln\left[\sum_{y \in \mathcal{Y}} \exp\left(\frac{\pi(y, \mathbf{x}_{t})}{\sigma_{\varepsilon}} + \beta \mathbb{E}\left[V_{\sigma}(\mathbf{x}_{t+1}) \mid y, \mathbf{x}_{t}\right]\right)\right].$$
 (17)

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

$$\frac{\pi(y_t, \mathbf{x}_t)}{\sigma_{\varepsilon}} = \mathbf{h}(y_t, \mathbf{x}_t)' \gamma, \qquad (18)$$

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:

$$\mathbf{h}(y_t, \mathbf{x}_t)' = (p_t \ \mathbb{E}[\min\{d_t, k_t\} | \mathbf{x}_t], -k_t, -\mathbb{E}[\mathbb{1}\{d_t > k_t\} | \mathbf{x}_t], -\mathbb{1}\{y_t > 0\}, -y_t),$$
(19)

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 .¹² 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 the linear-in-parameters structure of the expected profit $\pi(y_t, \mathbf{x}_t)$, we can represent the integrated Bellman equation in vector form as follows:

$$\mathbf{V}_{\sigma} = \ln \left[\sum_{y \in \mathcal{Y}} \exp \left(\mathbf{H}(y) \ \gamma \ + \ \beta \ \mathbf{F}_{\mathbf{x}}(y) \ \mathbf{V}_{\sigma} \right) \right].$$
(20)

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

 $^{^{12}}$ In the estimation, we discretize the state space using a k-means algorithm.

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

$$P(y|\mathbf{x}_t) = \frac{\exp\{\mathbf{h}(y,\mathbf{x}_t)' \gamma + \beta \mathbb{E}[V_{\sigma}(\mathbf{x}_{t+1}) \mid y, \mathbf{x}_t]\}}{\sum_{j=0}^J \exp\{\mathbf{h}(j,\mathbf{x}_t)' \gamma + \beta \mathbb{E}[V_{\sigma}(\mathbf{x}_{t+1}) \mid j, \mathbf{x}_t]\}},$$
(21)

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 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 as vector of CCPs \mathbf{P} and obtains the corresponding vector of values in the agent behaves according to these CCPs.¹³ In our Logit model, the *policy improvement* mapping has the following vector form, for any $y \in \mathcal{Y}$:

$$\mathbf{P}(y) = \lambda(y, \mathbf{V}) = \frac{\exp\{\mathbf{H}(y) \ \gamma + \beta \ \mathbf{F}_{\mathbf{x}}(y) \ \mathbf{V}\}}{\sum_{j=0}^{J} \exp\{\mathbf{H}(j) \ \gamma + \beta \ \mathbf{F}_{\mathbf{x}}(j) \ \mathbf{V}\}}.$$
(22)

The *valuation* mapping has the following form:

$$\mathbf{V} = \upsilon(\mathbf{P}) = \left[\mathbf{I} - \beta \sum_{y=0}^{J} \mathbf{P}(y) * \mathbf{F}_{x}(y)\right]^{-1} \left[\sum_{y=0}^{J} \mathbf{P}(y) * (\mathbf{H}(y) \ \gamma + euler - \ln \mathbf{P}(y))\right], \quad (23)$$

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

4.3 Estimation method

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 (Aguirregabiria 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:

$$Q(\mathbf{P},\gamma) = \sum_{t=1}^{T} \ln \psi \left(y_t, \mathbf{x}_t ; \mathbf{P}, \gamma \right), \qquad (24)$$

 $^{^{13}}$ See Puterman (2014) for a description of these three mappings in the context of a general dynamic programming problem.

where $\psi(.)$ is the policy iteration mapping defined by the composition of equations (22) and (23). In the first step of the 2PML method, we obtain a reduced form estimation of the vector of CCPs **P** using the following Kernel method: 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. The 2PML estimator is the vector $\hat{\gamma}$ that maximizes the pseudo likelihood function when $\mathbf{P} = \hat{\mathbf{P}}$. That is:

$$\widehat{\gamma} = \arg \max_{\gamma} Q(\widehat{\mathbf{P}}, \gamma) \tag{26}$$

Aguirregabiria and Mira (2002) show that this estimator is consistent and asymptotically normal with the same asymptotic variance as the full maximum likelihood estimator.

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 $\ln p_t$ and $\ln Q_{t-1}$.¹⁴ 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 48.¹⁵

Most of the computing time in the implementation of this two-step 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.

¹⁴For every store and product pair, we cluster the state variables $\ln Q$ and $\ln p$ 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.

¹⁵Note that variables $\ln d^e$ and σ^2 are indirectly clustered through $\ln Q$ and $\ln p$, as the sales forecasting equation determines the space of the state variables $\ln d^e$ and σ^2 . The discretized spaces of the state variables $\ln d^e$ and σ^2 contain around four hundred values each.

4.4 Parameter estimates

Figure 6 plots the empirical density across stores and products of our estimates of the four structural parameters, measured in dollar amounts.¹⁶ Table 6 presents the median from each of these distributions, as well as the median standard errors and t-statistics of the estimates. The median values of the estimates are \$0.0041 for the per-unit inventory holding cost, \$0.0171 for the stockout cost, \$2.9794 for the fixed ordering cost, and \$0.0338 for the per-unit ordering cost. 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, Yao, Duan, and Huo (2017), Bray, Yao, Duan, and Huo (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, 21 and 12.32, respectively), while a substantial fraction of the estimates of the stockout cost are quite imprecise (median t-ratio of 0.25).

	Median	Std. Dev.	Median	Median
	Estimate	Estimate	S.e.	t-stat.
γ^h : Per unit Inventory Holding Cost	0.0040	0.0031	0.0008	5.2331
γ^z : Stockout Cost	0.0179	0.3433	0.1477	0.2545
γ^f : Fixed Ordering Cost	2.9626	1.1403	0.2403	12.3417
γ^c : Per-Unit Ordering Cost	0.0336	0.0697	0.0285	1.4282
# of observed store-product pairs	3,076			
# of store-product pairs with structural estimates	2,840			

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

The empirical densities in Figure 6 show substantial heterogeneity across stores and products in the four parameter estimates. This is consistent with our descriptive evidence on the heterogeneity across stores in (S, s) thresholds and with the positive correlation between the two thresholds. However, part of the dispersion in parameter estimates comes from estimation error and not from true heterogeneity. For any parameter $\gamma_{i,j}$, let $\hat{\gamma}_{i,j}$ be its estimate, which comes

¹⁶More specifically, we first obtain the two-step PML esitmate 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.

from an estimator that is consistent and asymptotically normal with asymptotic variance $\sigma_{i,j}^2$. Using this asymptotic distribution, it is straightforward to establish the following relationship between the variances of $\hat{\gamma}_{i,j}$ and $\gamma_{i,j}$ across stores-products: $Var(\hat{\gamma}_{i,j}) = Var(\gamma_{i,j}) + \mathbb{E}(\sigma_{i,j}^2)$, where $\mathbb{E}(\sigma_{i,j}^2)$ is the mean across stores and products of the asymptotic variances $\sigma_{i,j}^2$. This formula shows that the ratio $r \equiv \mathbb{E}(\sigma_{i,j}^2)/Var(\hat{\gamma}_{i,j})$ captures the excess dispersion – or spurious heterogeneity – due to the estimation error, such that $Var(\gamma_{i,j}) = (1-r) Var(\hat{\gamma}_{i,j})$.

Based on the statistics in Table 6, the ratio between the median standard error and the standard deviation of the estimates is, in percentages: 25.8% for γ^h ; 43.0% for γ^z ; 21.1% for γ^f ; and 48.2% for γ^c . This implies a substantial over-estimation of the actual heterogeneity of cost parameters across stores and products, especially for the stockout and unit ordering costs.



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

$$\widehat{\gamma}_{i,j}^{*} = \overline{\widehat{\gamma}} + \left(1 - \min\left\{1 - \kappa, \frac{\widehat{\sigma}_{i,j}^{2}}{Var(\widehat{\gamma})}\right\}\right)^{1/2} \left(\widehat{\gamma}_{i,j} - \overline{\widehat{\gamma}}\right)$$
(27)

where $\hat{\gamma}_{i,j}$ is the original parameter estimate; $\hat{\sigma}_{i,j}$ is its standard error, $\bar{\hat{\gamma}}$ and $Var(\hat{\gamma})$ represent the mean and variance, respectively, in the empirical distribution of $\hat{\gamma}_{i,j}$ across stores and products; and $\kappa \in (0, 1)$ is a pre-specified constant that is arbitrarily close to zero. This estimator generates a distribution of estimates across stores-products that corrects for the spurious heterogeneity due to estimation error (i.e., for $Var(\hat{\gamma}_{i,j}) > Var(\gamma_{i,j})$).

Figure 7 presents the empirical distribution of parameter estimates across stores using the shrinkage estimator with $\kappa = 0.001$. The comparison of the distributions in figures 6 and 7 shows that estimation error implied very substantial excess of heterogeneity in the original estimates. This is reflected both in a larger variance and in thicker right tail. Nevertheless, after controlling for this excess dispersion issue, we find very substantial heterogeneity across store managers in their perceptions of cost parameters. For the rest of our empirical analysis, we use these shrinkage estimates.





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.

Table 7: Realized Inventory Management Costs to Revenue Ratios

	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 (%)	$\begin{array}{c} 0.3219 \\ 0 \\ 0.8677 \\ 0.2034 \end{array}$	$\begin{array}{c} 0.2139 \\ 0.0079 \\ 0.7203 \\ 0.1951 \end{array}$
Total Inventory Cost to Revenue Ratio (%)	1.4903	0.8451

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.¹⁷ According to our estimate, the total inventory cost to revenue ratio for the median store is approximately 1.5%. This would imply that the retail chain's cost of managing the inventories of their stores would represent approximately 10% of total costs, which entails that non-inventory related costs would account for approximately 90% of total costs (e.g. labour

¹⁷Of course, these expenses do not include the cost of merchandise.

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.



(a) Holding Cost to Revenue





(c) Fixed Ordering Cost to Revenue



(d) Variable Ordering Cost to Revenue



In Figure 8, 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%, 0.8%] for the inventory holding cost; [0.0%, 0.02%] for the stockout cost; [0.4%, 2.75%] for the fixed ordering cost; and [0.05%, 1.0%] 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 8 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 substantially 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.

4.6 Store characteristics and heterogeneity in cost parameters

It is reasonable to believe that a substantial part of the heterogeneity across stores in cost parameters is associated to the store itself and not to the local manager. Store characteristics such as its type according to LCBO classification of stores, its physical area, total product assortment, consumer traffic, or distance to the nearest distribution center can have an impact on store's unit costs. Furthermore, some parameters such as the stockout cost may depend on the socioeconomic characteristics of consumers living in the store's neighborhood. In this section, we investigate to what extent store characteristics and location characteristics explain heterogeneity in cost parameters.

Table 8 presents estimations results from regressions of each estimated cost parameter against the following store and location characteristics: LCBO's store type dummies (6 types); LCBO's regional market dummiest (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.

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 0.40, and 0.50, 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 significantly with assortment size, and are significantly higher for AAA stores relative to B, C, and D stores. In contrast, only 4.2% of the variation in unit ordering costs and 7.6% 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.

	γ^h	γ^{z}	γ^f	γ^c
Store Class	1	/	I	1
AA	-0.000397 (0.000976)	-0.0189 (0.0737)	-0.143 (0.153)	0.00817 (0.00767)
A	-0.000606 (0.000926)	-0.0150 (0.0697)	-0.150 (0.148)	0.0145^{**} (0.00700)
В	-0.00187^{**} (0.000940)	-0.0113 (0.0735)	$0.126 \\ (0.157)$	$\begin{array}{c} 0.0243^{***} \\ (0.00765) \end{array}$
C	-0.00376^{***} (0.000977)	-0.0543 (0.0811)	$\begin{array}{c} 0.833^{***} \\ (0.183) \end{array}$	0.0208^{**} (0.00948)
D	-0.00519^{***} (0.00100)	-0.0430 (0.0860)	$\begin{array}{c} 1.503^{***} \\ (0.199) \end{array}$	0.0221^{**} (0.0109)
$\ln(\text{Product Assortment Size})$	$\begin{array}{c} 0.000521^{***} \\ (0.000190) \end{array}$	-0.00193 (0.0239)	-0.498^{***} (0.0612)	0.00739^{**} (0.00358)
$\ln(\text{Population in City})$	$\begin{array}{c} -0.0000369 \\ (0.0000385) \end{array}$	0.00578 (0.00482)	$\begin{array}{c} 0.0287^{**} \\ (0.0144) \end{array}$	-0.000706 (0.000965)
ln(Median Income in City)	-0.000570 (0.000561)	0.00828 (0.0666)	0.0648 (0.153)	$\begin{array}{c} 0.0182^{**} \\ (0.00919) \end{array}$
Location dummies (25 regions, 4 districts) Product dummies (5 products)	YES YES	YES YES	YES YES	YES YES
R-squared Observations	$0.3971 \\ 2840$	$0.0764 \\ 2840$	$0.4982 \\ 2840$	0.0416 2840

Table 8: Regression of Cost Parameters on Store and Location Characteristics

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

(2) Robust standard errors in parentheses

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

Several factors may contribute to the residual term in the regressions in Table 8, including: sampling error in the estimates of structural parameters γ , and in the own regression in Table 8; store and location heterogeneity that is unobservable to us as researchers; and store managers' idiosyncratic perception of the cost. We interpret the substantial unexplained variation in these regressions as indicative that part of the heterogeneity comes from store managers' idiosyncrasy. For the counterfactual experiments that we present in Section 5, we interpret these residuals as coming from store managers' idiosyncratic perception of costs.

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 of 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 have 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 Table 8, we decompose this vector into two additive and orthogonal components: the part explained by the regressors, that we represent as $\hat{\gamma}_{i,j}^{exp}$; and the residual part, $\hat{\gamma}_{i,j}^{res}$. We interpret this residual component as idyosyncratic from the store manager, and construct a counterfactual scenario that removes this component $\hat{\gamma}_{i,j}^{res}$ 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 residuals component of the four cost parameters. This implies a total of 15, 850 experiments.

We implement each of these experiments 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.¹⁸ 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 the 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 after an order (i.e. *S* threshold), and inventory to sales ratio before an order (i.e. *s* threshold).

Figures 9 (for stockout frequency), 10 (for ordering frequency), and 11 (for inventory-to-sales ratio) summarize the results from these experiments. In each figure, the horizontal axis measures the value of the corresponding parameter $\gamma_{i,j}^{res}$, and the vertical axis measures the difference in the mean value of the outcome variable between the factual and the counterfactual scenario.

¹⁸Note 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.

For instance, in Figure 9(a), the horizontal axis represents $\gamma_{i,j}^{h,res}$, and the vertical axis measures $\Delta SOF_{i,j} = SOF_{i,j}^{factual} - SOF_{i,j}^{counter}$, where SOF denotes the stockout frequency.

Figure 9: Counterfactual Outcome: Stockout Frequency



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





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





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

Note that the counterfactual experiment of shutting down $\gamma_{i,j}^{h,res}$ to zero is equivalent to a change in parameter $\gamma_{i,j}^{h}$ from the counterfactual value $\gamma_{i,j}^{h,exp}$ to the factual value $\gamma_{i,j}^{h,exp} + \gamma_{i,j}^{h,res}$. Therefore, we can see the cloud of points in, say, Figure 9(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.

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 9 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. The sign of the effects of these parameters on the stockout frequency is exactly the opposite. Panels (a) to (d) in Figure 9 confirm the signs of these effects.

In Figure 10, 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 , and therefore the sign of the effects of this parameters on ordering frequency is exactly the opposite. Panels (a) and (c) confirm the sign of these effects. According to Blinder formula, the sign of effect 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.

Figure 11 illustrates the relationship between cost parameters and the inventory to sales ratio. Blinder's formula establishes that the two threshold S and s depend negatively on γ^h and positively γ^z . The sign of the effects of these parameters on the inventory sales ratio is the same as for the two thresholds. Panels (a) and (b) confirm these signs. According to Blinder's formula, the sign of the effects of γ^f and γ^c on the inventory 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	0.0004	0.0004
	(0.0005)	(0.0004)	(0.0005)	(0.0005)	(0.0005)	(0.0004)
Ordering Frequence	0.1826	0 1829	0 1826	0.1745	0.1751	0 1626
Ordering Prequency	0.1650	(0.1632)	(0.1030)	(0.11740)	(0.1404)	(0.1030)
	(0.1444)	(0.1464)	(0.1444)	(0.1372)	(0.1424)	(0.1336)
Inventory to Sales Batio	26.6300	23.2824	26.6309	28.0190	25.0897	21.7941
	(21.8715)	(14.8756)	(21.8722)	(28.3123)	(19.0763)	(12.7583)
		()	· · · ·	()	()	· /
Inventory to Sales Ratio After Order	31.4581	26.7637	31.4593	34.1110	28.1400	23.7675
	(30.2815)	(19.5448)	(30.2824)	(41.1025)	(23.5807)	(14.9385)
Inventory to Sales Ratio Before Order	21.6489	17.7833	21.6499	24.2655	19.0300	15.6574
	(21.2277)	(12.8067)	(21.2288)	(31.1512)	(15.6194)	(8.6556)
		2 2205	0.0010	1 0 0 0 0	0.0011	10 51 50
Change in Total Inventory Cost (%)	—	-3.6285	0.0013	1.2608	-6.9611	-12.5173
		(12.0103)	(0.0119)	(30.6882)	(13.0466)	(18.6898)

Table 9: Counterfactual Experiments Removing Residual Component of Cost Parameters

It is of interest to consider what is the average effect across stores and products of shutting down the store manager specific component in costs. Table 9 presents these average effects for each of the four cost parameters and for the combination of the four.¹⁹ Removing the manager component in all four inventory costs generates: a decrease in the mean ordering frequency of 2 percentage points, from 18.4% to 16.4%; a decrease in inventory to sales ratio of 4.8 days of average sales, from 26.6 to 21.8 days; a decrease in the lower *s* threshold of 6 days, from 21.6 to 15.6 days; and a decrease in the S - s gap of 1.6 days, from 9.8 to 8.2 days.

Store managers' idiosyncratic perception of costs has a very substantial effect on inventory management at the aggregate level of the firm, with a 6-day increase in waiting time between two orders, a decrease in the average order amount equivalent to 1.6 days of average sales, a 22% decrease in the inventory to sales ratio, but a negligible effect on the frequency of stockouts. Accordingly, if the idiosyncratic component of costs is a biased perception by store managers,

¹⁹Note that, by construction, the residual component $\gamma_{i,j}^{res}$ has mean zero and is orthogonal to the explained component $\gamma_{i,j}^{exp}$. 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.

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 9 presents the effect of removing $\gamma_{i,j}^{res}$ on total inventory management cost calculated using $\gamma_{i,j}^{exp}$ but not $\gamma_{i,j}^{res}$. We find that, on average, this cost declines by 12.5%. 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's information about storelevel demand, inventories, and cost parameters. The experiments that we present below are based on the following assumptions.

- First, we assume that there is an information delay from stores to headquarters. Storeproduct level information regarding sales is processed by headquarters with a weekly delay. This assumption relies on the institutional details we describe in Section 2.1.
- 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 $\gamma_{i,j}^{exp}$ 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 $\gamma_{i,j}^{res}$ as store managers' behavioral biases and not as "true" costs. Under the decentralized system, store managers make decisions as if the cost parameters were $\gamma_{i,j}^{exp} + \gamma_{i,j}^{res}$, but our measure of their profits is based only on $\gamma_{i,j}^{exp}$. The evaluation of profits under this assumption provides an upper bound for the (profit) gains from centralization. Alternatively, we could assume that $\gamma_{i,j}^{res}$ 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 trade-off 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 $\gamma_{i,j}^{res}$. 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.

Table 10: Decentralized vs. Centralized Profits: Average Daily Profit Per-Store Per-Product (in CAD)

	Mean	Pct. 10%	Pct. 25%	Median	Pct. 75%	Pct. 90%
Centralized Solution (\$)	54.93	9.75	18.28	43.91	82.21	116.13
Decentralized Solution (\$)	54.64	9.75	18.15	44.01	81.92	116.35
Gains in Profit from Decentralization (\$)	-0.29	-2.29	-1.01	-0.20	0.26	1.45
Gains in Profit from Decentralization $(\%)$	-0.41	-3.52	-2.26	-0.73	1.02	3.27
Change in Inventory Cost from Decentralization $(\%)$	25.45	-5.52	0.10	6.30	23.01	56.97

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

Figure 12: Change in Daily Profit From Decentralization (\$)



We depict the results of this experiment in Table 10, and Figures 12 and 13. 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 10 presents means, medians, and several percentiles for the profit per store-product under the centralized and decentralized systems and for the gains from decentralization. To have a better perspective of the implications

Figure 13: Change in Daily Inventory Cost From Decentralization (%)



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 2012.²⁰ At the aggregate level, decentralization has a negative impact on LCBO profits. It implies a 0.4% decline in profits, which represents \$6.8 million in annual profits for the retail chain. The effect on the median store is also negative: -0.7%. This 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 10: on average, decentralization increases total inventory costs by 25.5%.

The evidence on the mean and median effects in Table 10 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 10 and Figure 12 show large heterogeneity in the impact of decentralization, with both very substantial losses and gains: percentiles 10 and 90 are -3.5% and 3.3%, respectively. Therefore, although centralization would generate positive gains in total profits for the retail chain relative to the existing decentralized structure, the distributional effects of these gains are very substantial.

 $^{^{20}\}mathrm{According}$ 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.

Figure 13 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., 6.3%), but most striking is the long right tail of this distribution, that implies a 25.5% increase in average inventory costs.

6 Conclusion

Retail chains are complex organizations with multiple divisions and teams, each with its own decision rights. Store managers play an important role in some retail chains. A store manager can collect and process information just-in-time about her own store. This information at the store level is more homogeneous and manageable than at the level of the whole chain. The transfer of this information from stores to headquarters, and especially the processing of information at the headquarters, can generate delays of days or even weeks, which can have an important negative impact on decision making and profits. On the other hand, store managers can have heterogeneous skills, motivations, and can put different levels of effort. A centralized decision system that uses only the best managers in the firm can minimize the negative effects of store managers' heterogeneous skills.

In this paper, we investigate this trade-off between centralized and decentralized decision making in the context of inventory management in a large retail chain. Using a unique dataset with daily information on inventories, sales, prices, and stockouts at the level of individual stores and products, we estimate a dynamic structural model of store managers' inventory decisions. Based on the principle of revealed preference, we obtain separate estimates for every store and product of four cost parameters: per unit inventory holding cost, stockout cost, fixed ordering costs, and per unit ordering costs. We find very substantial heterogeneity across stores in these cost parameters. Part of this heterogeneity is explained by observable store and location characteristics, but the residual heterogeneity is substantial and we attribute it to store managers' idiosyncratic information/perceptions.

We use the estimated model to implement a counterfactual experiment that evaluates the effect of centralizing inventory management at LCBO. For this counterfactual experiment, we assume that the store managers' idiosyncratic component of the costs is a behavioral bias and does not represent true cost. This working assumption provides an upper bound to the gains from centralization. We find that a centralized inventory management system would yield a 0.4% increase in annual profit for LCBO. This modest effect is the result of combining two substantial effects with opposite signs. The negative impact on profits of losing the *just-in-time* information

from store managers is more than compensated by the large reduction in ordering and storage costs from eliminating store managers' behavioral biases and heterogeneous skills (i.e., 25.5% on average, 6.3% for the median store). Furthermore, the gains/losses from centralization are very heterogeneous across stores in the retail chain, including an important fraction of stores with substantial losses from adopting centralization. This distributional effect may have implications in the decision of adopting an organizational structure that – all other things being equal – would increase the company's overall profit. Distributional conflict within multi-divisional companies and its relation to organizational structure has been studied in theoretical models by Inderst, Müller, and Wärneryd (2007), and it is an interesting topic for further empirical research.

References

- ADDA, J., AND R. COOPER (2000): "Balladurette and Juppette: A discrete analysis of scrapping subsidies," *Journal of Political Economy*, 108(4), 778–806.
- AGUIRREGABIRIA, V. (1999): "The dynamics of markups and inventories in retailing firms," The Review of Economic Studies, 66(2), 275–308.
- AGUIRREGABIRIA, V., D. ERSHOV, AND J. SUZUKI (2016): "Estimating the Effects of Deregulation in the Ontario Wine Retail Market," *Working Paper*.
- AGUIRREGABIRIA, V., AND A. MAGESAN (2020): "Identification and estimation of dynamic games when players' beliefs are not in equilibrium," *The Review of Economic Studies*, 87(2), 582–625.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): "Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models," *Econometrica*, 70(4), 1519–1543.
- ALONSO, R., W. DESSEIN, AND N. MATOUSCHEK (2008): "When does coordination require centralization?," *American Economic Review*, 98(1), 145–179.
- ARROW, J., T. HARRIS, AND J. MARSCHAK (1951): "Optimal Inventory Policy," *Econometrica*, 19(3), 250–272.
- ARTHUR, D., AND S. VASSILVITSKII (2007): "k-means++: The advantages of careful seeding," Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA), p. 1027–1035.
- ATTANASIO, O. (2000): "Consumer durables and inertial behaviour: Estimation and aggregation of (S, s) rules for automobile purchases," *The Review of Economic Studies*, 67(4), 667–696.
- BERTOLA, G., L. GUISO, AND L. PISTAFERRI (2005): "Uncertainty and Consumer Durables Adjustment," *Review of Economic Studies*, 72(4), 973–1007.
- BLINDER, A. (1981): "Retail Inventory Behavior and Business Fluctuations," *Brooking Papers* of Economic Activity, 2(1), 443–505.
- BRAY, R., AND I. STAMATOPOULOS (2021): "Menu Costs and the Bullwhip Effect: Supply Chain Implications of Dynamic Pricing," *Operations Research*.
- BRAY, R., R. YAO, Y. DUAN, AND J. HUO (2017): "Ration Gaming and the Bullwhip Effect," *Operations Research.*

- BRAY, R., Y. YAO, Y. DUAN, AND J. HUO (2019): "Ration gaming and the bullwhip effect," Operations Research, 67(2), 453–467.
- CHO, S., AND J. RUST (2010): "The flat rental puzzle," *The Review of Economic Studies*, 77(2), 560–594.
- DELLAVIGNA, S., AND M. GENTZKOW (2019): "Uniform Pricing in US Retail Chains," *NBER* Working Paper 23996.
- DENARDO, E. (1981): Dynamic Programming: Models and Applications. Prentice Hall PTR, USA.
- DESSEIN, W., L. GARICANO, AND R. GERTNER (2010): "Organizing for synergies," American Economic Journal: Microeconomics, 2(4), 77–114.
- DORASZELSKI, U., G. LEWIS, AND A. PAKES (2018): "Just starting out: Learning and equilibrium in a new market," *American Economic Review*, 108(3), 565–615.
- EBERLY, J. (1994): "Adjustment of consumers' durables stocks: Evidence from automobile purchases," *Journal of Political Economy*, 102(3), 403–436.
- ELLISON, S., C. SNYDER, AND H. ZHANG (2018): "Costs of Managerial Attention and Activity as a Source of Sticky Prices: Structural Estimates from an Online Market," *NBER Working Paper 24680.*
- FOOTE, C., E. HURST, AND J. LEAHY (2000): "Testing the (S,s) Model," American Economic Review, 90(2), 116–119.
- GOLDFARB, A., AND M. XIAO (2011): "Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets," *American Economic Review*, 101(7), 3130– 61.

GU, J., AND R. KOENKER (2017): "Empirical Bayesball remixed: Empirical Bayes methods for longitudinal data," *Journal of Applied Econometrics*, 32(3), 575–599.

HADLEY, G., AND T. WHITIN (1963): Analysis of inventory systems. Prentice-Hall.

^{— (2019): &}quot;Transitory shocks, limited attention, and a firm's decision to exit," *Working Paper*.

- HALL, G., AND J. RUST (2000): "An empirical model of inventory investment by durable commodity intermediaries," *Carnegie-Rochester Conference Series on Public Policy*, 52, 171– 214.
- HART, O., AND B. HOLMSTROM (2010): "A theory of firm scope," The Quarterly Journal of Economics, 125(2), 483–513.
- HITSCH, G., A. HORTAÇSU, AND X. LIN (2021): "Prices and promotions in US retail markets," *Quantitative Marketing and Economics*, pp. 1–80.
- HORTAÇSU, A., F. LUCO, S. PULLER, AND D. ZHU (2019): "Does strategic ability affect efficiency? Evidence from electricity markets," *American Economic Review*, 109(12), 4302– 42.
- HORTAÇSU, A., O. NATAN, H. PARSLEY, T. SCHWIEG, AND K. WILLIAMS (2021): "Organizational structure and pricing: Evidence from a large us airline," *National Bureau of Economic Research, Working Paper*, 29508.
- HORTAÇSU, A., AND S. PULLER (2008): "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.
- HUANG, Y., P. ELLICKSON, AND M. LOVETT (2020): "Learning to set prices," Available at SSRN 3267701.
- INDERST, R., H. M. MÜLLER, AND K. WÄRNERYD (2007): "Distributional conflict in organizations," *European Economic Review*, 51(2), 385–402.
- KRYVTSOV, O., AND V. MIDRIGAN (2013): "Inventories, markups, and real rigidities in menu cost models," *Review of Economic Studies*, 80(1), 249–276.
- LIQUOR CONTROL BOARD OF ONTARIO (2012): "Annual Report 2011-2012," https://www.lcbo.com/content/dam/lcbo/corporate-pages/about/pdf/2011_2012.pdf.
- (2013): "Annual Report 2012-2013," https://www.lcbo.com/content/dam/lcbo/ corporate-pages/about/pdf/LCB0_AR12-13-english.pdf.
- ——— (2016): "An Introduction to Web Based Ordering & SOS (Store Order Services)," LCBO PSM Trade Symposium.
- MARSCHAK, J., AND R. RADNER (1972): Economic Theory of Teams. Yale University Press.

- MIRAVETE, E., K. SEIM, AND J. THURK (2020): "One markup to rule them all: Taxation by liquor pricing regulation," *American Economic Journal: Microeconomics*, 12(1), 1–41.
- MOSSER, P. (1991): "Trade inventories and (S, s)," *The Quarterly Journal of Economics*, 106(4), 1267–1286.
- PUTERMAN, M. (2014): Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons.
- RUST, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.
- (1994): "Structural estimation of Markov decision processes," *Handbook of Econometrics*, 4, 3081–3143.
- SCARF, H. (1959): "The optimality of (S,s) policies in the dynamic inventory problem," in Mathematical Methods in the Social Sciences, ed. by K. Arrow, S. Karlin, and P. Suppes, chap. 13. Stanford University Press, Stanford.

7 Appendix

7.1 Estimates of Sales Forecasting Equation

Table 11 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 11 reports the three quartiles in the distribution across stores of parameter estimates and standard errors for the coefficient of logprice, 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 around 5% 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.

	Distr	ibution of Esti	mates
	1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile
	Est.	$\frac{-}{Est.}$	$\frac{1}{Est.}$
	(st.err.)	(st.err.)	(st.err.)
SKU #67 – Smirnoff Vodka		()	()
$\ln(p)$	-2.9544	-1.5502	0.0481
	(1.3823)	(1.7332)	(2.3553)
$\ln(Q)$	0.3725	0.5275	0.6934
	(0.0687)	(0.0949)	(0.1610)
lpha	0.3302	0.4241	0.6199
	(0.0363)	(0.0559)	(0.1026)
$Pseudo-R^2$	0.0575	0.0744	0.0874
SKU $\#117$ – Bacardi Superior White Rum			
$\ln(p)$	-6.8605	-4.2215	-2.0114
	(2.0144)	(2.5312)	(3.5763)
$\ln(Q)$	0.3416	0.4848	0.6317
	(0.0785)	(0.1133)	(0.1987)
lpha	0.2859	0.3571	0.5341
	(0.0364)	(0.0578)	(0.1187)
Pseudo-R ²	0.0562	0.0725	0.0850
SKU $#340380$ – Two Oceans Sauvignon Blanc			
$\ln(p)$	-6.4681	-5.2248	-3.7396
	(0.8554)	(1.0824)	(1.4750)
$\ln(Q)$	0.1614	0.3231	0.5272
	(0.0846)	(0.1227)	(0.2185)
lpha	0.4444	0.7146	1.3043
	(0.0465)	(0.0824)	(0.1840)
Pseudo-R ²	0.0387	0.0522	0.0677
SKU $\#550715$ – Forty Creek Select Whisky			
$\ln(p)$	-5.5620	-3.2317	-0.7215
	(1.9387)	(2.4236)	(3.2657)
$\ln(Q)$	0.0942	0.2564	0.3991
	(0.0898)	(0.1298)	(0.2040)
lpha	0.1888	0.2546	0.3866
	(0.0338)	(0.0533)	(0.0977)
Pseudo-R ⁻	0.0450	0.0653	0.0794
SKU $\#$ 024544 – Yellow Tall Shiraz Red	F 0079	9 7440	0.7154
$\ln(p)$	-3.09(3)	-3.(449)	-2.(154)
$\ln(\Omega)$	(1.0007)	(1.3303)	(1.9912)
$\operatorname{III}(Q)$	(0.2343)	0.401(0.0000 (0.1009)
	0.0072	(0.0929) 0.6133	(0.1003 <i>)</i> 1 0/03
ά	(0.4100)	(0.0133	(0.1493)
$\mathbf{P}_{sourdo} \mathbf{R}^2$	$ \begin{array}{c} (0.0335) \\ 0.0325 \end{array} $	0.0300)	(0.1403) 0.0577
i seudo-49	0.0020	0.0404	0.0011

Table 11:	Sales Fo	recasting	Equation	– Negative	Binomial	Model