

Learning to Set Prices

Yufeng Huang, Paul B. Ellickson and Mitchell J. Lovett*

January 10, 2020

Preliminary and comments welcome!

Abstract

Do entrants in new markets learn to set optimal prices based on demand conditions? Focusing on the privatized Washington State liquor market, this paper presents novel evidence of learning to set prices, and investigates the significance and nature of such learning. We first descriptively show that entrants learn from realized sales quantity, and learn about customer preferences in the new market. Next, we structurally estimate demand and wholesale prices, combining aggregate data on quantity and price, microdata on customer compositions, and cost data from liquor-control states. Our structural estimates suggest that pricing mistakes are costly but transitory: initial prices lead to a 9% loss in profit compared to optimal prices, but retailers learn to correct these mistakes and quickly converge to close to optimality. In addition, initial prices resemble retailers' misinformation about Washington's unique customer composition. Nevertheless, prior experiences in different liquor markets help retailers set better initial prices, implying that experiences learned from one market is transferable to other (even different) markets. Broadly, in contrast to anecdotes and conjectures that pricing decisions in practice are set by costs and heuristics, our finding suggests that managers learn from experiences to develop a nuanced, sophisticated understanding of demand.

Keywords: Pricing strategies; Learning; New entrant behavior; Liquor markets

*University of Rochester Simon Business School. We gratefully acknowledge numerous comments and suggestions from Kristina Brecko, Christopher Conlon, Ulrich Doraszelski, Ronald Goettler, Brett Gordon, Avery Haviv, Jihye Jeon, Przemyslaw Jeziorski, Jeanine Miklos-Thal, Kanishka Misra, Sarah Moshary, Harikesh Nair, Xiliang Lin, Marc Rysman, Greg Shaffer, Kosuke Uetake, Thomas Wollman, Chenyu Yang, and Hongsong Zhang, as well as seminar and conference participants at Northwestern University, 2018 NYU IO Day, 2017 SHUFE IO Conference, Stanford University, 2018 SICS Conference, 2019 QME Conference, University of Chicago, and University of Rochester. We thank Zhe Hong for excellent research assistance. We thank the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business who provide access to data owned by The Nielsen Company (US), LLC. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Nielsen retains copyright of the data: Copyright © 2018 The Nielsen Company (US), LLC. All Rights Reserved.

1 Introduction

“Pricing is a big question mark, for everyone entering the spirits business in Washington... I sure don’t know what we’ll charge the consumer. There is going to be a lot of scrambling...”

– Alan Johnson, CEO of BevMo!¹

Pricing is a key decision firms make when adapting to new market conditions. Setting new prices when managers face significant demand uncertainty can be quite challenging. To meet this challenge, quantitative marketing has a long tradition of providing normative guidance for value-based pricing strategies (Reibstein and Gatignon, 1984; Chintagunta et al., 2003; Ellickson et al., 2019). Yet, these optimal pricing strategies seem to be at odds with actual practices. Recent empirical evidence documents pricing anomalies that depart from canonical optimal-pricing models (DellaVigna and Gentzkow, 2019), and, in particular, managers do not respond to substantial market changes (Arcidiacono et al., 2019). Further, numerous interviews and anecdotes find that managers often describe their pricing practices as based on routines and heuristics (Noble and Gruca, 1999). What explains the gap between models and practices? One possible reconciliation is that managers do not explicitly adopt the formal language and procedures of quantitative models, yet implicitly adapt their routines and heuristics over time to approach optimal prices.

In this paper, we examine this hypothesis by studying whether and how entrants learn to set prices in a new market. We present new empirical evidence that entrants in a new market start far from optimal decisions but learn about demand and gradually adjust towards fully-optimizing behavior. We show that the key mistakes are associated with systematic customer-segmentation differences in the new market vis-a-vis other states, and the process of price adjustments is consistent with learning about these differences. We demonstrate that prices gradually stabilize to a point close to full-information optimizing behavior, suggesting that managers are not locked into suboptimal heuristics or routines, but can improve their pricing practices by learning from information

¹Source: Paul Gregutt, “BevMo! Ramps It Up in Washington State”. From www.paulgregutt.com.

obtained from market experiences.

We study the newly-privatized Washington State liquor market, where existing retail outlets are allowed for the first time to sell liquor, and we leverage two aspects of this setting to overcome key empirical challenges. First, we need to isolate learning about demand from a myriad of other factors that arise when firms enter a new market. In the Washington liquor market, we demonstrate that customers have stable liquor preferences, liquor purchases do not drive store visits, and retail competition over liquor purchases is negligible. These aspects of the institutional setting allow us to focus on retailers learning about consumer demand for products. Second, the researcher typically does not observe costs, making it difficult to directly measure profits. In the Washington market, we observe wholesale price data both from before market privatization and within the sample period in a neighboring, liquor-control state that help to provide measures of profitability. Together, these features allow us to study the pricing behavior and its impact on firm profitability, while carefully isolating learning from confounding factors.

We begin by documenting sizable and heterogeneous price movements in the first two years after the privatization of liquor sales that suggest retailers indeed adapt to the new market conditions. In the first year, the median price drops by around 10% and remains stable afterward. Importantly, these price changes differ systematically by product category. We demonstrate dramatic price changes for bourbon and rye, the category where Washington customers differ the most from other states' customers. These price movements are consistent with retailers learning about demand, and in particular, Washington's distinct composition of customer segments by product category.

We present novel descriptive evidence supportive of firm learning. First, we show that retail prices for a product respond to lagged demand shocks for the same product and that the rate of response to these shocks declines over time. This pattern is consistent with retailers learning *from* the new information contained in sales – i.e., they respond to sales shocks until they accumulate enough experience to know demand well. Second, novel to the literature, we show that, across products, the correlation between observed prices and average quantities first increases and then stabilizes over time. This finding suggests that retailers learn *about* the customer types that prefer

different products: i.e., they are able to identify products that customers are willing to pay more (less) for and set higher (lower) prices correspondingly. These analyses also demonstrate that most of the learning occurs early in the sample. Later in the sample, firm strategies do not systematically change as a result of learning.

Next, we investigate the breadth of initial mistakes, the nature of these mistakes, and the process by which firms correct them. To do so, we estimate a structural model of demand and cost primitives, imposing minimal ex-ante assumptions on the optimality of retailers' pricing decisions. We estimate a random coefficient demand model (Berry et al., 1995) incorporating standard aggregate- and micro-level moments. The resulting parameter estimates and output accord well with existing evidence from related liquor markets (Conlon and Rao, 2015; Miravete et al., 2018). Following our descriptive evidence, we estimate firms' marginal costs based on the assumption that firms have reached full information in the last half-year of the sample. To validate this assumption, we compare our estimated retailer costs for the beginning of the privatization period with the relevant observed wholesale prices. We show that these cost estimates and wholesale prices correspond closely, indicating that our assumption of optimal prices at the end of the observation period is valid. Importantly, this analysis also demonstrates that decisions become optimal (or very nearly optimal) within a few years. Hence, in this setting, if suboptimal managerial heuristics do exist, they are only transitory and diminish quickly through learning from experience.

We then use the model to conduct a series of counterfactual analyses that produce three additional key findings. In these exercises, we compare actual prices against the model-implied optimal prices for the earlier privatization period. First, we establish the size of initial pricing mistakes and the speed of retailer adaptation. In the median, initial prices are 9% too high compared to the optimum, and this gap leads to a 9% loss of variable (gross) profits. To improve profits, retailers would have lowered prices in exchange for higher sales volume. However, despite these early mistakes, retailers learn quickly: they lower prices and recover one-third of the profit gap within a quarter and another one-third within the next six quarters. Consistent with our descriptive evidence on prices responding to new information, high-volume product prices adapt more rapidly than low-

volume product prices, indicating that higher information content leads to more rapid adjustment. Hence, although the scope of pricing mistakes at the start of the market is sizable, retailers adapt quickly to correct these mistakes based on informative experiences.

Second, we consider the nature of the pricing mistakes. As demonstrated by our demand estimates, the customer composition (by product category) in the Washington market is distinct from other states. The initial pricing mistakes closely match what these nuanced differences in customer composition would predict. In particular, retailers price categories too closely to the demand conditions in existing markets, which indicates that pricing is responsive to demand conditions, and contradicts the idea that managers set prices based on simplistic heuristics that ignore demand variations (Noble and Gruca, 1999; DellaVigna and Gentzkow, 2019). Yet, these initial mistakes due to misinformation about the market conditions suggest that managerial practices allow retailers to learn from experience about nuanced customer segmentation distinctions in a new market. This finding implies sophisticated pricing behavior beyond what managers can or do describe in interviews.

Third, we investigate the role of prior experience on initial pricing decisions. For local retailers that only operate in Washington or Oregon, the privatized Washington market is the first market where they sell liquor. In contrast, national chains previously sold liquor in other states and therefore have some knowledge about the demand for liquor. Importantly, in our setting, this prior experience comes from markets providing misinformation about the customer composition in Washington. If firms accumulated knowledge in a relatively superficial manner, then the experience in previous markets would lead the national retailers to be worse off. On the contrary, we find that the experienced retailer does better. Local retailers make sizable initial mistakes and incur 12-15% profit loss relative to the optimum, whereas the experienced retailer prices at only 6% below the optimal profit. Thus, we find that experience does not lead to a superficial understanding of demand, nor does it lock the firm into prices that serve their existing customer segments at the expense of new markets. Instead, our results imply that learning to set prices must involve a more fundamental reckoning with the underlying determinants of demand that allow accumulated

knowledge to transfer between markets with different characteristics. Again, this finding suggests firms learn from experience to develop a nuanced, sophisticated understanding of demand.

Contribution and related literature. The most direct contribution is to the firm learning literature. We present new evidence of firm learning and what firms learn about. We show that retailer strategies are at first suboptimal and that they adapt to informative experiences (demand shocks) by adjusting prices to increasingly capture demand. The closest related paper is Doraszelski et al. (2018), who show that, after the opening of the frequency response market (in the UK electricity network), prices appear random at first and evolve to patterns that can be rationalized as a stationary equilibrium. They focus on the adaptation of equilibrium play whereas our paper focuses on firm learning about demand. Our paper is also closely related to Jeon (2017) and Covert (2015), who study firm learning and investment decisions in, respectively, the container shipping industry and the hydraulic fracturing industry. Our paper complements these two papers in that they structurally estimate how firms' investment decisions react to demand signals whereas we focus on empirically describing how pricing decisions adapt toward optimality. Also related to our paper, Hitsch (2006) and Dixit and Chintagunta (2007) structurally estimate learning about specific aspects of demand and make optimal exit decisions. In contrast to these papers that measure the information value of quantity signals under specific model assumptions, our context allows us to better isolate what firms learn about (the composition of customer segments that buy different products), and how experience plays a role in this learning process. In this sense, our work also relates to the literature applying experience or learning curves (Argote and Epple, 1990; Argote et al., 1990; Benkard, 2000) as a reflection of accumulation of human capital or improvements in organizational routines. We add to this literature new evidence that ties such learning more closely to the underlying information environment as well as new evidence about knowledge transfer.

Our results also speak to the broader literature on pricing practices. Earlier studies that use surveys and interviews to learn about the ways managers make pricing decisions (see, e.g., Kaplan et al., 1958 and Noble and Gruca, 1999) find that managers rely on suboptimal heuristics such as

“cost-plus” pricing, and suggest firms are bounded rational (Simon, 1955). Bounded rationality is also one explanation to some of the pricing anomalies documented by the recent empirical literature: For example, Arcidiacono et al. (2019) show lack of price responses to large demand shocks and DellaVigna and Gentzkow (2019) show lack of price variations across geographic markets. Similarly, Adams and Williams (2017) and Hitsch et al. (2019) document uniform or zone pricing, but interpret their findings either as firms’ rational responses to market conditions or as equilibrium strategies that soften competition. We show that, when entering a new market, firms make initial mistakes arising from the unique characteristics of that market, yet adapt prices as they incorporate nuanced information from experience, and ultimately converge to optimal pricing.

Broadly, this paper is related to the strategy literature on management effectiveness, firm heterogeneity, and dynamic capabilities. In a survey-based approach, Bloom et al. (2017), among others, show that differences in management practices explain a sizable fraction of the productivity differences across firms. Using a structural approach, Goldfarb and Xiao (2011) and Hortaçsu et al. (2019) show that there is heterogeneity in observed strategies that reflect potential limitations in managerial practices or decision-making. In contrast, Teece et al. (1997) focus on dynamic capabilities that the firm develops that allow it to respond to environmental change, including such capabilities focused on improving marketing decisions such as price (Day, 2011; Morgan et al., 2009). We provide an empirical demonstration that firms are able to learn from experience and that such learning allows the firm to enter new markets with an improved position (King and Tucci, 2002). This latter point also speaks to the boundary conditions of the notion that incumbents may be saddled with poor incentives that focus them on meeting existing customers’ needs to the detriment of serving new markets (Chandy and Tellis, 2000; Christensen, 2013).

Finally, our paper also contributes to the research and evidence in recent studies about the retail liquor industry. Seo (2016) studies consumer store choice before and after the privatization of Washington State liquor market and quantifies the welfare impact of one-stop shopping. Illanes and Moshary (2018) leverage Washington State’s 10,000 square-foot minimum required retail space for liquor vendors, as a regression-discontinuity design, and study the effect of en-

try on prices and product assortments. Conlon and Rao (2015), Aguirregabiria et al. (2016), and Miravete et al. (2018) study different aspects of liquor regulation policy. Conlon and Rao (2015) study the “post-and-hold” policy as a collusive instrument for the wholesalers and investigate the welfare improvements (and redistribution) of an alternative tax policy. Aguirregabiria et al. (2016) investigate counterfactual regulation, tax, or competition regimes in the Ontario wine market, highlighting the importance of spatial differentiation. Miravete et al. (2018) study the welfare impact of state-imposed constant retail markup and find that the single-markup policy decreases but also re-distributes consumer welfare.

The rest of the paper is organized as follows. Section 2 describes the context, data sources, and sample construction. Section 3 provides descriptive statistics about price movements in the privatized market, and key evidence of retailer learning. Section 4 then estimates a structural model and backs out demand and costs. After the structural estimation, Section 5 describes firms’ learning process from the structural estimates. Section 6 concludes the paper.

2 Context, Data, and Sample Construction

Before June 2012, in Washington State, a state-owned chain operated by the Washington State Liquor Control Board (WSLCB) had exclusive rights to off-premise liquor sales. The law imposed a fixed markup of 51.9% over the wholesale price. Starting on June 1, 2012, this state-owned chain was replaced by licensed, privately-owned retailers under Initiative 1183. This initiative required that retail licenses only be issued if the store has at least 10,000 square feet of floor size, and prioritized the issuance of retail liquor licenses for “existing grocery premises licensed to sell beer and/or wine” (See Article (3)(b) in I-1183). The initiative also mandated a retailer licensing fee of 17% of liquor revenue and an increase in retail off-premise sales taxes from 10% to 20.5%. In addition to these higher taxes, excise tax remained at \$3.77 per liter. These taxes and fees gave Washington state the highest retail liquor taxes in the nation.

Data. Our primary data source is the Nielsen Retail Measurement Services (RMS) Dataset. For our study, we focus on the liquor category in Washington State, and primarily the period from June 2012 (the month of privatization) until the last quarter of 2016.

We supplement the Nielsen data with auxiliary datasets. First, in some descriptive analyses, we compare Washington to states that have had as long as several decades of private retail grocery liquor sales, which serve as a form of placebo test. We identify fifteen such states (in descending order of total liquor sales volume): California, Arizona, Louisiana, Texas, New Mexico, Nevada, Nebraska, South Dakota, Colorado, Arkansas, Delaware, Maryland, North Dakota, and Washington, D.C. Second, we use public retail prices from Washington and Oregon to help estimate and validate wholesale prices in our structural analysis. Firms in Washington before privatization and Oregon throughout the sample period practiced fixed markup policies.² These policies allow us to calculate wholesale prices from the observed retail prices. Finally, we also use the Nielsen Consumer Panel Data between December 2009 and December 2016 to examine consumer behavior in greater detail and to construct micro-moments (Petrin, 2002) for use in the structural demand estimation.

Sample construction. We apply three filters in constructing our sample. First, we narrow our focus to a category within liquor, the broad whiskey category, that is sizeable, has little substitution with other categories (e.g., vodka), and contains the most product diversity, including whiskey, bourbon, scotch, and rye. This filter simplifies our analysis and allows us to focus on firm learning. This filter yields a dataset of 6,288,941 observations at the UPC-retailer-store-week level in Washington State, containing 724 unique UPCs (product name - size) and 635 unique product names.

Second, we also restrict our attention to the stable set of stores during our sample period to avoid store entry and exit unrelated to the liquor category. We select stores that sell a positive quantity in at least 95% of all weeks. This filter selects 561 out of 625 stores and removes 5.9% of

²For the state of Oregon, ORS 471.745 mandates a fixed markup. Various sources indicate that the markup is set at 106%. See Governor's Task Force, 2003, "Final Report on the Oregon Alcohol Beverage Industry" (page 25). Also see The Romain Group, LLC, 2014, "Written Comments on Draft Ballot Title for Initiative Petition No. 57" (page 2).

the observations from the overall sample.

Third, we filter products that enter and exit during the period to focus on “core assortments,” which contain the bulk of liquor revenues and profits. Within a given retailer, we select those products that first appear before December 2012, last appear after March 2016, and remain present for at least 25 weeks. This filter selects 276 out of 724 UPCs and eliminates newly introduced products, discontinued products, and products that are only occasionally or seasonally offered. Although it might appear that we eliminate many UPCs in this step, the products we drop out only account for 15.8% of the observations and 11.4% of the total revenue from the previous filter. We discuss the choice of focusing on these core products in more detail in Web Appendix B.1. After these sample selection steps, our sample contains 4,985,621 observations, from 276 products, six retailers, and 561 stores. Further, we focus our structural analysis on (1) the most popular size, 750ml bottles, (2) products with high enough sales to ensure precise market share measurements, and, for the supply side analysis, (3) the period of stable demand for the three largest retailers, which represent 80% of sales revenue. We discuss these considerations in detail in Web Appendix C.

Unit of analysis and categorization. Following the spirit of the analysis in DellaVigna and Gentzkow (2019) and Hitsch et al. (2019), we demonstrate that, empirically, retail prices are uniform across the state at the product level at a given point in time (see Web Appendix A). Motivated by this finding, we will focus on retailer decisions at the product level in Washington. We also categorize products in the broader whiskey category into four sub-categories representing important observable heterogeneity in products: (1) bourbon (and rye), (2) Canadian whiskey, (3) scotch (and Irish whiskey), and (4) “other” whiskey that do not fall into the first three categories.

3 Learning about Demand: Descriptive Evidence

In this section, we present descriptive evidence for firm learning, focusing on how prices vary over time and how they vary in ways that reflect information about demand. We begin by documenting

that while category-level prices initially increase too much over pre-privatization levels and gradually go down, product-level prices adjust in different directions by sizable amounts, and eventually stabilize years after the market opens. These patterns suggest that retailers might be learning about demand both at the category level and at the product level.

We then present two new pieces of descriptive evidence of learning. First, we show that prices adjust according to realized demand shocks when retailers start selling liquor in Washington, and that, over time, they adjust less and less to such shocks. Second, we show that the price of a product increasingly reflects its demand fundamentals (e.g., whether the product is highly popular or caters to a price-insensitive customer segment). This evidence suggests that retailers gain information about demand fundamentals, reflected in the prices they set. We also present placebo tests where we replicate these descriptive exercises in other states where retailers have sold liquor for a long time before 2012, and demonstrate that the evidence of learning does not appear in these “placebo” states.

Next, we discuss plausible features of demand that firms might learn about in the Washington market in Section 3.4. We focus on two plausible explanations that are consistent with observed changes in prices: Retailers might have to learn about the price sensitivities for liquor in Washington’s new, exorbitant tax regime, or learn about Washington’s distinct customer composition by product category. Finally, we address possible alternative explanations in Section 3.5 and in Web Appendix B.

3.1 Price changes after the privatization of liquor sales

We first examine how the whiskey and the whiskey sub-categories prices in Washington change over time, and whether these changes are consistent with learning about demand. We focus our analysis on the 750ml bottle size and 53 products that are available in Washington before and after the privatization as well as in other states. In Figure 1, we show that the pre-tax, volume-weighted average prices in other states generally increase over time, possibly as a result of inflation or cost changes. Washington pre-tax prices before the privatization are much lower than prices in other

liquor-privatized states, but their changes over time are similar. Upon privatization, the average price increases by 38%, from \$14.86 to \$20.49, to a level similar to the average pre-tax prices in other states. However, within the next two years, the average price level in Washington decreases and eventually resumes the same time trend as in other states.

[Insert Figure 1 about here]

Having established the overall pattern, we now examine how the price paths differ across categories. Such differences might be expected because of the differences in customer composition across categories. To control for (mild) cost changes we focus our analysis on the ratio of average price of each product (product-retailer-half year) for Washington versus other states, $\tilde{p}_{jrt} = \bar{p}_{jrt}^{\text{WA}} / \bar{p}_{jt}^{\text{other}}$, and normalize this ratio by its initial value, $(\tilde{p}_{jrt} - \tilde{p}_{jrt_0}) / \tilde{p}_{jrt_0}$.

We show in Figure 2 that bourbon prices experience significant declines after the privatization. The median price of this category declines by almost 10% by 2014, and prices of 25% products drop by at least 20%. We find similar patterns for scotch and Irish whiskey, although the magnitude is not as pronounced. In contrast, prices for Canadian whiskey and other whiskey remain relatively stable. We later show that Washington's customer composition differ from other states and that, in particular, the bourbon category disproportionally draw lower-income customers. Although not conclusive, the overall pattern and category differences are consistent with retailers that face initial demand uncertainty and adjust prices as they learn about demand, and in particular, the unique customer composition of categories in Washington state. We proceed with more-formal tests of learning.

[Insert Figure 2 about here]

3.2 Do prices adjust based on past quantity shocks?

If retailers learn about stable demand through the information contained in sales quantity shocks, they will initially adjust prices in the direction of demand shocks, but cease doing so after they have learned. To test these two hypotheses, we run a regression on retailer-product prices. We

aggregate time to the monthly level to reduce the effect of temporary price promotions. Denoting j as a product, r as a retailer and t as a month, we estimate a linear model of the current price on 1-month lagged quantity, controlling for current quantity and lagged prices:

$$\begin{aligned} \log(p_{jrt}) = & \beta_{\tau} \log(q_{jrt-1}) + \\ & \rho \log(p_{jrt-1}) + \alpha^{-1} \log(q_{jrt}) + \psi_{jr} + \phi_t + \eta_{jrt}. \end{aligned} \quad (1)$$

Our key parameter of interest is the sensitivity of the current price to q_{jrt-1} , the units sold for product j by retailer r in the previous month, $t - 1$. We allow β_{τ} to take a different value for each half-year. Learning would predict that these sensitivities should be initially positive, but then decline in magnitude over time. Our controls capture stable product-retailer characteristics (ψ_{jr}), common variations within a month (ϕ_t), serial correlation in prices (ρ), and the inverse demand (α^{-1}). For estimation, we take first differences and correct for the mechanical endogeneity bias in dynamic linear models, using Arellano and Bond (1991) instruments. To instrument the first lag of price, we use the third lagged price difference, $\Delta \log(p_{jrt-3})$, which guards against potential serial correlation in η_{jrt} , for instance, arising from retailers being unwilling to change price right after a previous price change due to menu costs.

Figure 3 plots the estimates of β_{τ} and corresponding confidence intervals. We find that in Washington, prices respond to the previous month's sales quantity positively and significantly in the early periods. We also find that the effect decreases over time until it is less than a third of the initial response. Hence, retail prices incorporate past sales shocks, but decreasingly so. This finding is consistent with retailers learning about demand by incorporating information contained in the realized sales quantities. This pattern for Washington State stands in sharp contrast to what we find for the states where retailers already had extensive experience in the liquor market. In these states, the responsiveness to sales shocks is much smaller, centered close to zero, and most coefficients are statistically insignificant. Hence, for retailers with greater experience, prices do not respond to past sales shocks. In Web Appendix Table 8, we present further details of the estimation

results that support the use of the instruments and controls.

[Insert Figure 3 about here]

3.3 Do prices increasingly reflect demand fundamentals?

If retailers start with incorrect estimates about demand but learn over time, these estimates, and the prices they set based on these estimates, should increasingly reflect (true) demand fundamentals. Prices for products that have surprisingly high demand will rise to reflect that higher demand, and prices for surprisingly low demand products will fall. The end result is that prices and the demand levels of products should become more highly correlated as retailers learn about demand. Hence, learning should be reflected in a pattern of increasing correlations between prices and the underlying demand for the product.

This intuitive argument can be illustrated with a simple model with zero marginal costs and profit maximization under linear demand,

$$q_{jt} = \beta_j + \alpha_j p_{jt} + \varepsilon_{jt}. \quad (2)$$

Assuming the ε_{jt} is mean zero, serially independent, and not known when setting price, the static optimal price is $p_{jt}^* = -\frac{\beta_j}{2\alpha_j}$. However, if retailers are uncertain about β_j and α_j and have noisy estimates of these demand primitives, then their static optimal price will be $\hat{p}_{jt} \neq p_{jt}^*$. Learning about the demand primitives from realizations of q_{jt} will lead to a sequence of \hat{p}_{jt} , that on average move toward p_{jt}^* . This movement will increase the correlation between price and the underlying demand primitives, β_j and α_j . As a result, the relationship between price and quantity will increasingly reflect this positive (endogenous) correlation.

These arguments motivate a new descriptive test of firm learning about demand. The idea is to estimate the cross-sectional relationship (i.e., ignoring the product fixed effects) between prices and sales quantity in each period, and evaluate how this estimated relationship changes over time. We expect the correlation between prices and quantities to reflect underlying product quality or

positioning (Trajtenberg, 1989; Berry, 1994), and as a result, cross-sectional variations in prices and quantity between products would estimate a price-sensitivity parameter that is biased upward. Instead of using instruments to recover demand, our test examines whether learning intensifies this endogeneity bias: If learning occurs, we would expect the correlation between quantity and price to increase over time, driven by the enlarging endogeneity bias. Otherwise, this correlation should be stable.

To implement this test, we regress quantity on price at the product-retailer-week level. We estimate via ordinary least squares (OLS) inserting retailer-week fixed effects $\bar{\beta}_{rt}$,

$$q_{jrt} = \bar{\beta}_{rt} + \alpha_{\tau} p_{jrt} + \eta_{jrt}. \quad (3)$$

The regression coefficient α_{τ} pools over half-year time windows (denoted τ) and products. Hence, this coefficient will both the true relationship and any bias due to the correlation between prices and the unobserved β_j and α_j . We report the estimated α_{τ} 's and their standard error for the Washington sample in Figure 4. We also examine retailers in other states as a placebo test, but with one change: we include retailer-state-week fixed effects.

In Washington, we find that retail prices are increasingly correlated with demand intercepts, reflected by the way $\hat{\alpha}_{\tau}$ increases between 2012 and 2015. If the underlying price sensitivities are stable over time, this finding suggests that retailers set prices with more and more information about demand, so that the price sensitivities are biased more and more towards zero. In addition, $\hat{\alpha}_{\tau}$ stabilizes in the second half of the sample, suggesting that most of learning occurs in the first two years after the market opens and there is no systematic learning that changes how prices reflect demand fundamentals.

In contrast to the finding in Washington, we find that other states' $\hat{\alpha}_{\tau}$ are stable over time. This finding is consistent with experienced retailers in the existing liquor markets do not further learn about demand in systematic ways. The placebo and main plots both support the conjecture that retailers learn about demand in the newly-privatized Washington market.

[Insert Figure 4 about here]

3.4 Discussion: What might retailers learn about?

What about the demand in Washington might retailers learn over time? Although we discuss this point in detail in Section 5.5, we offer two conjectures now. First, consumers in the post-privatization Washington market might be more sensitive to the pre-tax prices because they need to pay for much higher prices. The post-privatization Washington levies 20.5% sales tax and \$3.77/liter excise tax, which are exorbitant compared to the median sales tax at 6% and the median excise tax at \$0.80/liter, among other liquor-privatized states (see also Web Appendix Figure 5 for a comparison of tax regimes across states). If retailers do not fully account for potential price sensitivity differences, their initial prices might be too high for all products. The way average prices change after the privatization (Figure 1) is consistent with this conjecture.

Second, Washington consumers might have different tastes for individual products. Consistent with Figure 2, retailers who learn about these tastes over time might adjust prices in different directions for different products. In particular, we show using the Homescan data that Washington high-income consumers have distinct preferences over product categories (see Figure 5). Relative to high-income consumers in other states, high-income consumers in Washington purchase less bourbon (and rye), less scotch (and Irish whiskey), and much more Canadian whiskey. We speculate this pattern might be related to the proximity to Canada and distance from the main domestic liquor origins. This difference changes the composition of customers who purchase each category: As a result, bourbon and scotch/Irish whiskey face more low-income consumers, and Canadian whiskey faces more high-income consumers. Hence, relative to a generally-declining price trend, retailers who learn about category-level customer compositions would adjust bourbon and scotch/Irish whiskey prices downward, and Canadian whiskey prices upward.

[Insert Figure 5 about here]

3.5 Alternative explanations

The descriptive evidence suggesting learning about demand might also be consistent with alternative explanations: (1) changes in other retailer strategies, (2) changes in consumer demand, or (3) changes in (competitive) market structure. We provide evidence that these alternative explanations are not of first-order importance in our setting. This section outlines the main findings, and Web Appendix B presents further details.

First, it is important to consider whether other retailer strategies –in particular promotion and assortments– evolved over the sample period and contributed to the price changes. We show that both promotion frequency and depth are stable during the sample period (see Web Appendix B.1). Thus, we can approximate retailer decision-making as setting a shelf price (i.e., price including promotion). Further, as we document in Section 2, we focus on 276 "core" assortments in the empirical analysis, which are sold throughout the sample period and account for 89% of retailer revenue. Therefore, although retailers do adjust assortments over time, the changes pertain primarily to low-revenue products that have a limited influence on retailer pricing. Thus, we focus our inquiry on the pricing of core products.

Second, it is important to rule out systematic changes in consumer behavior post-privatization. If consumers gradually learn about the set of products or search for low-price products, retailers might adjust prices to capture such changes in consumer behavior. To address this possibility, we show in Figure 6 that total sales are stable over time in the sample period, consistent with the conjecture that demand does not change much in the sample. Later in the paper, we use our structural demand estimates to confirm further that consumer demand can be well-approximated by a stable function. We present details of this analysis in Web Appendix B.2.

[Insert Figure 6 about here]

Third, it is also important to rule out concerns over changes in market structure post-privatization. While Figure 6 suggests that we do not observe shifts in market positions between the set of retailers in our sample, we discuss two possible channels that could lead to instability of the market.

First, do retailers compete with each other in the liquor market? We show in Appendix Section B.3 that liquor customers are part of the grocery customer-base (instead of a stand-alone market), and that demand for a given product has negligible substitution between local retailers. These findings suggest that competition between chains is not prevalent in the liquor market. Second, does the upstream market change over time? Although we do not have data on wholesale contracts, industry publications reveal that two national wholesalers dominated the liquor wholesale market, and the market structure did not change with privatization. Further, wholesale prices from the Washington state-owned chain (reflected in their retail prices because of constant markups) indicates no discontinuity at the privatization of the wholesale market. This lack of change is consistent with the continuity in the channel and in sharp contrast with what happens in June 2012, when the retail market is privatized, leading immediately to a new channel of distribution.

4 Model: Demand and Costs

We characterize the primitives of demand and costs with the goal of computing counterfactual, full-information prices as a normative benchmark. Such a benchmark will allow us to investigate further the scope of learning and initial mistakes, and the probe into the nature of these mistakes as well as firms' processes of correcting them.

Demand for liquor is characterized by the random coefficient logit model, and estimated via standard nested-fixed point methods that incorporate micro-moments to identify heterogeneous demand parameters (Berry et al., 1995; Nevo, 2001; Petrin, 2002). We then estimate costs using the last six months of the sample (when we assume learning is complete) and present evidence that the model-estimated costs match well with costs data in Washington (pre-privatized) and Oregon.

4.1 The demand for liquor

Consumer i in market m in month t comes to retailer r to buy groceries, and derives utility from purchasing liquor j :

$$u_{ijrmt} = \gamma_{ki} + \alpha_i p_{jrmt} + x_{jrmt} \beta + \delta_{jrm} + \lambda_{rt} + \xi_{jrmt} + \varepsilon_{ijrmt}. \quad (4)$$

In the above, γ_{ki} represents category k and household i specific intercept, where k takes 1 (bourbon and rye), 2 (Canadian whiskey), or 3 (Irish whiskey and scotch). The parameter α_i represents household-specific price coefficients. These parameters capture heterogeneity in the tastes for liquor types and sensitivities to prices across households. x_{jrmt} are time-varying indicator variables for whether or not the product is on feature or display, or the feature/display status is unknown. δ_{jrm} are product-retailer and retailer-market fixed effects capturing tastes for different products, which could differ across shoppers going to different retailers or are in different markets. λ_{rt} are retailer-time fixed effects capturing changes over time in the demand for liquor in grocery stores, or changes in market positions for different retailers (in the liquor market). ξ_{jrmt} are unobserved characteristics or demand shocks. ε_{ijrmt} are type-1 extreme value utility shocks. If the consumer does not buy any liquor in the given trip, her utility is normalized to $u_{i0rmt} = \varepsilon_{i0rmt}$.

The consumer chooses among products in a given retailer, i.e., from the choice set J_{rmt} . This choice set only includes products from the focal retailer, assuming that the liquor category does not drive store choice, and hence, that the retailers act as monopolists over their own store traffic. As noted earlier, this assumption is motivated by multiple corroborating observations including that (1) consumers' grocery shopping patterns after liquor deregulation are stable, (2) when a consumer makes a liquor purchase, she typically spends more on other grocery items than liquor on the trip, and (3) sales of a product does not respond to variations in local availability of the same product. With these assumptions, the market share within a retailer-market is an integral of logit choice

probability over random coefficients

$$s_{jrmt} = \int s_{ijrmt} dF(\alpha_i) = \int \frac{\exp(\gamma_{ki} + \alpha_i p_{jrmt} + x_{jrmt} \beta + \delta_{jrm} + \lambda_{rt} + \xi_{jrmt})}{1 + \sum_{j' \in J_{rmt}} \exp(\gamma_{k'i} + \alpha_i p_{j'rmt} + x_{j'rmt} \beta + \delta_{j'rmt} + \lambda_{rt} + \xi_{j'rmt})} dF(\alpha_i, \gamma_i). \quad (5)$$

In Section 3.4, we show that Washington consumers potentially have different preferences towards the subcategories of liquor products compared to other states. Such differences might be an important driver of pricing mistakes. To capture this aspect in the model, we parameterize the household's preferences towards liquor subcategory k as a function of log household income,³

$$\gamma_{ki} = \gamma_k \log(y_i), \quad (6)$$

noting that the average category-level intercept is absorbed by product intercepts. We parameterize the price coefficient as a function of log household income and an independent normal random draw v_i ,

$$\alpha_i = \alpha_0 + \alpha_1 \log(y_i) + \alpha_v v_i. \quad (7)$$

4.2 Identification

Price coefficient. We estimate model parameters by a set of moments enforcing that demand shocks ξ_{jrmt} are conditional mean zero, given instruments z_{jrmt} (including non-price covariates, fixed effects, and excluded instruments for price and random coefficients):

$$\mathbb{E}[\xi_{jrmt} | z_{jrmt}] = 0. \quad (8)$$

One reason to instrument for prices is that, despite the inclusion of fixed effects, retailers might set prices based on private information not directly controlled for by these covariates: for example, the retailer might set higher prices for certain products in markets or time periods with high demand, in which case prices will be correlated with ξ_{jrmt} . Another reason is that the price data are averaged

³This income-sub category interaction term follows, but enriches, the specifications in Conlon and Rao (2015) and Miravete et al. (2018).

across weeks in a month and stores in a retailer/3-digit zipcode, potentially averaging across stores or weeks (even days) with different prices and creating an attenuation bias. To address these concerns, we construct price instruments similar to Conlon and Rao (2015) and Miravete et al. (2018). For each product carried by each retailer, we construct average prices of the product carried by all other retailers and across all states other than Washington, and use them as the instrument for retail prices in Washington. The prices in other states likely capture wholesale price variations that are common across states but do not correlate with demand shocks in Washington (after controlling for the above fixed effects) or have independent measurement errors. One example of wholesale price co-movement is that prices of Scotch whiskey move with the USD-GBP exchange rate, as illustrated in Figure 11 in the Web Appendix.

Random coefficients. We identify the random coefficients by combining the instruments suggested by Berry et al. (1995) with additional micro-moments (Petrin, 2002). First, we count the number of products available in each retailer-market-month. Variations in the market shares of the focal product in response to changes in the number of products can help to identify the non-proportionate substitution to other products versus to the outside option captured by category-intercepts γ_{ki} .

Second, we construct five sets of micro-moments, using the Nielsen Homescan panel data, to help identify how the category intercept and price coefficient vary with log income. Specifically, we divide annual household income (in \$1,000s) into three bins I_b : $[0, 42.5]$, $(42.5, 85]$, $(85, \infty)$. For Washington households in Homescan who visit the six focal retailers, these three income bins create three roughly equal-sized groups.

Next, for each income bin, we compute three moments: the average probability of buying liquor among retailer visits, the average price paid among liquor purchases, and the share of the three major categories of liquor among purchases. For each set of parameters θ , we match the theoretical moments to the sample analogs. The first set of moments match the average probability

of choosing the inside good:

$$\bar{s}_{rmt}^b = \frac{1}{N_b} \sum_{i \in I_b} \sum_{j \in J_{rmt}} s_{ijrmt}(\theta). \quad (9)$$

The second set of moments match the average price paid conditional on purchase of liquor:

$$\bar{p}_{rmt}^b = \frac{1}{N_b} \sum_{i \in I_b} \frac{\sum_{j \in J_{rmt}} p_{jrmt} \cdot s_{ijrmt}(\theta)}{\sum_{j \in J_{rmt}} s_{ijrmt}(\theta)}. \quad (10)$$

The third to fifth set of moments match, the share of bourbon (and rye), Canadian whiskey, and scotch (and Irish whiskey) among liquor purchases.

$$\bar{s}_{rmt}^{k,b} = \frac{1}{N_b} \sum_{i \in I_b} \frac{\sum_{j \in k} s_{ijrmt}(\theta)}{\sum_{j \in J_{rmt}} s_{ijrmt}(\theta)}. \quad (11)$$

In the above notation, N_b is the number of income draws falling into bin b , $k = 1, 2, 3$ is the product category, and $j \in k$ represents product j falling into category k . As in Petrin (2002), we match the observed overall and category-level purchase probabilities and purchase prices to the simulated ones, yielding our micro-moments.

4.3 Estimation results

Table 1 reports parameter estimates for the mean and standard deviation of price coefficients. We control for product-retailer, retailer-market, retailer-year and month fixed effects but these are not directly reported in the table. We also estimate the model without household-level coefficients and without micro-moments as in Berry (1994).

In the random coefficient logit model (“main spec.”), we find considerable heterogeneity in price sensitivities and in category utility (i.e. the intercept). The 5th percentile price sensitivity is -0.432 and the 95th percentile is -0.194 – the former is more than twice as large in magnitude as the latter. Part of this heterogeneity is driven by income, indicating that high-income consumers are less price-sensitive. Consistent with Conlon and Rao (2015), we also find that high-income consumers derive lower utility from the liquor category despite being less price sensitive.

[Insert Table 1 about here]

We measure the in-sample fit of the model using the R-squared for the mean utility projection, which is inverted from applying the Berry et al. (1995) contraction mapping on observed market shares given the nonlinear coefficients. We find that the model fits the data well, explaining 84% of the share variation. After estimation, we calculate the extent to which the estimated ξ_{jrmt} are auto-correlated. When imposing an AR(1) structure, i.e.

$$\xi_{jrmt} = \rho \xi_{jrmt-1} + \iota_{jrmt}, \quad (12)$$

we find that $\hat{\rho} = 0.651$ (standard error = 0.002), indicating the within shocks, ξ_{jrmt} , take a half-year to dissipate below 10% of their original influence.

Implied elasticities and markups. We next compute implied elasticities using our demand estimates. In table 2, we present the own- and cross-elasticity matrix for six products sold by retailer 32 in June, 2016. We find that elasticities are increasing in magnitude with price: the implied price elasticities of these example products range between -1.98 and -4.53.

Across all retailers and products, we find that the average elasticity at observed prices is -4.00. For products below a \$15 average price, the average elasticity is -2.59. For products above \$15, the elasticity is -4.79. These numbers are very close to Miravete et al. (2018) who find elasticities in the Pennsylvania liquor market to be -2.9 for cheap products and -4.9 for expensive products. Intuitively, this finding comes from the fact that consumers who are less sensitive to price (and thus are the main customers for high-end liquor) value the liquor category lower as a whole and thus have limited willingness to pay. Thus, retailers have lower percentage margins for high-price products.

Based on those elasticities, we compute the implied margin as a share of price if retailers set prices with full information about the demand parameters (the details of which will be shown in the next section). In the example below, for product 3, the retailer prices it at \$14.53 and gains 30.7% of the price as gross margin.

[Insert Table 2 about here]

4.4 Wholesale price (marginal costs for the retailers)

State-level uniform pricing with full information. In this section, we outline a price-setting model where the retailer sets prices for all of its products given full information on demand. In Section 3, we presented evidence that retailer's appear to stop responding to new information and adjusting prices to incorporate demand fundamentals by early 2015. We take the strategy of imposing full-information optimal pricing for this latter part of the data and then validating this assumption with a comparison to pre-privatization Washington wholesale prices. We estimate the marginal cost function for the period April to September 2016 and translate these cost functions back in time using Oregon wholesale price data. We then evaluate the quality of our wholesale price estimates for the initial period in Section 5.1 and find that they match well with external measures.

To recover the cost function, we assume that retailer r sets prices for its products with the restriction that the price must be uniform for each product across all markets in Washington. When setting prices, the retailer knows demand primitives $\mathcal{D}_r = \left(\{ \delta_{jrm}, \lambda_{rt} \}_{j,m,t}, \alpha, \beta, \gamma \right)$, but not the realized demand shock, $\xi_{jrm t}$. Instead, they take the projected demand shock from its one-month lagged value, $\hat{\xi}_{jrm t} = \hat{\rho} \xi_{jrm t-1}$.⁴

Letting $\tilde{s}_{jrm t}$ denote the implied market shares when demand shock innovations are set to $\tilde{\iota}_{jrm t} = 0$. The retailer, as a multi-product and multi-market monopolist, chooses the vector of prices, p_{rt} , to maximize the total profit in month t ,⁵

$$\pi_{rt}(p_{rt}) = \sum_{j \in J_{rmt}} \sum_{m \in M_r} ((1-f) p_{jrt} - c_{jrt}) \cdot \tilde{s}_{jrm t}(p_{rt}) \cdot h_{rmt}, \quad (13)$$

⁴Accommodating the full solution to forecasting the impact of $\hat{\xi}_{jrm t}$ that accounts for the unobserved $\iota_{jrm t}$ requires a computational complex integration over many random variables (all products). Hitsch (2006) sets the entire $\hat{\xi}_{jrm t} = 0$ and finds it does not affect the results meaningfully. We relax this assumption, setting $\iota_{jrm t} = 0$.

⁵One might wonder whether static profit maximization is a reasonable assumption. Following Hendel and Nevo (2006b), we test for consumer stockpiling and find little evidence in support of consumer stockpiling (see Appendix Table 4).

where M_r is the set of markets the retailer operates in, h_{rmt} is the exogenous market size of market m for retailer r in month t (which is local population times the retailer's share of grocery revenue share in the market), c_{jrt} is the wholesale price (marginal cost) for product j in t , and $f = 0.17$ is the share of gross revenue levied by the state, which is excluded from the list price. This profit maximization problem leads to the first-order condition such that for all j ,

$$\sum_m (1-f) \tilde{s}_{jrm} (p_{rt}) h_{rmt} + \sum_{j'} \sum_m ((1-f) p_{j'rt} - c_{j'rt}) \frac{\partial \tilde{s}_{j'rm} (p_{rt})}{\partial p_{jrt}} h_{rmt} = 0. \quad (14)$$

Using matrix notation, and inverting prices, reveals that:

$$p_{rt} = \frac{c_{rt}}{1-f} - (\Delta_{rmt})^{-1} \sum_m \tilde{s}_{rmt} (p_{rt}) \cdot h_{rmt} \quad (15)$$

where the j, j' th element of Δ_{rmt} is $\frac{\partial (\sum_m \tilde{s}_{jrm} h_{rmt})}{\partial p_{jrt}}$.

We impose this optimality condition in the last six months of the sample and calculate implied markup $(\Delta_{rmt})^{-1} \tilde{s}_{rmt}$ over “effective cost” $c_{rt}/(1-f)$ based on demand estimates. We use these markups to back out both wholesale prices $c_{rt} = \{c_{jrt}\}_{j \in J_{rt}}$ and margin as $(1-f)(\Delta_{rmt})^{-1} \tilde{s}_{rmt}$ divided by price (see Table 2 for examples).

We then project these wholesale price estimates to the rest of the sample. To do so, we assume that

$$\log(c_{jrt}) = \bar{c}_{jr} + \tau_{k(j)y(t)} + \omega_{jrt}. \quad (16)$$

The log wholesale price is composed of a component that is constant over time, but varies across products and retailers (\bar{c}_{jr}), and one that captures common category-year time variation ($\tau_{k(j)y(t)}$). In addition, ω_{jrt} are mean-zero cost shocks. We estimate \bar{c}_{jr} as the average log wholesale price using the only the last six months of the 2016 sample. Separately, we estimate $\tau_{k(j)y(t)}$ using the implied Oregon wholesale prices for the period before 2015 and normalizing $\tau_{k(j)y(t)} = 0$ for the period of 2015-2016. Together, these estimates allow us to form wholesale price estimates for each product and time period.

We find that these wholesale prices have large dispersion across different products, but such dispersion is expected given the large amount of vertical differentiation in the category. In Year 2016, the 5th percentile of wholesale prices (among product-retailer pairs) is \$4.09, the median is \$12.23, and the 95th percentile is \$31.38. More details are presented in the Web Appendix: Figure 13 visualizes the distribution of estimated average wholesale prices \bar{c}_{jr} , and Table 7 presents estimates of wholesale price time trends from the Oregon wholesale price data. These estimates suggest that the wholesale price trends reflect a combination of inflation and exchange rate variations.

5 Do retailers learn, and how quickly?

In this section, we delve deeper into the size of mistakes, the speed of learning, what affects these mistakes and learning, and conclude with what retailers are learning about. We start the section by presenting evidence that retailers approach optimal pricing behaviors by the end of our sample period. We then turn to using our structural model estimates to simulate optimal prices for the rest of the sample period. By doing so, we construct a benchmark against which to compare the observed pricing behaviors. We use this comparison to calibrate the size of pricing mistakes and when those mistakes are reduced to approximately optimal prices. We then show a series of elaborations that demonstrate that the information content of experience matters, that experience in prior liquor markets matters, and what retailers are learning about from experience.

In most of this section, we use the structural model to simulate a counterfactual in which retailers have full information about demand and set prices optimally. This counterfactual serves as a benchmark to evaluate how close the observed retailer behavior is to perfect-information behavior. To solve for model-implied optimal prices we use our supply-side estimates to compute the implied marginal costs, \hat{c}_{jrt} , as predictions of equation (16) setting the error term ω_{jrt} to zero. We solve for product-specific optimal prices for each of the three retailers in each month, assuming

each product’s price is uniform within the state.⁶ After obtaining p_{rt}^* , we compute the implied total profit for retailer r in month m across all products and markets from the retailer’s profit function (equation (13)).

We contrast these optimal full-information profits against the profit implied by the observed prices. We measure the differences in observed prices and profits relative to the full-information benchmark: We construct a measure “%price gap” capturing the relative differences between the observed prices and the optimal prices implied by our estimates, or

$$\%price\ gap = \frac{P_{jrt}^* - P_{jrt}}{P_{jrt}^*}. \quad (17)$$

We also examine the percent profit gap obtained at observed and model-implied optimal prices, $\frac{\sum_r (\pi_{rt}(p_{rt}^*) - \pi_{rt}(p_{rt}))}{\sum_r \pi_{rt}(p_{rt}^*)}$, where the profit function $\pi_{rt}(\cdot)$ is evaluated at our estimates of demand and costs.

Throughout this section (and the next one), we focus on three grocery retailers that were stable throughout our observation period—Retailer 158, 182, and 32—and that account for 80.7% of revenues. Retailers 152 and 182 operate in other states, but only ones where liquor cannot be sold in grocery stores, meaning that they have no direct experience with this category. Retailer 32, in contrast, has extensive experience selling liquor in the other states in which it operates, just not in Washington.

5.1 Do retailers approach optimal prices through learning?

Do retailers learn about demand within the sample period, i.e., within four years after entering the market? Our estimation of wholesale prices assumes that retailers are capable of setting full-information optimal prices by the end of the sample. To see whether this assumption is reasonable,

⁶Specifically, we jointly solve for p_{rt}^* as the implied full-information optimal price vector, defined as a solution of the first-order condition defined by equation (15). The demand shocks are set to $\hat{\beta} \hat{\xi}_{jrmt-1}$, the predictable part of the demand shock, noting that the innovation term ι_{jrmt} is very small. To find the optimal prices, we iterate on equation (15) until p_{rt}^* converges for all products in each retailer-month. We also use an optimizer to directly maximize the joint profit and find the solution of optimal prices is robust.

we compare our estimates to direct measures of wholesale prices for a subset of products from the pre-privatization period. Before mid-2012, Washington State applies a fixed, 51.9% markup above the wholesale price for all products. One should note that a fixed markup rule will lead to shading of wholesale prices in cases where the fixed markup differs from optimal retail markups. However, comparing the model-implied wholesale prices to observed ones should still be instructive, particularly for products that have post-privatization retail markups that are similar to the pre-privatization fixed markup. To do so, we back out the wholesale price from published retail prices using data between March to May 2012, hand-match 65 products with size 750ml that are offered both before and after the privatization, and compare these prices to the corresponding model-implied ones.

If, by the end of the sample period, retailers have learned about demand and set prices optimally using this information, our estimated wholesale prices should be similar to the observed ones in all periods. We show in Figure 7 that their estimated wholesale prices are close to those backed out from the state data. The median wholesale price is \$12.1 in the privatized market, almost identical to the median wholesale price observed from the pre-privatization era at \$12.4. In addition, across products, the model-implied- and state- wholesale prices have a correlation coefficient of 0.99. Further, visual inspections suggest that the pre-privatization wholesale prices and the model-implied ones are closely related, but do have a slight rotation. The rotation is consistent with what one might expect from wholesale price shading arising from the fixed markup rule imposed during the pre-privatization period. As apparent in the figure, higher-priced products are affected more by this rotation than lower-priced products. As a result, the mean wholesale price is \$14.0 in the privatized market and \$13.0 in the pre-privatization era. Four premium whiskeys show cost differences larger than \$4, but they account for less than 0.5% of the total sales volume. For the products that make up most of the revenue in the market, the estimated wholesale prices seem to be within a small margin of error.

As a further validation test, we also compare the implied retailer wholesale prices with state wholesale prices in Oregon at various time periods. We find they are similar (see Appendix Figure 14), though the rotation that particularly affects the higher-priced products is more pronounced as

one would expect due to the lower fixed-percent-markup rule in Oregon. Altogether, these results provide external validity that the model does a good job recovering wholesale prices that represent marginal costs for the retailers. Importantly, this validation also implies that, after operating in the market for a few years, retailers behave in ways consistent with the full-information optimal pricing.

[Insert Figure 7 about here]

5.2 The size of mistakes and speed of learning

Figure 8 presents, across all retailers for each quarter,⁷ the gap between the observed and optimal prices and total profits. Prices in the first quarter differ markedly from the model-implied optimal prices, resulting in considerably lower profit in this period. In the median, prices are 9.3% higher than the optimal level with large dispersions (quartiles are [0%, 21%]), and such difference implies 9.3% lower profit compared to the full-information optimum. These price and profit differences in the initial period suggest that limited information about the new market is consequential in setting prices – the scope of learning about demand is high.

The gap between observed and optimal prices decreases over time, and particularly fast in the early periods. Already by the second quarter of the sample, prices are much closer to the optimum than at the start of the market: the median %price gap is 5.8%, quartiles are [-4%, 15%]. As a result, the profit gap shrinks to 6.2% – a third of the lost profits are recovered from learning in one quarter. By early 2016, the profit gap closes by 7.4 percentage points, which is 80% of the initial profit gap. This further reduction in profits results less from average price changes and more from shrinking the variation in the price gap distribution. The price gap interquartile range shrinks by 16 percentage points to [-2%, 3%] by 2016. This finding suggests that retailers learn about demand and improve pricing, and do so at a faster rate at the beginning than the end. Such learning brings significant improvements in the variable profits.

⁷Here, we define as quarter as the periods of June to August, September to November, December to February, and March to May. We shift the definition by a month to keep June 2012 in the first quarter of the sample.

[Insert Figure 8 about here]

5.3 The role of information

In section 3.2 we document that whereas Washington prices adjust based on realized quantity shocks in the period shortly after the privatization, other states' prices do not. This evidence suggests that realized quantity shocks carry information about demand and retailers learn by observing and reacting to this information from experiences in the new market.

If firms learn from observed market outcomes, one might imagine that experiences might contain different amounts of informative content. We consider the volume of sales as a reflection of the information content contained in the experiences from selling a product. The idea is that learning should occur more quickly for products with high sales volume than low sales volume because of this more informative experience. To examine this possibility, we divide products into high- and low- sales quantity groups. To avoid confounding these groups with learning, we construct a median quantity split within the retailers based on the last six months of sales in our observation period (after the firms learned demand).

Figure 9 presents %price gaps by high-volume and low-volume products. Initial prices appear to be further away from the optimum for high-volume products than for low-volume ones. However, high volume products' prices more rapidly update, moving closer to the optimum at a faster pace than lower volume products. In the first six months, prices dramatically adjust downwards, with the median price moving from 13% to only 2% above the optimal level. In addition, the distribution of price gaps compresses in this period, with retailers first correcting for the extreme pricing mistakes (such as over-pricing by >30%): In the first six months, the inter-quartile range of pricing mistakes compresses from [1%, 26%] to [-4%, 8%] of optimal prices.

In contrast, for low-volume products initial mistakes are slightly smaller than high-volume products, but learning appears to be much slower. Six months in the market, the median price gap adjusts from 9% to 5% and the inter-quartile range of mistakes changes from [1%, 23%] to [-2%, 13%]. More than a quarter of products remain priced 10% above the optimum until around

mid-2014, two years after the privatization. Low volume products, however, actually face stronger incentives to learn in terms of the impact of a price change, i.e., they have higher price elasticities. Thus, the slower learning for low volume products suggests that the pace of learning is more likely to be related to lacking information rather than managerial inattention (Arcidiacono et al., 2019).

[Insert Figure 9 about here]

5.4 The role of experience

The above overall patterns include Retailer 32, which has experience selling liquor in other states, and Retailer 158 and 182, which are new to the liquor market. One might expect retailers' prior experience selling liquor in other markets to play an important role in shaping their initial prices and learning rates in this market. Retailers with prior experience in other geographic markets might be able to transfer that knowledge to Washington, setting prices closer to the optimum at the start of the market. In contrast, retailers that have never sold liquor in other states might be less informed about demand, thus setting prices further away from the optimum, and have more to learn.

[Insert Figure 10 about here]

Retailers completely new to the liquor market. Figure 10 presents the distribution of %price gaps and the %profit gaps over time for each retailer. Focus on the two retailers with no prior liquor market experience, Retailers 158 and 182. At the start of the market, they set much higher prices than the full-information optimum: 13.8% and 16.3% higher in the median and with sizable dispersion across products. As a result, total variable profits at the observed prices are 11.8% and 14.9% lower than the optimum. The lack of experience is consequential for new retailers.

Over time, prices adjust towards the optimum. By the second quarter of 2014 for the two inexperienced retailers, the median price gap reduces to 2.2-0.0%, and the interquartile range reduces by more than half. After 1.5 years in the new market, these changes convert to a 7.9 and 11.7 percentage point decrease in the profit gaps, respectively (or 66-79% relative to the initial profit gap). This finding implies sizable gains from learning about demand within a short period of operating

in the market. Over a longer horizon, prices continue to improve, and that leads to further, but diminishing improvements in profits. By early 2016, prices are so close to optimal that retailers gain 99.5% of the optimal profit. Therefore, whereas initial mistakes are significant, the rate of learning is high for retailers who are completely new to the liquor market.

Although both inexperienced retailers learn relatively fast, their pricing patterns are distinct. Retailer 158 gradually decreases prices during the first six quarters in the new market, so that the median prices stabilize at 6.6% above the optimum. During the first two quarters of 2014, this retailer sharply drops prices so that media price is only 2.2% above the optimum, leading to sizable profit improvements. In contrast, Retailer 182 makes sharp price adjustments immediately after the first quarter in the market, so that the median price is at the optimum, closing about half of the initial profit gap. We conclude that, whereas both of the inexperienced retailers show similar initial mistakes and eventually achieve full-information optimal pricing, the rate of achieving optimal prices is heterogeneous even among inexperienced retailers. Because of the patterns of prices, we speculate that these differences likely reflect firm specific routines and managerial practices rather than differences in the initial beliefs .

The experienced liquor retailer that is new to Washington market. In contrast, Retailer 32 has operated in liquor markets in other states for many years. In the first quarter, this retailer sets prices 5.3% higher than the optimum, resulting in a profit gap of 6.3%. Although the experienced retailer still makes some pricing mistakes when the market opens, the magnitude of these mistakes is smaller than those of the inexperienced retailers. Over time, the experienced retailer adjusts prices closer to the full-information optimum. In particular, in the first year of the sample, the median prices reduce to the level of the optimum, and this change implies a reduction in the profit gap by 2.4 percentage points (38% of its initial value). By 2016, the price distribution narrows, and the profit gap reduces to between 1% and 4% of the optimal profit.⁸ Although the experienced retailer makes smaller initial mistakes suggesting that it is better-informed about demand than the

⁸It also appears that Retailer 32 practices category-wide promotions, especially since 2014. Because high-low prices are not fully captured by the model, one should not necessarily interpret the spikes in the profit gap as ex post mistakes.

inexperience retailers, there is still some meaningful scope for learning.

5.5 The nature of misinformation

Thus far, our evidence suggests that retailers eventually reach optimal pricing, that they do so as a result of responding to new information from experiences in the market, and that retailers are able to transfer knowledge gained from experience in other markets to the new market in order to generate superior initial performance. Although the underpinnings of our analysis suggest that information plays a central role, the specific patterns of price adjustment also suggest that differing managerial routines and practices between retailers might limit or shape the way (e.g., frequency and magnitude of price changes) such information is used. We now aim to provide a deeper understanding of what the retailers lack in their initial information about the market, and what it is they learn about demand.

In Section 3.5, we have discussed alternative explanations and focused our attention on two possible sources of misinformation. One possibility is high liquor taxes. Retailers have little knowledge about the overall price elasticities at such high tax levels, and, as a result, set incorrect prices for all products. Another possibility is Washington's distinct customer composition by product category (Figure 5 in Section 3.4). Notably, high-income customers in Washington have strong preferences for Canadian whiskey, relative to bourbon, rye, Irish whiskey, or scotch. Not knowing this customer distribution will lead to systematic pricing mistakes by product category.

Do either or both of the two conjectures explain retailers' initial prices? While we find that initial prices are too high, Figure 11 shows that pricing mistakes clearly differ across categories. Whereas all retailers over-price almost all bourbon and rye products by a significant margin, pricing mistakes seem to be smaller on Irish whiskeys and scotch, and almost non-existent for Canadian whiskey. These category-specific pricing mistakes are consistent initial misinformation about the composition of Washington liquor customers by product category.

[Insert Figure 11 about here]

To formally quantify these intuitions, we take a simple approach and test which of these two hypotheses –learning about price sensitivities or customer composition– can better explain retailers’ initial prices. Our approach assumes that retailers are static profit maximizers who each set prices to maximize profit conditional on its *perceived* demand. Specifically, we assume that the perceived demand follows the demand model in Section 4.1 up to differences in the utility function parameters from equations (6)-(7). We allow retailers to have two types of biased beliefs: (1) the distribution of consumer price sensitivities (α_0 and α_1), or (2) the preferences towards product categories at different income levels (γ_k).

In Web Appendix D, we document details of the calibration exercise. We also report the calibrated parameters and model fit. Figure 12 summarizes how well each hypothesis explains the initial pricing mistakes. If we were to attribute all pricing mistakes as misinformation about the price sensitivity distribution independent of product categories (i.e., biased α_0 and α_1), we would explain about 18-33% of the initial profit gap. On the other hand, if we were to attribute all pricing mistakes as mis-information about the customer composition by product category (i.e., biased γ_k), we would rationalize about 53%-62% of observed profit gap, showing that misinformation about the customer composition explains the pricing patterns much better.⁹

[Insert Figure 12 about here]

The first three columns of Table 3 summarizes the difference between perceived and estimated γ_k 's, by retailer and category. First, we find that retailers make sizable and systematic mistakes for bourbon and rye, and for Canadian whiskey. Our structural demand model estimate for γ_1 (for bourbon and rye) is -0.30, reflecting that Washington high-income consumers do not favor bourbon. In contrast, retailer-perceived γ_1 , however, are between 1.36 and 1.65 for all retailers, suggesting that retailers set initial bourbon prices far above the optimum. Relative to reality, these initial prices reflect that *all* retailers anticipate a significantly higher fraction of high-income customers to patron this category. Conversely, the perceived γ_2 (for Canadian whiskey) are all negative compared to the structural demand estimates, suggesting that *all* retailers believe that Canadian whiskeys should be

⁹We also confirm that these hypotheses imply prices that align well with the data, shown in Appendix Figure 17.

positioned at the lower-income segment.

[Insert Table 3 about here]

Second, *why* do retailers make systematic mistakes about the customer composition? The last two columns of Table 3 shows the differences in quantity shares between other states and Washington, derived from Figure 5 earlier. For example, conditional on liquor purchase, high-income consumers in other states, i.e., those with household annual income above \$85,000, purchase 9.8% more of bourbon and rye, and 17.1% less Canadian whiskey, than Washington high-income consumers. Such differences are much less dramatic among the households with lower income. These purchase patterns are consistent with the direction of initial pricing mistakes by category. If retailers set initial prices by looking at the customer composition in other states, they would have anticipated more high-income customers in the bourbon and rye category, and consequently, overprice this category relative to others. This evidence suggests that what retailers learn about demand is nuanced customer segmentation that is distinct in the new market. Learning about such segmentation differences suggests retailers attend closely to the customers that buy products and adjust prices in response to new information about these segment preferences. If this is the case, these retailers are effectively implementing a deep level of sophistication in the way that they respond to complex new market conditions.

This difference in customer composition between Washington and other states presents additional insight about the role of experience. In section 5.4, we show that the experienced retailer initially performs considerably better than the inexperienced retailers, presumably by transferring knowledge gained from previous market experience. Because the previous markets' customer composition are wrong, we learn something about the nature of this knowledge transfer. The transfer needs to be more than simple price levels or relative price arrangements. It needs to be more than simply demand levels or even aggregated demand response. Instead, this evidence suggests that firms learn about general knowledge about pricing that allow them to adapt to new market conditions, possibly understanding deep fundamentals about what drives differences in customer demand across geographies. Although such learning and knowledge might be reflected in the

kind of human capital accumulation in learning-by-doing models (e.g., Argote and Epple, 1990; Benkard, 2000), this level of sophistication also suggests deeper dynamic capabilities (Teece et al., 1997; Day, 2011). Overall, these findings are consistent with managers engaging in sophisticated pricing and learning behaviors beyond what managers can or do describe in interviews.

6 Conclusion

In this paper, we empirically study whether, and how quickly, entrants to new markets learn about demand and adapt towards fully-informed pricing. Focusing on the privatized Washington State liquor market, we document large price changes across a broad range of products and present novel evidence showing that these changes result from retailer learning about demand: prices change in response to realized demand shocks and adjust to better reflect demand primitives. We then investigate the nature of initial pricing mistakes and the process by which firms correct them, with the assistance of a structural model estimated using a combination of aggregate-level data on quantity and price, micro-data on customer compositions, and direct cost data from liquor-control states. We deliver three findings: First, retailers do learn and converge to fully-informed pricing by the end of the sample, but initial pricing mistakes are consequential and account for an average 9% loss in variable profits. Second, prior experience affect retailers' starting positions and products with rich sales information correct prices at a quicker pace. Third and finally, initial pricing mistakes are most consistent with misinformation about customer composition by product category in Washington vis-a-vis other states.

Structural models are used to make predictions of market outcomes under a counterfactual policy shock, where one makes crucial assumptions on the ways firms re-optimize under the new, counterfactual equilibrium. Canonical models assume firms are fully-rational optimizers with rational expectations – as oracles who can forecast what happens in the new market. Yet empirical evidence in many fields often reveals much slower price adjustment in macroeconomics (Bils and Klenow, 2004), pricing anomalies and violations of the efficient-market hypothesis in finance (Lo,

2017), departures from profit maximization in industrial organization and marketing (Arcidiacono et al., 2019; DellaVigna and Gentzkow, 2019), and managers self-reports of using heuristic-based decision-making (Noble and Gruca, 1999). Understanding the nature of these departures from the canonical model will affect how we interpret, or modify, these structural models. Speaking to this broad question, our evidence in this context suggest that retailers behave consistent with models of rational (or optimizing) agents, but depart from canonical models in the information sets they possess. In the short run, information acquisition and processing might be costly and such costs drive firm behavior away from the predictions of a canonical, full-information model. However, in the long run, firm behavior seem to converge to a fully-optimizing model.

References

- Adams, Brian and Kevin R Williams (2017), 'Zone pricing in retail oligopoly'.
- Aguirregabiria, Victor, Daniel Ershov and Junichi Suzuki (2016), 'Estimating the effects of deregulation in the ontario wine retail market'.
- Arcidiacono, Peter, Paul B Ellickson, Carl F Mela and John D Singleton (2019), 'The competitive effects of entry: Evidence from supercenter expansion', *Available at SSRN 3045492* .
- Arellano, Manuel and Stephen Bond (1991), 'Some tests of specification for panel data: Monte carlo evidence and an application to employment equations', *The review of economic studies* **58**(2), 277–297.
- Argote, Linda and Dennis Epple (1990), 'Learning curves in manufacturing', *Science* **247**(4945), 920–924.
- Argote, Linda, Sara L Beckman and Dennis Epple (1990), 'The persistence and transfer of learning in industrial settings', *Management science* **36**(2), 140–154.
- Benkard, C Lanier (2000), 'Learning and forgetting: The dynamics of aircraft production', *American Economic Review* **90**(4), 1034–1054.
- Berry, Steven, James Levinsohn and Ariel Pakes (1995), 'Automobile prices in market equilibrium', *Econometrica: Journal of the Econometric Society* pp. 841–890.
- Berry, Steven T (1994), 'Estimating discrete-choice models of product differentiation', *The RAND Journal of Economics* pp. 242–262.
- Bils, Mark and Peter J Klenow (2004), 'Some evidence on the importance of sticky prices', *Journal of political economy* **112**(5), 947–985.
- Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron S Jarmin, Megha Patnaik, Itay Saporta-Eksten and John Van Reenen (2017), 'What drives differences in management?'.
- Chandy, Rajesh K and Gerard J Tellis (2000), 'The incumbent's curse? incumbency, size, and radical product innovation', *Journal of marketing* **64**(3), 1–17.
- Chintagunta, Pradeep K, Jean-Pierre Dubé and Vishal Singh (2003), 'Balancing profitability and customer welfare in a supermarket chain', *Quantitative Marketing and Economics* **1**(1), 111–147.
- Christensen, Clayton M (2013), *The innovator's dilemma: when new technologies cause great firms to fail*, Harvard Business Review Press.
- Conlon, Christopher T and Nirupama Rao (2015), 'The price of liquor is too damn high: Alcohol taxation and market structure'.
- Covert, Thomas R (2015), 'Experiential and social learning in firms: the case of hydraulic fracturing in the bakken shale'.

- Day, George S (2011), ‘Closing the marketing capabilities gap’, *Journal of marketing* **75**(4), 183–195.
- DellaVigna, Stefano and Matthew Gentzkow (2019), ‘Uniform pricing in us retail chains’, *The Quarterly Journal of Economics* **134**(4), 2011–2084.
- Dixit, Ashutosh and Pradeep K Chintagunta (2007), ‘Learning and exit behavior of new entrant discount airlines from city-pair markets’, *Journal of Marketing* **71**(2), 150–168.
- Doraszelski, Ulrich, Gregory Lewis and Ariel Pakes (2018), ‘Just starting out: Learning and equilibrium in a new market’, *American Economic Review* **108**(3), 565–615.
- Ellickson, Paul B, Mitchell J Lovett and Bhoomija Ranjan (2019), ‘Product launches with new attributes: A hybrid conjoint–consumer panel technique for estimating demand’, *Journal of Marketing Research* **56**(5), 709–731.
- Goldfarb, Avi and Mo Xiao (2011), ‘Who thinks about the competition? managerial ability and strategic entry in us local telephone markets’, *The American Economic Review* **101**(7), 3130–3161.
- Hansen, Lars Peter (1982), ‘Large sample properties of generalized method of moments estimators’, *Econometrica: Journal of the Econometric Society* pp. 1029–1054.
- Hendel, Igal and Aviv Nevo (2006a), ‘Measuring the implications of sales and consumer inventory behavior’, *Econometrica* **74**(6), 1637–1673.
- Hendel, Igal and Aviv Nevo (2006b), ‘Sales and consumer inventory’, *The RAND Journal of Economics* **37**(3), 543–561.
- Hitsch, Günter J (2006), ‘An empirical model of optimal dynamic product launch and exit under demand uncertainty’, *Marketing Science* **25**(1), 25–50.
- Hitsch, Günter J, Ali Hortacsu and Xiliang Lin (2019), ‘Prices and promotions in us retail markets: Evidence from big data’.
- Hortaçsu, Ali, Fernando Luco, Steven L Puller and Dongni Zhu (2019), ‘Does strategic ability affect efficiency? evidence from electricity markets’, *American Economic Review* **109**(12), 4302–42.
- Illanes, Gastón and Sarah Moshary (2018), ‘Market structure and product variety: Evidence from a natural experiment in liquor licensure’.
- Jeon, Jihye (2017), ‘Learning and investment under demand uncertainty in container shipping’.
- Kaplan, Abraham David Hannath, Robert F Lanzillotti et al. (1958), Pricing in big business: a case approach., Technical report.
- King, Andrew A and Christopher L Tucci (2002), ‘Incumbent entry into new market niches: The role of experience and managerial choice in the creation of dynamic capabilities’, *Management science* **48**(2), 171–186.

- Lo, Andrew W (2017), 'Efficient markets hypothesis', *The New Palgrave Dictionary of Economics* pp. 1–17.
- Miravete, Eugenio J, Katja Seim and Jeff Thurk (2018), 'Market power and the laffer curve', *Econometrica* **86**(5), 1651–1687.
- Morgan, Neil A, Douglas W Vorhies and Charlotte H Mason (2009), 'Market orientation, marketing capabilities, and firm performance', *Strategic management journal* **30**(8), 909–920.
- Nevo, Aviv (2001), 'Measuring market power in the ready-to-eat cereal industry', *Econometrica* **69**(2), 307–342.
- Noble, Peter M and Thomas S Gruca (1999), 'Industrial pricing: Theory and managerial practice', *Marketing science* **18**(3), 435–454.
- Petrin, Amil (2002), 'Quantifying the benefits of new products: The case of the minivan', *Journal of political Economy* **110**(4), 705–729.
- Reibstein, David J and Hubert Gatignon (1984), 'Optimal product line pricing: The influence of elasticities and cross-elasticities', *Journal of marketing research* **21**(3), 259–267.
- Seo, Boyoung (2016), 'Firm scope and the value of one-stop shopping in washington state's deregulated liquor market'.
- Simon, Herbert A (1955), 'A behavioral model of rational choice', *The quarterly journal of economics* **69**(1), 99–118.
- Teece, David J., Gary Pisano and Amy Shuen (1997), 'Dynamic capabilities and strategic management', *Strategic Management Journal* **18**(7), 509–533.
URL: <http://www.jstor.org/stable/3088148>
- Trajtenberg, Manuel (1989), 'The welfare analysis of product innovations, with an application to computed tomography scanners', *Journal of political Economy* **97**(2), 444–479.

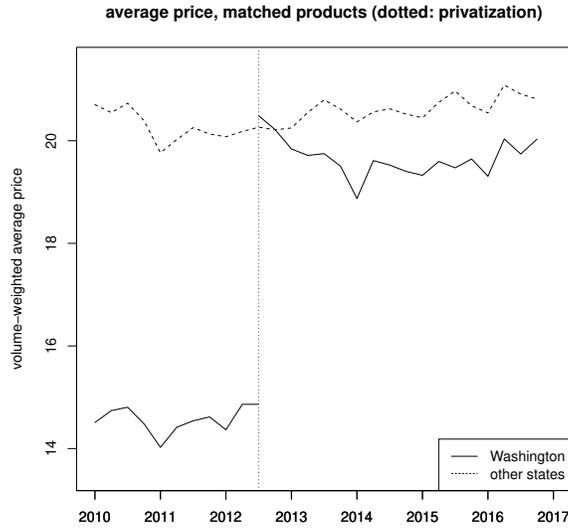


Figure 1: Changes in average price over time, Washington and other states

Notes: The left panel shows weighted-average prices for the same set of products over time, before and after liquor privatization. We use total sales quantity in the second half of 2016 in Washington as volume weight, and fix the weight over time.

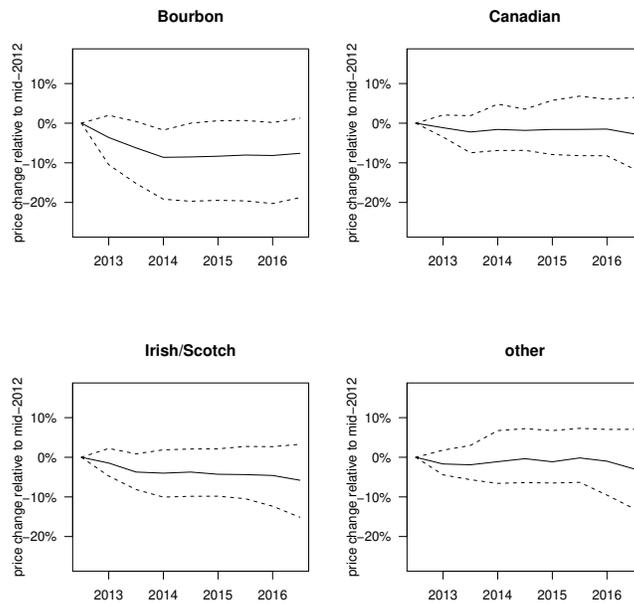


Figure 2: Distribution of price changes (mean, 25th and 75th percentile) by category

Notes: The solid lines represent mean price changes relative to the initial price for each product. The dashed lines represent 25th and 75th percentile. Each panel focuses on one category of products.

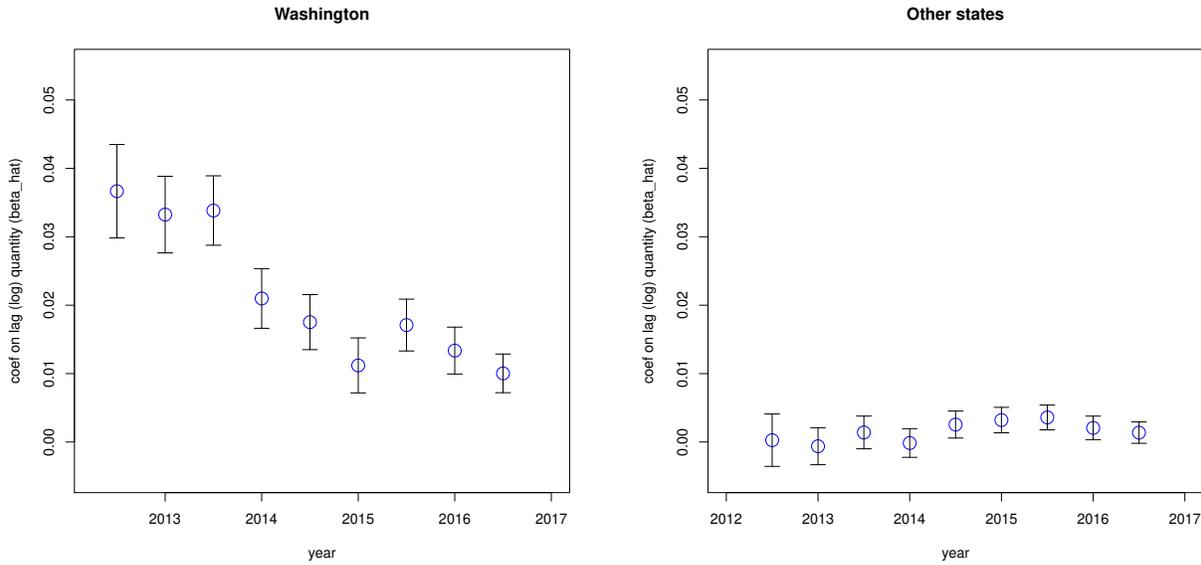


Figure 3: The response of current price to lagged quantity, Washington (left) and other states (right)

Notes: Estimates β_τ from Table 8. See table notes for more detail. Confidence intervals are two standard error.

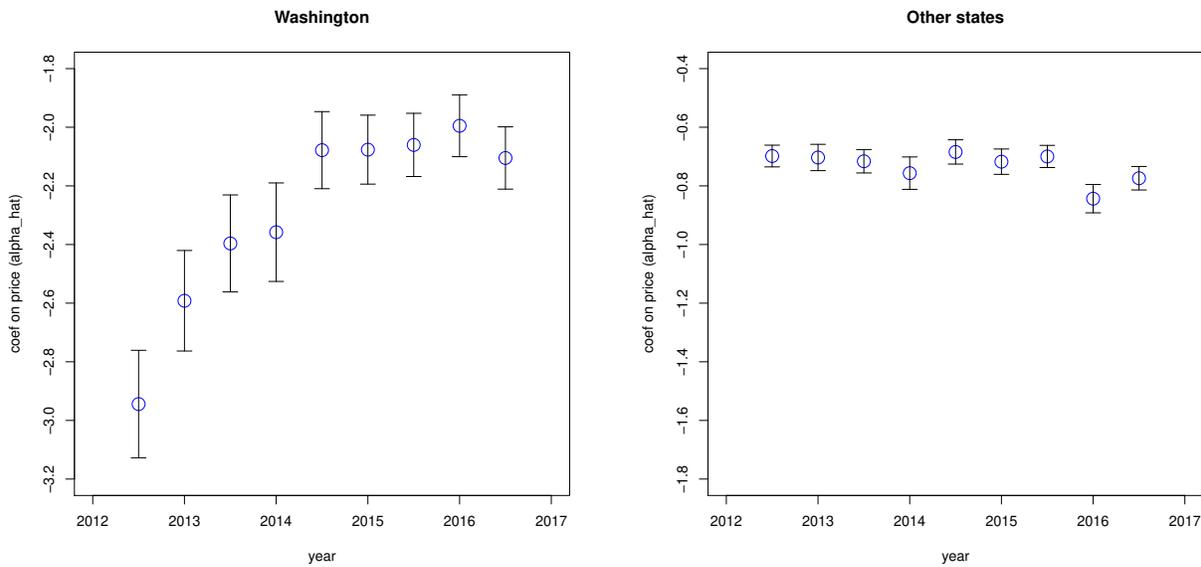


Figure 4: Coefficient estimates of α_τ from equation (3), Washington (left) and other states (right)

Notes: We report coefficient estimates of α_τ from equation (3) where τ represents half-year periods. The regressions control for retailer-state-week fixed effects but leave product fixed effects in the residual to explicitly capture the covariance between price and product fixed effects in the coefficient estimates. Confidence intervals are two standard errors.

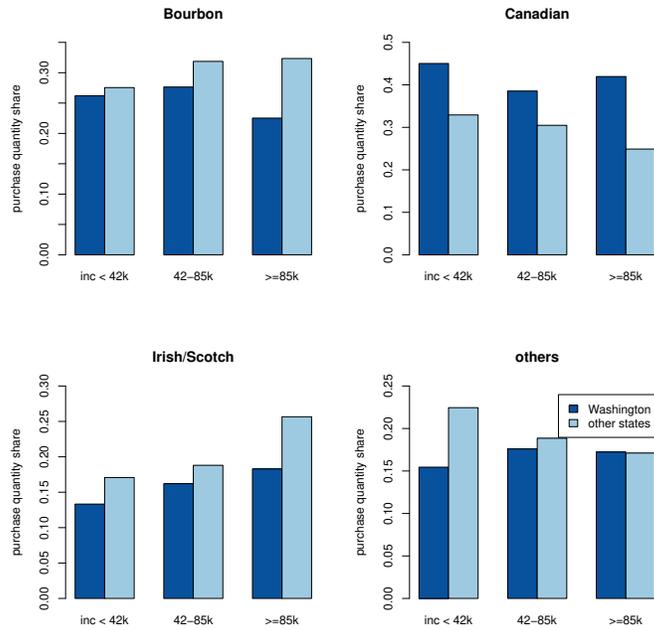


Figure 5: Share of purchases by category and income group

Notes: Probability of purchasing each product category, by income group in Washington and other states. For example, the first dark-blue bar reads: for consumers with annual income lower than \$42,000, they purchase bourbon in 26% of liquor-purchase occasions.

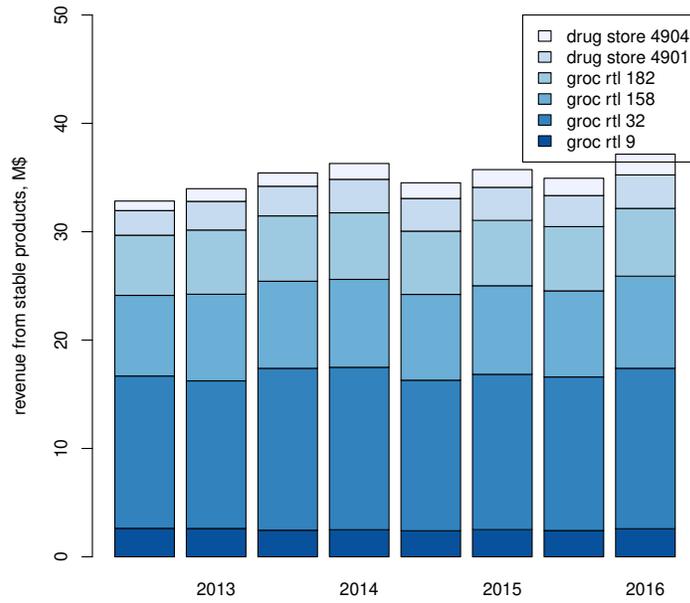


Figure 6: whiskey sales revenue decomposition, across retailers

Notes: Decomposition of liquor sales revenue across six focal retailers in Washington State, focusing on the set of core products.

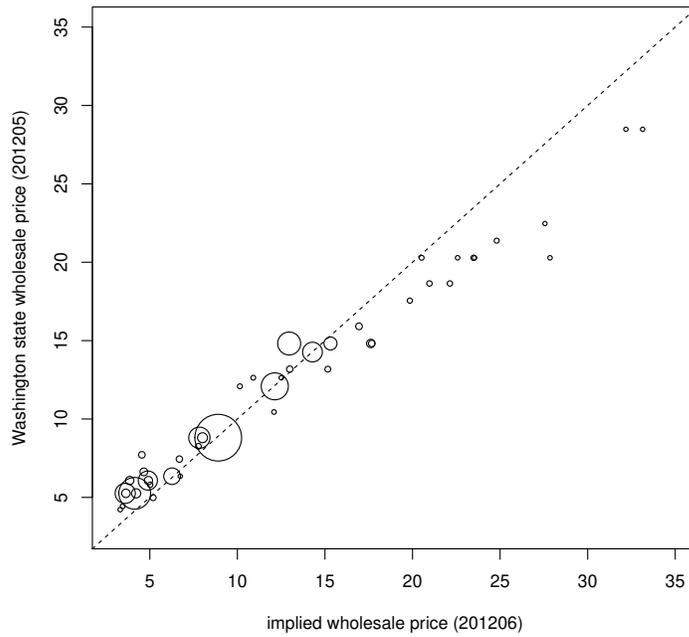


Figure 7: Comparison between estimated and observed wholesale prices

Notes: Comparison between model-estimated and observed wholesale prices. Estimated wholesale prices focuses on the period of June 2012. Observed wholesale prices come from March to May, 2012, backed out from the published price data from the Washington state-owned liquor chain. Circle size is 0.3 plus a term proportional to the average sales quantity for the product, post privatization. The dashed line is the 45-degree line on which the two measures are exactly equal.

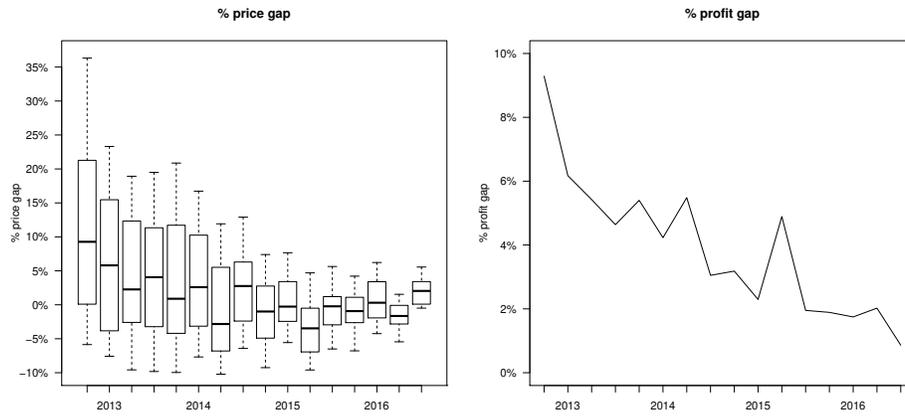


Figure 8: Percentage price and profit gap: pooled across retailers

Notes: Left panel: distribution of the percentage price gap (positive or negative) between observed prices and full-information optimal prices, where the distribution is across products and retailers. Right panel: percent total profit gap between observed prices and full-information optimum.

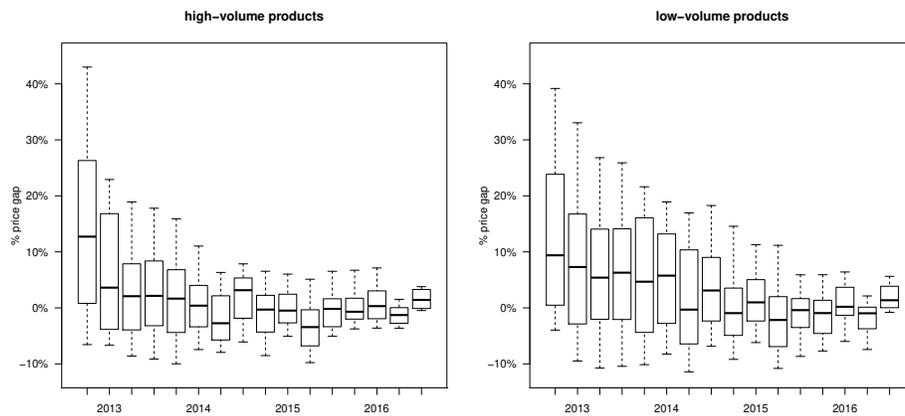


Figure 9: %price gaps by high- and low-volume products

Notes: The distribution of percent price gap for products of which the average sales volume is above or below median. The average sales volume are measured using data in the last six months of the sample.

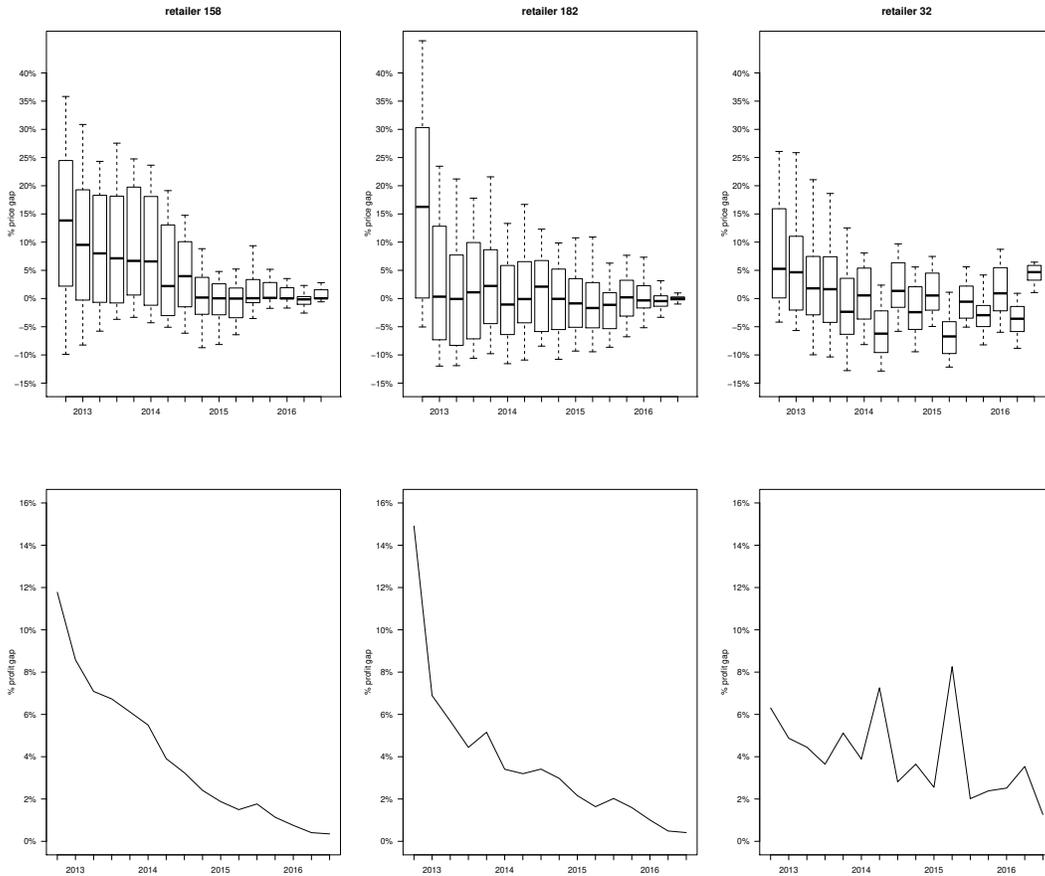


Figure 10: Percentage price and profit gap: by retailer

Notes: Top panels: distribution of the percentage price gap (positive or negative) between observed prices and full-information optimal prices, where the distribution is across products. Bottom panels: percent total profit gap between observed prices and full-information optimum. Each column represents a retailer. Retailer 32 is the experienced retailer.

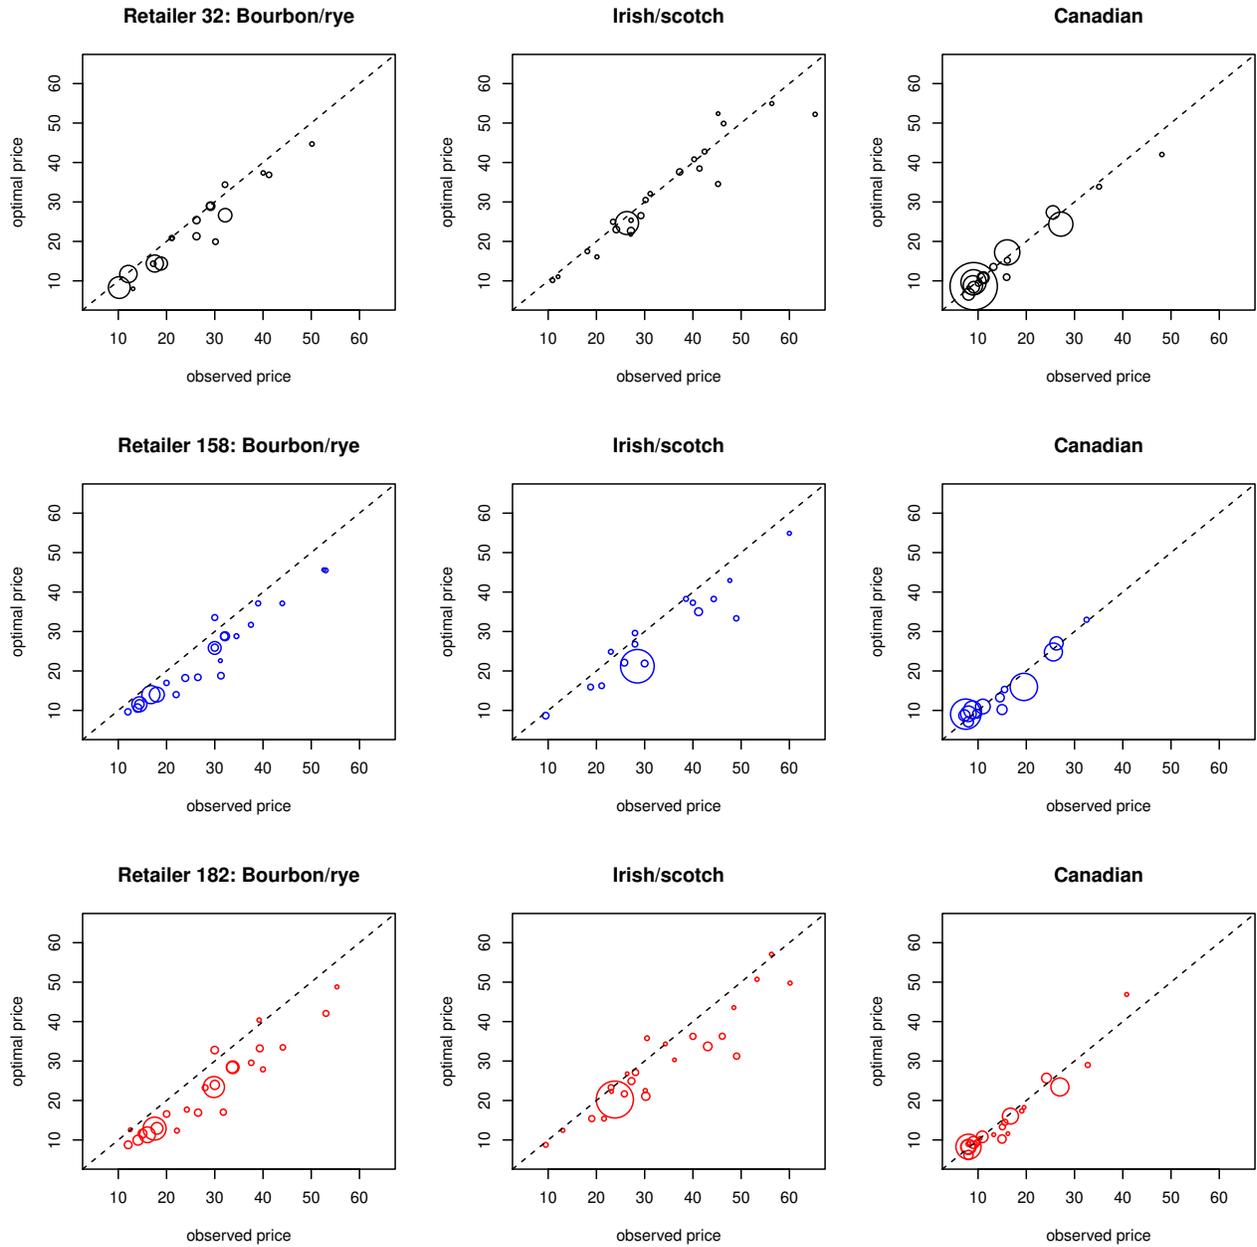


Figure 11: Distribution of initial observed and optimal prices by retailer and category

Notes: These figures plot observed prices in the first quarter of the sample against full-information optimal prices. Circle size is proportional to quantity weights, i.e., the average sales quantity shares in the last year of the sample.

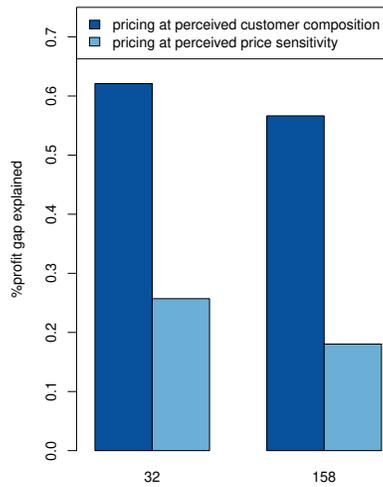


Figure 12: Share of profit gap explained by optimal pricing under two different perceived demand

Notes: This figure presents the share of %profit gap explained across the two different hypotheses about optimal pricing under perceived demand. For each retailer, the dark blue bar (left) represents share of profit gap explained if retailers set optimal prices with biased belief about customer category preferences, and the light blue bar (right) represents the share of profit gap explained if retailers set optimal prices with biased price sensitivity distributions. We calibrate these perceived demand parameters as described by Section D.

Table 1: Demand parameter estimates

	Main spec.	Berry (1994)	price (first stage)
price (α_0)	-0.833 (0.002)	-0.171 (0.015)	
price \times log (income) (α_1)	0.121 (0.002)		
std. dev. of price coef. (σ_v)	0.053 (0.014)		
bourbon/rye \times log (income) (γ_1)	-0.301 (0.116)		
Canadian \times log (income) (γ_2)	0.773 (0.100)		
Irish/scotch \times log (income) (γ_3)	0.011 (0.414)		
feature or display	0.154 (0.014)	0.145 (0.013)	-0.278 (0.037)
missing feature/display	0.239 (0.004)	0.237 (0.009)	-0.096 (0.025)
average price in other states			0.104 (0.005)
number of products			-0.004 (0.003)
product-retailer and retailer-market FE (δ_{jrm})	X	X	X
retailer-year and month FE (λ_{rt})	X	X	X
Number of observations	174,299	174,299	174,299
R-squared (linear part)	0.839	0.821	0.971

Notes: This table reports parameter estimates of the demand side. The first column reports estimates and standard error of the main specification. The second column reports estimates of a Berry (1994) logit model. The third column reports the first stage for price in the Berry (1994) logit model. The F-statistics for the two excluded instruments is 268.85.

Table 2: Example of implied elasticities and markups

Product	Price	% Margin	Elasticity of: 1	2	3	4	5	6
1	8.01	0.448	-1.978	0.007	0.007	0.007	0.007	0.007
2	11.50	0.353	0.002	-2.550	0.002	0.002	0.002	0.002
3	14.53	0.307	0.017	0.020	-2.955	0.024	0.025	0.029
4	21.29	0.234	0.003	0.004	0.004	-3.963	0.006	0.007
5	26.76	0.207	0.004	0.005	0.006	0.008	-4.551	0.010
6	29.08	0.211	0.002	0.003	0.004	0.005	0.006	-4.526

Notes: Elasticity and implied markup for six products (these products are picked because of the differences in prices), sold by retailer 32 in June, 2016. The elasticity table reads: 10% decrease in price of product 1 will increase its sales by 19.78% and decrease the sales of product 2 by 0.07%.

Table 3: Retailer-perceived category preferences and observed consumer category choice

	perceived $\tilde{\gamma}_k$ – estimated γ_k			Δ share: Other States – WA	
	Retailer 32	Retailer 158	Retailer 182	income < 85k	income \geq 85k
Bourbon/rye	1.063	1.242	1.345	0.031	0.098
Canadian whiskey	-1.689	-1.991	-2.339	-0.102	-0.171
Irish whiskey/scotch	-0.297	0.458	0.117	0.032	0.074

Notes: This table shows the difference between retailers’ perceived demand primitives (i.e., those that rationalize initial prices under an optimal static-pricing model) and the estimated ones. The table also contrasts these differences with gaps in sub-category shares between other states and Washington. Column 1-3 present the gap between perceived and estimated category-income interaction term, by category and retailer. For example, the first column reads: Retailer 32’s initial prices can be rationalized as optimal prices in a market where γ_1 (for bourbon and rye) is 1.063 higher than the true demand, γ_2 (for Canadian whiskey) is 1.689 lower, and γ_3 (for Irish whiskey and scotch) is 0.297 lower. Column 4-5 are differences in category shares conditional on purchase, calculated from the Homescan data (Figure 5). For example, Column 5 reads: for households with at least 85,000 annual income (i.e. the top 1/3 of the income distribution), they purchase 9.8% more bourbon/rye in other states than in Washington, 17.1 less Canadian whiskey, and 7.4% more Irish whiskey and scotch.

A Evidence for uniform pricing

A.1 Example

Figure 1 follows DellaVigna and Gentzkow (2019) and visualizes price variations over time and across markets for the top-selling product across all six retailers, Fireball Canadian whiskey. We take the log price for the product in each 3-digit-zipcode in each month, normalize it by its mean within each retailer, and plot the normalized log price deviations in color grids of 0.1. Dark greys show markets or months with higher prices, and light greys show markets or months with lower prices.

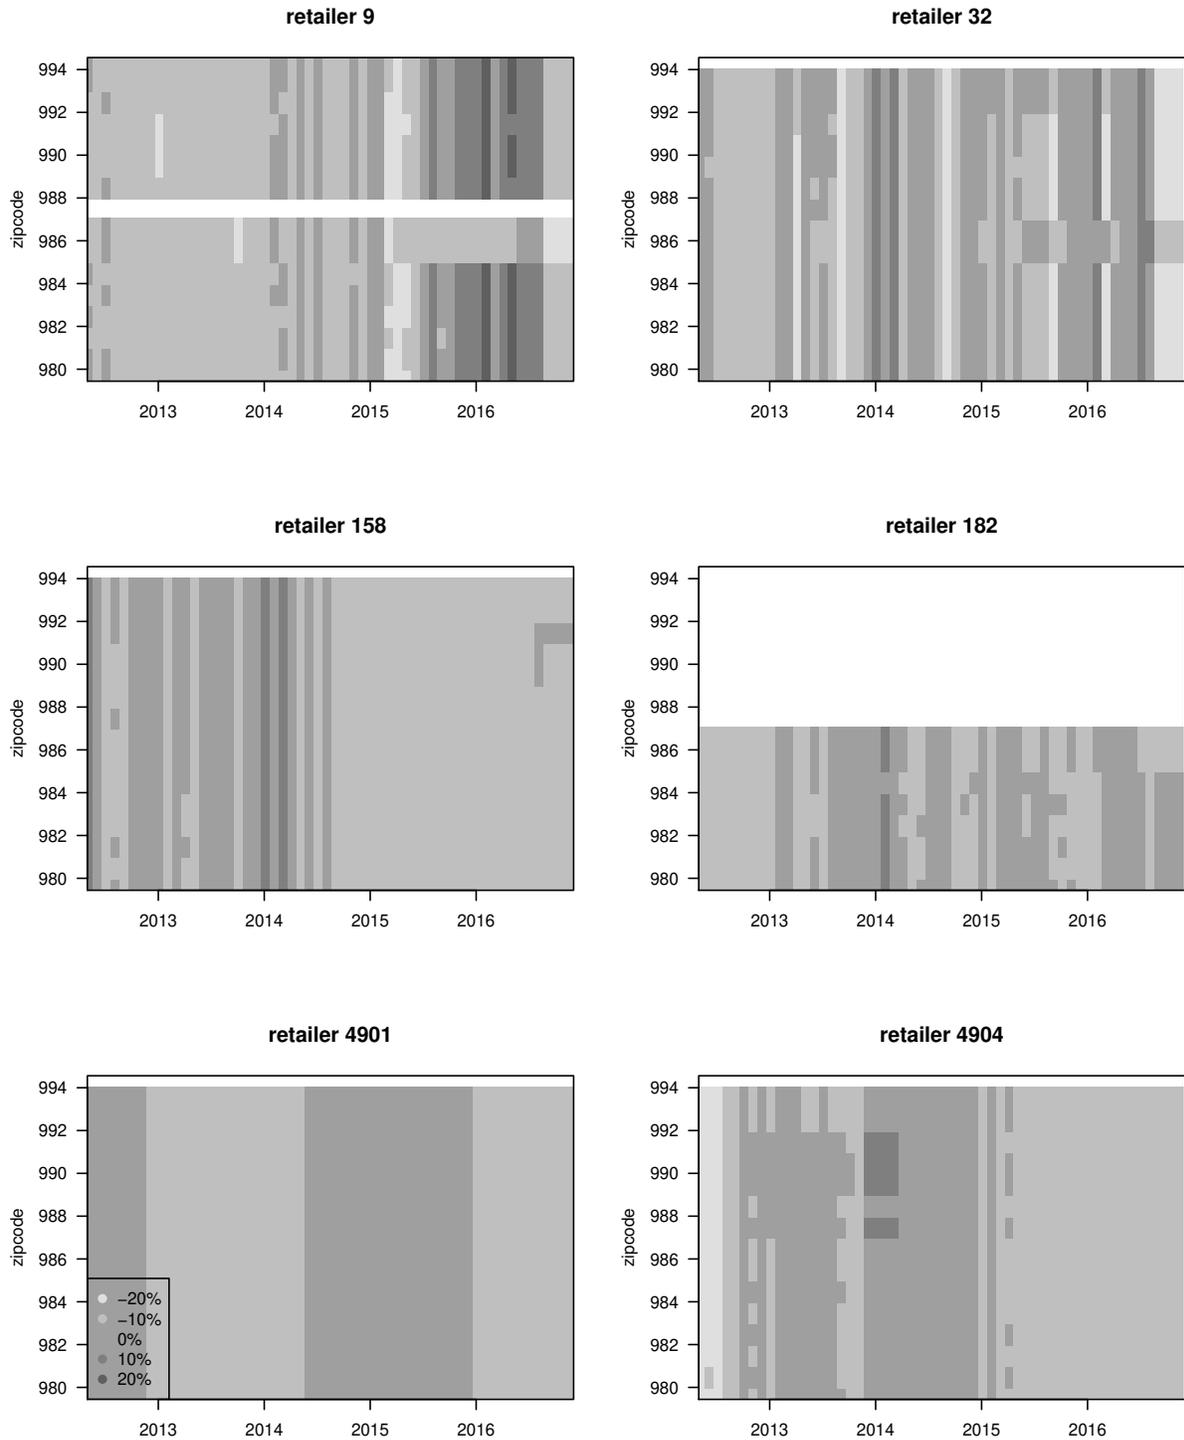
For this product, we find that prices change over time in correlated ways across markets. Prices for Retailer 4901 is the most uniform as there are no visible differences across markets at the same time. Even for Retailer 9 or 32 where there are some price variations across markets in a given month, price increases or discounts are usually synchronized across most markets. The example shows that prices are close to uniform in Washington State and suggests that each retail chain makes pricing decisions at the state level. We check and find similar patterns for other products, and show formal evidence of state-level pricing in Section A.2.

A.2 Test of uniform pricing

We formally examine the degree of price dispersion for given products across stores and chains, and show that prices seem to be set on the chain level. To examine the cross-sectional aspects, we first estimate simple specifications of log prices on various levels of fixed effects (FEs):

$$\log(p_{jrm}) = \alpha_X + \varepsilon_{jrm}, \quad (18)$$

where, with an abuse of notation, the subscripts of α_X takes product level (i.e. α_j), product-retailer level (α_{jr}) or product-retailer-market level (α_{jrm}) respectively. Because liquor products are naturally vertically differentiated, we do not interpret differences in the price levels across



Appendix Figure 1: Price variations over time and across markets: example

Notes: This figure follows DellaVigna and Gentzkow (2019) and visualizes price variations over time and across markets for Fireball Canadian whiskey across all six retailers. We take log price for the product in each 3-digit-zipcode in each month, normalize by the mean within the retailer, and plot the deviation in 0.1-grid. White represents that the product (or the retailer) is not present in the market.

products. Instead, conditional on the estimated product FEs $\hat{\alpha}_j$, we focus on whether additional FEs at product-retailer or product-retailer-market levels help explain the “left-over” variations in log prices. Specifically, to examine the extent to which product-retailer FEs explain price variations, we calculate both the incremental amount of price variations explained by product-retailer FEs, $SSR^{jr} - SSR^j$, and the total variations not explained by the product FEs, $SST - SSR^j$.¹⁰ The ratio between the two,

$$\frac{SSR^{jr} - SSR^j}{SST - SSR^j}, \quad (19)$$

measures the amount of price variations explained by adding product-retailer FEs on top of the product FEs. We calculate the incremental fit in the same way for the model with product-retailer-market FEs.

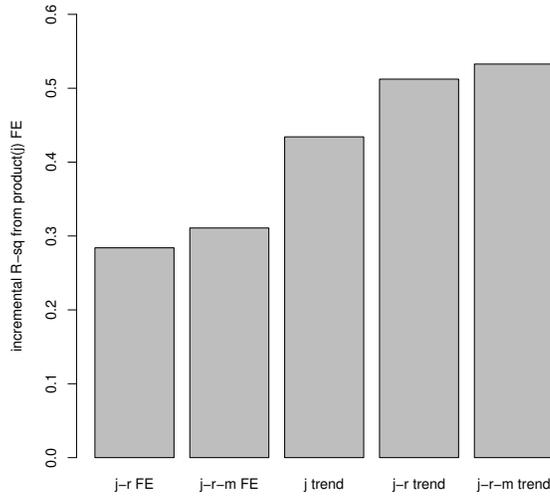
Relative to having product FEs only, we find that adding product-retailer level FEs explains 30.9% of the “left over” price variations. This number indicates that there are large differences in the price levels across retailers for the same product. However, adding the product-retailer-market level fixed effects only captures 1.7% (percentage points) additional price variations, suggesting that the meaningful cross-sectional price variations happen at the product-retailer level.

Next, we estimate a series of regressions with product-retailer-market FEs, but with linear time trends that vary at different levels. This is to say, we estimate a class of regressions

$$\log(p_{jrm}) = \alpha_{jrm} + \beta_X \cdot t + \varepsilon_{jrm}, \quad (20)$$

where we allow the coefficient on time, β_X , to vary at the product, product-retailer, or product-retailer-market level. Measuring incremental fit in the same way, we find that product-level time trend explains 48.7% left-over variations from the model with only product-retailer-market FEs, implying a 16.1% incremental fit than without the product-level trend. Product-retailer trends explain 6.2% more of the variation, while product-retailer-market trends only capture an additional 1.4%. This is to say, both price levels and price variations over time occur at the product-retailer

¹⁰Where SST is the sum of squared of the dependent variable (as a measure of total variation), and SSR is the sum of squared of the regression fit (as a measure of model-predicted variation).



Appendix Figure 2: Incremental explanatory power of co-variates to price variations

Notes: This figure reports incremental fit measures defined in (19), across log price regressions with different sets of fixed effects: product(j)-retailer(r), product-retailer-market(m), and from this version, adding product-time trend, product-retailer time trend and product-retailer-market time. The benchmark regression we use to compare fit is one with only product fixed effects.

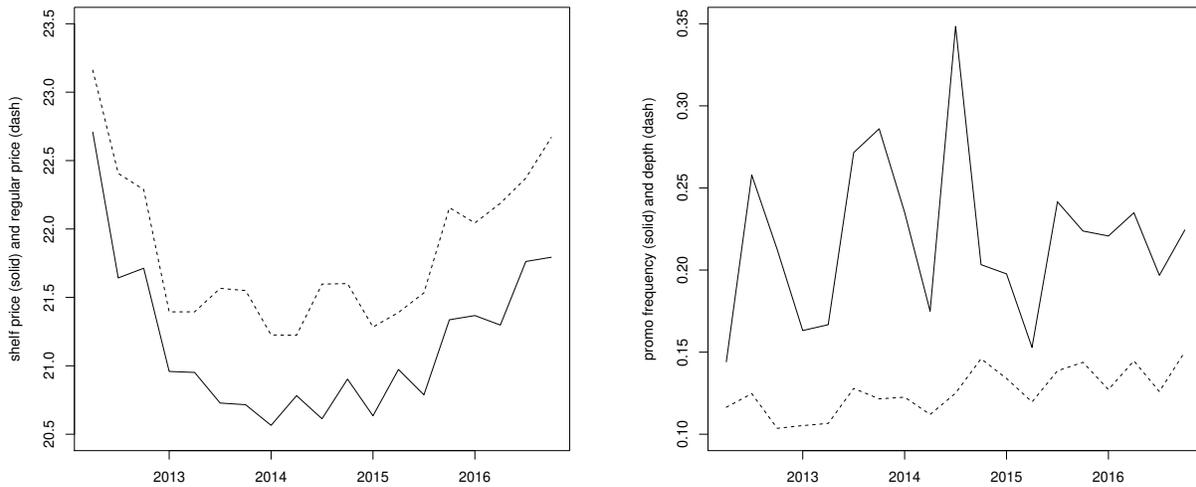
level.

B Alternative explanations and robustness checks: detail

Section 3.5 discussed evidence that shows that a few alternative explanations are not first-order in the context of Washington State liquor market. This section presents detailed evidence mentioned in the main text of this paper.

B.1 Changes in retailer promotion policies or product assortments

Promotion policies. It is also important to consider whether promotion policies evolved over the sample period, perhaps contributing to or explaining the price changes. To this end, we first define regular price as the 90th percentile of prices for a given retailer-market-product in a given quarter, and define promotion instance as when the distance between list price and regular price



Appendix Figure 3: Left: shelf price and regular price (dash), Right: promotion frequency and depth (dash)

Notes: Left panel: solid line is the average shelf price; dashed line is the average regular price defined as the 90% quantile of shelf price in each quarter. Right: solid line is promotion frequency defined as the share of products having at least 5% differences between regular and shelf prices; dashed line as promotion depth defined as the percentage differences between shelf price and regular prices conditional on product being on promotion.

is beyond 5% below the regular price (and promotion depth accordingly). Promotion as defined occurs at a frequency of 22.2%. When a product is on promotion, the average promotion depth is 12.8% of the regular price (25/75th quantiles at 7.2% and 17.1%).

Appendix Figure 3 presents time trends in shelf price, regular price, promotion frequency and promotion depth. We find that the main variations in promotion frequency seem to be seasonal variations, and promotion depth increases from about 12% to 14% over the sample period. As a result, the gap between regular and shelf price roughly stays stable in the sample period, whereas the main variations of interest seem to be the variations in regular prices over time. For simplicity, and also because the shelf price is clearly defined by the data while the regular price relies on additional definitions, we model the shelf price.

Product assortments. Alternatively, one might expect retailers to make decisions about which assortments to carry and that such decisions could also be an outcome of learning. We examine whether the decision of which products to carry changes over time and gauge the overall magnitude of assortment changes. As mentioned in Section 2, among the 724 products (as product name - size combinations, sizes can be 375ml, 750ml or 1750ml), only 276 products are “core” in the sense that they are carried by the retailer between privatization and the end of 2016. The “non-core” products represent 15.8% of the sample and account for 11.4% of the overall revenue. Appendix Figure 7 shows these core products account for most of the revenue in all time periods.

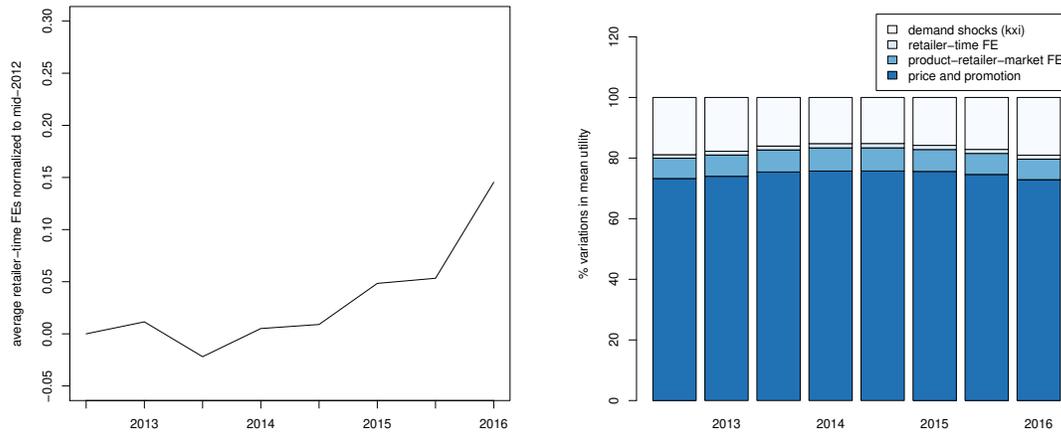
Although firms may adjust assortments over time, we conclude that these decisions are only relevant for low-revenue products. Thus, we focus our inquiry around the pricing of core products.¹¹

B.2 Changes in consumer demand

A potential concern in this context is that, if consumer demand changes over time after the privatization, the researcher would then need to explicitly model these changes on the demand side and the retailers might instead be learning about demand as moving (as opposed to stable) objects.

Based on our estimates, we investigate whether demand can be well-approximated by a stable function. We first examine whether the estimated time fixed effects imply that demand evolves over time. The left-hand side panel of Appendix Figure 4 reports the average implied time fixed effects by half-year windows. We construct these time fixed effects as estimated retailer-time fixed effects λ_{rt} , having first taken revenue-weighted average across retailers, and then normalized to the level in 2012. We find that the time effects are small: the maximum changes implied by these fixed effects are 0.14 in mean utility terms, which, at the average price coefficient, implies a \$0.22 change in willingness to pay. These changes are negligible compared to the amount of price variation in the data, suggesting that demand at the aggregate is stable and consumers are unlikely to learn about the existence of the category (or many products) over time.

¹¹Hitsch (2006) discusses firms’ choices to discontinue unexpectedly low-volume products.



Appendix Figure 4: Time trend (left) and model fit (right)

Notes: Left panel: the time effects are retailer-time intercepts averaged across retailers and across months within half-year time windows, and normalizing to the level of 2012. Right panel: Decomposition of variance in the predicted mean utility into observables (price and promotion), product-retailer and retailer-market FE, retailer-year and month FE, and the error term ξ_{jrm} . The decomposition exercise is done by half-year windows. Estimates come from the main specification presented in Table 1.

One additional concern is that demand might change beyond what the model captures. We address this concern in two ways. First, one might in particular worry about changes in consumers' price-search patterns over time. If consumers search in the early part of the sample, demand might appear to be less price-sensitive which is not captured by the model. To test whether the slope of demand varies over time, we estimate log-log (constant-elastic) demand model without price instruments, and in an alternative version, allow for the price coefficients to vary between the period before 2014 and the period after. As we show in Appendix Table 4, there are very little changes in the price elasticities between the two periods. In fact, we find that demand is 0.077 (about 4% in relative magnitude) *more* price-elastic the early period, inconsistent with the conjecture that consumer search for prices in the early part of the market, and does not rationalize the high initial prices set by the firms. However, note that an alternative explanation to this finding is that firms set more-accurate prices and the issue of price endogeneity intensifies over time.

Appendix Table 4: log-log demand estimates without instruments

	(1)	(2)	(3)
log price	-1.753 (0.018)	-1.730 (0.018)	-1.588 (0.020)
log price · 1 (year ≤ 2014)		-0.077 (0.006)	
duration since previous promotion			0.131 (0.008)
duration squared			0.017 (0.012)
feature or display	0.150 (0.011)	0.149 (0.011)	0.148 (0.012)
missing feature/display	0.205 (0.008)	0.204 (0.008)	0.193 (0.008)
product-retailer and retailer-market FE (δ_{jrm})	X	X	X
retailer-year and month FE (λ_{rt})	X	X	X
Number of observations	174,299	174,299	174,299
R-squared	0.854	0.855	0.856

Notes: This table reports parameter estimates of a log-log model without price instruments (column (1)), and contrasts that against a model where the price elasticities vary between the period before 2014 and after (column (2)), and a model that includes duration since previous price promotion (column (3)). Duration is the number of months since previous promotion multiplied by 4.5 and divided by 100 – to match the specification in Hendel and Nevo (2006b).

Second, we examine model fit by half-year time windows. As shown on the right-hand side panel of Appendix Figure 4, the model explains the data well throughout the sample: R-squared for the mean utility is at least 0.8 and remains stable throughout 2012 to the first half of 2016.¹² These results suggest that demand can be well-approximated by a stable function, implying that we can indeed think of firm learning as tracking a stable object.

We further show, in Appendix Figure 9, that several measures of consumer behavior using the Nielsen Homescan data are stable in the early periods after the privatization.

B.3 Changes in market structure

We next examine whether the market structure across retailers change over time. First, we show that there are few dedicated liquor trips and almost all liquor-shopping trips involve significant grocery expenditures. The median share of liquor expenditure per trip is 0.1% among all trips, or 28.7% among trips with liquor purchases. Appendix Figure 9 shows the distribution of liquor expenditure shares among trips with liquor purchases. This finding suggests that liquor shoppers are part of the regular grocery-shopper population, instead of being a new and isolated consumer segment over which the firms might compete.

Second, we test for whether grocery chains compete over liquor buyers. Specifically, denote j as a product (pooled across sizes), r as a retail chain, z as a 5-digit zipcode, and t as a month. We estimate a linear model of log sales quantity on the availability of the product in the same chain and in other chains in the market:

$$\log(q_{jrt}) = \beta_1 n_{\text{store}}_{jrt} + \beta_2 n_{\text{store}}_{j,-r,zt} + \delta_{jrz} + \psi_{jt} + \varepsilon_{jrt}. \quad (21)$$

Here, n_{store}_{jrt} is the number of stores in chain r within the 5-digit zipcode z where product j is available, and $n_{\text{store}}_{j,-r,zt}$ is the number of stores in other chains (among the six focal

¹²As we discuss in Section 4.4, we find new entry and many new products in the second half of 2016, the last half-year window within our sample period. We do not analyze this part of the data on the supply side to keep the analysis simple.

chains) where j is available. Nielsen RMS data only identify store location up to 3-digit zip-code level, which is too large to measure spatial substitution. We proxy chain location using the most frequently-appearing household 5-digit zipcode among shoppers to the chain.¹³ In addition, we control for product-retailer-zipcode fixed effects and product-time fixed effects. With these controls, and also keep in mind that we focus on products that are available at the start and at the end of the sample, variations in the number of stores carrying the product comes from store entry and exit (see Figure 6) and from stores not carrying the product in a subset of weeks (likely, these are due to stockouts). In addition, we also test whether promotion in other chains in the market reduce sales in the focal chain. In a similar model as (21), we estimate the effect of share of stores on price promotion (defined as price at least 5% below the regular price), for the focal chain and for other chains.

As reported in Appendix Table 5, we find that if the product is available in one more store in the same chain, sale quantity increases by 26%. This finding is coherent with Seo (2016), in that local availability is important to demand. However, availability in one more store in other chains do not affect sales quantity in the focal chain. Similarly, our estimates on price promotion suggest that the share of stores where the product is on promotion does not affect sales quantity in the focal chain. These results suggest that liquor sales do not substitute between chains. These evidence further suggest that substitution between chains for a liquor product is not detectable in our sample. Therefore, is it a reasonable approximation that grocery retailers set liquor prices for existing customers as if they are a monopolist on this customer base.

Third, to complement the above findings, we show that there are limited overlap for grocery shoppers among the set of focal retailers. Among household-year and among the set of focal retailers, the median share of expenditure at the primary chain (i.e. the chain where the household spends the most in that year) is 80%. In contrast, the median share of expenditure at the secondary chain (i.e. chain with the second-highest expenditure) is 17%. Given the limited shopper overlap, and that few grocery shoppers buy liquor, it seems unlikely that chains compete in liquor prices.

¹³Our approach follows an earlier version of Illanes and Moshary (2018).

Appendix Table 5: Sales quantity on availability and promotion of other retailers

	(1)	(2)
#stores carrying the product, own chain (β_1)	0.2578 (0.0063)	
#stores carrying the product, other chains (β_2)	0.0002 (0.0012)	
#share of stores on promotion, own chain		0.0797 (0.0056)
#share of stores on promotion, other chain		0.0019 (0.0066)
product-retailer-market (δ_{jrm})	X	X
product-month (ψ_{jt})	X	X

Notes: The effect of product availability and promotion, in the own chain and in other chains, on sales quantity.

In addition, we show in Appendix Figure 12 that these numbers do not change much over time and in particular show no discontinuity at the point of liquor privatization, further supporting the conjecture that the opening of the liquor market does not result in significant changes in grocery traffic patterns.

C Implementation details for the structural model

Sample selection. In structural estimation, we restrict attention to whiskey products with the size of 750ml, thus focusing on 176 products (out of 276 from grocery retailers) that take 63.6% overall grocery liquor revenue. Focusing on one size alleviates the necessity of having to model non-IIA substitution between sizes. Miravete et al. (2018) characterize substitution between different categories (e.g. whiskey or Vodka) and sizes with random coefficients on these characteristics.

We also restrict the set of inside good to products that have sold at least 2,500 bottles in total, and have prices below \$80. This selection criteria leaves us with 98 products but this set of products occupy 98.7% of the total revenue (98.4% of sample size after the previous sample reduction).

In our supply-side analysis, we further focus on three large grocery retail chains and the period before July 2016. We observe entry of a new large retail chain in the last half-year period of the data. Thus, we exclude this period to maintain a stable environment throughout the sample. The

three large grocery retailers capture 80.7% of the liquor market (in revenue shares). For the three excluded retailers, two of them are drug stores (Retailer 4901 and 4904), and one of them is a grocery chain (Retailer 9) with a small share of the Washington liquor market. We exclude this grocery chain because it closed a large fraction of its stores during our sample period.

Because our demand-side model contains rich time fixed effects and retailer fixed effects, we include all six retailers and all time periods when estimating the demand side. In past versions of the paper, we varied the range of time periods and retailers and find that demand estimates are stable.

Aggregation across stores and weeks. The original data is on the level of product-retailer-store-week. We aggregate the data to product-retailer-market-month (market defined as 3 digit zip code) for two reasons. First, liquor is a slow-moving product and there are often weeks where products have 1 or 2 unit sales. It is entirely plausible that some products have zero sales in some stores for considerable number of periods. In these cases, a random coefficient logit model is not well-defined because $\log(s_{jrm}) = \log(0)$. Aggregating the sales will considerably alleviate the problem of zero shares. Second, while there are little price variations across stores within a market, we do average over variations in prices across weeks within a month. However, because liquor is a storable product, promotional elasticities could reflect forward purchasing rather than regular price elasticities used for setting long-run price levels (Hendel and Nevo, 2006a). Using monthly data helps us to focus on long-run rather than short-run price response.

Fixed effects. In implementation, we control for all product-retailer and retailer-market fixed effects, instead of the full product-retailer-market fixed effects. Controlling for too many fixed effects will eliminate much statistical power and risk overfitting the data. Likewise, we include retailer-year fixed effects and common month fixed effects. In total, we have 412 product-retailer and retailer-market fixed effects and 79 retailer-year and month fixed effects (relative to 174,299 observations).

Market size definition. We measure market size in the follow way. We take the population in the market,¹⁴ multiply it by the share of total grocery expenditure among the set of focal retailers in the given market-month, and multiply this by 2 to allow each person to purchase at most 2 bottles of liquor a month. With the above definition, the median market share is 0.00009. The median outside good share is 0.985 and the minimum outside good share is 0.812.

Estimation and inference. We estimate model parameters θ via iterative generalized methods of moment (GMM). We first stack all moments denoted $g(\theta) = \begin{pmatrix} g_1(\theta) \\ g_2(\theta) \end{pmatrix}$, where g_1 represents moments from the aggregate data

$$\mathbb{E}[g_1(\theta)] = \mathbb{E}[\hat{\xi}(\theta) \cdot z] = 0 \quad (22)$$

and g_2 moments from micro data

$$\mathbb{E}[g_2(\theta)] = \mathbb{E} \begin{bmatrix} \bar{s}^b(\theta) - \bar{s}^b \\ \bar{p}^b(\theta) - \bar{p}^b \end{bmatrix} = 0 \quad (23)$$

and \bar{s}^b and \bar{p}^b come from the household panel. The GMM minimizes the quadratic function of moments given the weighting matrix W :

$$\mathbb{E}[g(\theta)]' \cdot W \cdot \mathbb{E}[g(\theta)]. \quad (24)$$

We estimate θ with a two-step GMM algorithm. We start with the identity matrix as the initial value of W , and estimate θ by minimizing (24). Then, we take the previous estimate of $\hat{\theta}$ to compute $\hat{W} = \mathbb{E}[g(\hat{\theta}) \cdot g(\hat{\theta})']$. Finally, we use the updated \hat{W} to estimate θ . We find no visible difference between the one-step and two-step GMM estimates.

Following Hansen (1982) and Petrin (2002), we compute the asymptotic variance-covariance

¹⁴We linearly interpolate the each pair of annual population levels to obtain the monthly population levels.

matrix of the parameters,

$$V = (\Gamma\Gamma')^{-1} \cdot (\Gamma W \Gamma') \cdot (\Gamma\Gamma')^{-1}, \quad (25)$$

where Γ is the Jacobian matrix $\frac{\partial g(\theta)}{\partial \theta}$. Unlike Petrin (2002), the upper off-diagonal of the Jacobian is non-zero because the aggregate moments $g_1(\theta)$ are informative about the random coefficients (due to “BLP” instruments).

D Calibration details for Section 5.5

This section provides technical details for Section 5.5, where we calibrate retailer optimal pricing under different forms of mis-information to the initial prices. Specifically, we examine two hypothesis: learning about price sensitivities or customer composition. Our approach assumes that retailers are static profit maximizers who each set prices to maximize profit conditional on its *perceived* demand. We assume that the perceived demand follows the demand model in Section 4.1 up to differences in the utility function parameters, which come in one of two ways. First, retailer r might perceive that all parameters are the same as the estimates except for two parameters: the baseline price sensitivity $\tilde{\alpha}_0^r \neq \hat{\alpha}_0$ and the interaction term with income $\tilde{\alpha}_1^r \neq \hat{\alpha}_1$. This is a way to capture that retailers might not have full information about the distribution of consumer price sensitivity in the new Washington market with high taxes. Second, to capture different beliefs about the composition of customers in each category, we specify that retailer r might perceive the preferences for category k as follows

$$\gamma_{ki} = \tilde{\gamma}_k^r \log(y_i), \quad (26)$$

which shares the same structure as (6) but $\tilde{\gamma}_k^r \neq \hat{\gamma}_k$. For example, our estimated demand shows that high-income Washington customers prefer Canadian whiskey ($\hat{\gamma}_2 = 0.773$) rather than bourbon or rye ($\hat{\gamma}_1 = -0.301$), consistent with Figure 5 which shows that such preference profile is different than but different from other states. If retailer r perceives Washington customers as having similar

Appendix Table 6: Calibrated perceived demand parameters

Perceived demand different at Retailer	Estimates	Price sensitivity			Customer composition		
		32	158	182	32	158	182
Perceived baseline price sensitivity $\tilde{\alpha}_0^r$	-0.833	-2.275	-0.999	-0.737			
Perceived price $\times \log(y_i)$ $\tilde{\alpha}_1^r$	0.121	0.364	0.169	0.134			
Perceived bourbon/rye preference $\tilde{\gamma}_1^r$	-0.301				0.762	0.941	1.044
Perceived Canadian whiskey preference $\tilde{\gamma}_2^r$	0.773				0.158	-0.143	-0.492
Perceived Irish/scotch preference $\tilde{\gamma}_2^r$	0.011				0.027	0.782	0.441
%profit gap: observed prices		0.066	0.126	0.132	0.066	0.126	0.132
%profit gap: optimal prices at perceived demand		0.017	0.023	0.044	0.041	0.071	0.070
% of the profit gap explained		0.257	0.180	0.329	0.621	0.566	0.532

Notes: Column 1 reports parameter estimates (same as Table 1) and the mean and standard deviation of implied price sensitivities. Columns 2-4 and 5-7 report calibrated perceived demand parameters assuming, respectively, retailers' perceived demand is different in either the baseline price sensitivity and the interaction between price and log income, or the interaction term between category indicator and log income. Row 1-5 report the calibrated parameters that minimize the sum of squared log price differences defined in (27). Row 6-7 report the % gaps between full-information optimal prices and, respectively, observed prices and optimal prices under the perceived demand. Finally, row 8 reports the ratio of row 8 to row 7, which captures the share of observed price gaps explained by optimal pricing under perceived demand.

preferences than other states, its initial pricing decisions could be captured by different $\tilde{\gamma}_k^r$'s.

We calibrate the perceived demand by matching model-implied optimal prices, under a candidate set of *perceived* demand parameters, to observed prices in the first quarter of the sample. Under each of the two conjectures, we can calibrate the perceived demand parameters $\tilde{\theta}^r$ by finding the minimal squared distances between observed prices and model-implied optimal prices at the perceived demand (for example, $\tilde{\theta}^r$ can be $(\tilde{\alpha}_0^r, \tilde{\alpha}_1^r)$). To find $\tilde{\theta}^r$ we solve

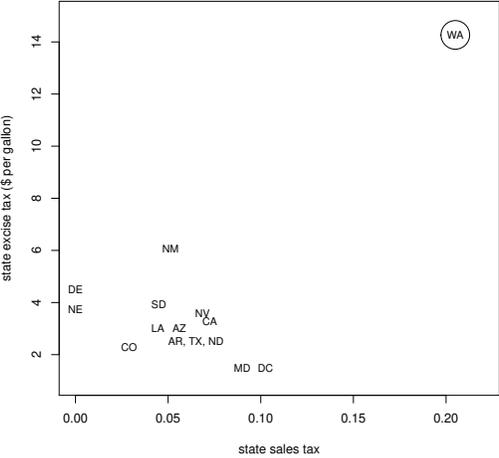
$$\tilde{\theta}^r \in \arg \min_{\theta} \sum_{j \in J_r} w_{jr} (\log(p_{jrt}) - \log(p_{jrt}^*(\theta)))^2 \quad (27)$$

where w_{jr} is a fixed quantity weight computed by the average sales quantity at the end of the sample, and $p_{jrt}^*(\cdot)$ is the solution of the first-order conditions

$$\sum_m (1-f) \tilde{s}_{jrm}^{\tilde{\theta}^r} h_{rmt} + \sum_j \sum_m ((1-f) p_{jrt} - c_{jrt}) \frac{\partial \tilde{s}_{jrm}^{\tilde{\theta}^r}(p_{rt})}{\partial p_{jrt}} h_{rmt} = 0. \quad (28)$$

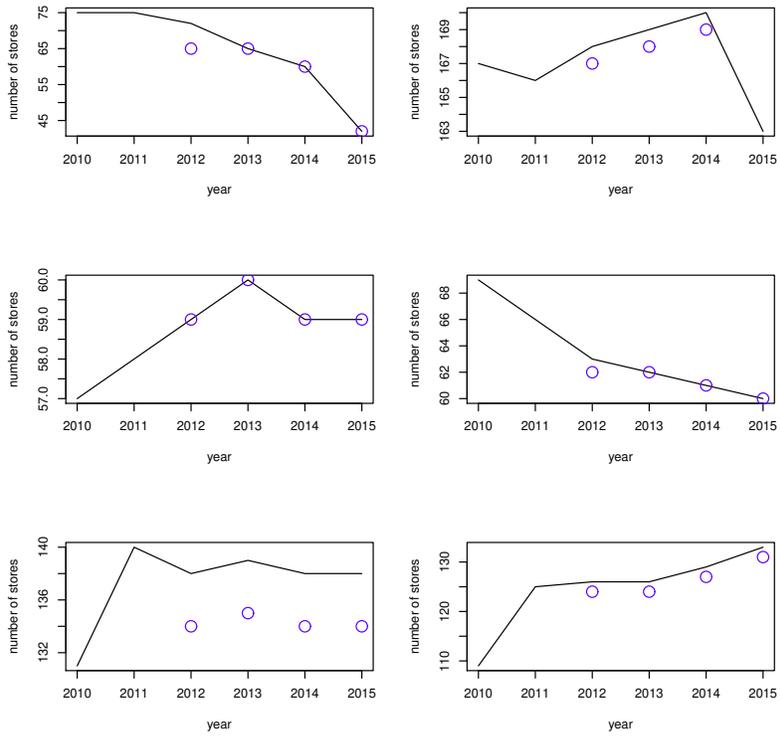
Here, $\tilde{s}_{jrm}^{\tilde{\theta}^r}$ captures the demand system (4)-(7) with parameter $\tilde{\theta}^r$. In this way, we empirically pin down retailers' perceived demand that best rationalizes the observed set of prices.

E Additional tables and figures



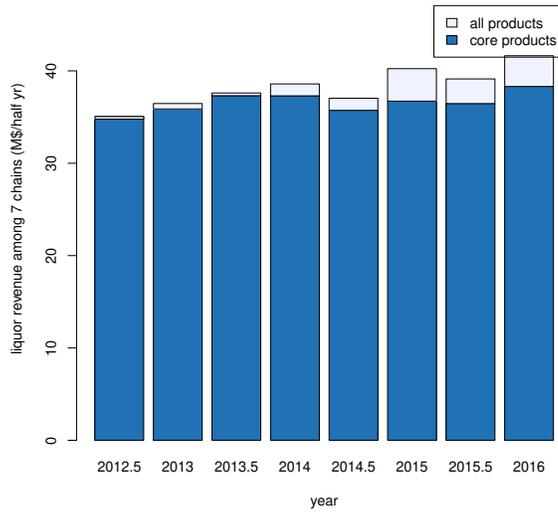
Appendix Figure 5: Distribution of liquor sales and excise taxes across states

Notes: Distribution of liquor sales and excise taxes imposed by states. Washington is highlighted by the circle. Arkansas, Texas and North Dakota have similar taxes and are grouped by the label “AR, TX, ND” for visual clarity.



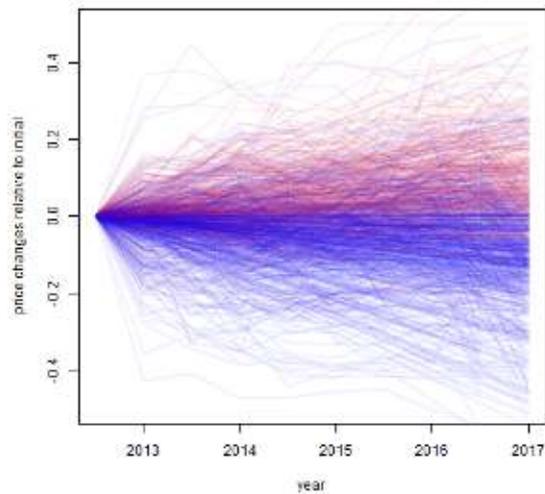
Appendix Figure 6: Number of stores selling grocery (solid) or liquor (dots)

Notes: Solid: number of distinct store IDs for each chain selling grocery. Dots: number of distinct store IDs selling liquor. For measures of stores selling liquor, we cross-checked with the number of license holders reported by WSLCB and find identical results. The 6 chains are chain 9 (top-left), 32 (top-right), 158, 182, 4901, 4904.



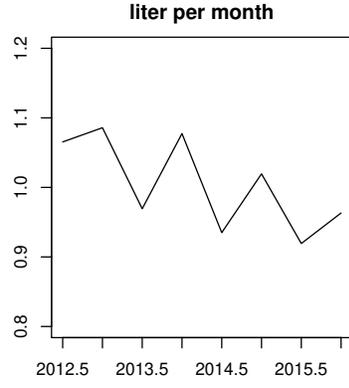
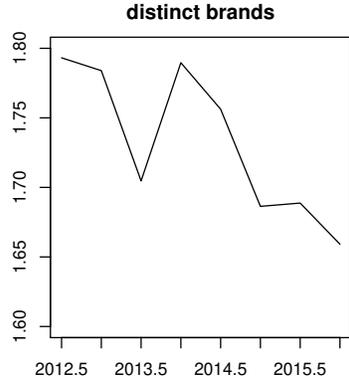
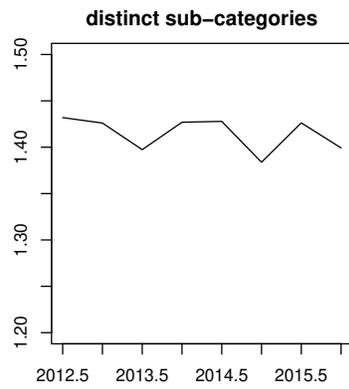
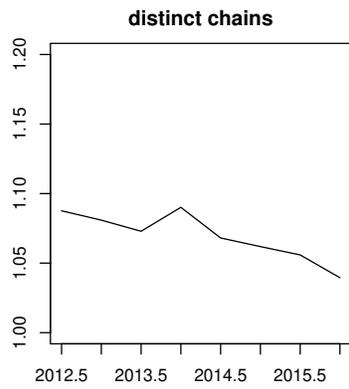
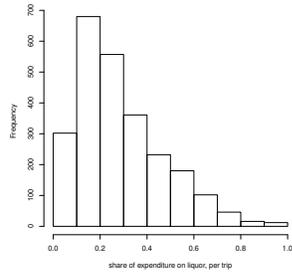
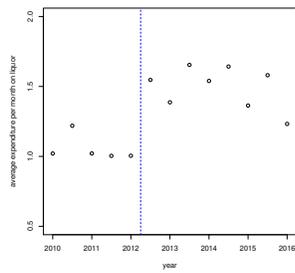
Appendix Figure 7: Total liquor sales revenue over time

Notes: Sales revenue for the focal six chains from the liquor category in half-year intervals. White: across all products. Blue: across core products defined in Section 2.



Appendix Figure 8: Price changes over time for each retailer-product (initial value = 0)

Notes: The y-axis represents the relative differences between price and initial price for a given product. Blue = final price lower than initial price. Red = final price higher than initial price.



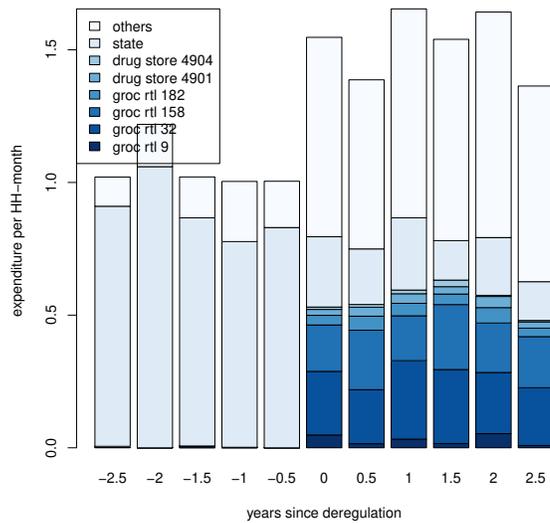
Appendix Figure 9: Descriptive evidence for consumer behavior in the WA liquor market

Notes: These panels present additional descriptive evidence from the household panel. In the top-left, we present changes in liquor expenditure before and after the privatization. For this panel, we estimate linear regression of log liquor expenditure, $\log(\text{expd} + 1)$, on a set of half-year dummies, for all consumers in Washington state in the household panel. The figure reports these regression coefficients. We re-defined half-years to “December to May” and “June to November” in order to align with the timing of the policy change. In the top-right, we present share of liquor expenditure in a trip to grocery store. The bottom panels are additional measures of varieties: the number of distinct chains, liquor product types, brands, and bottles of liquor per month, for trips with liquor purchases.

Appendix Table 7: Estimates of the wholesale price trend

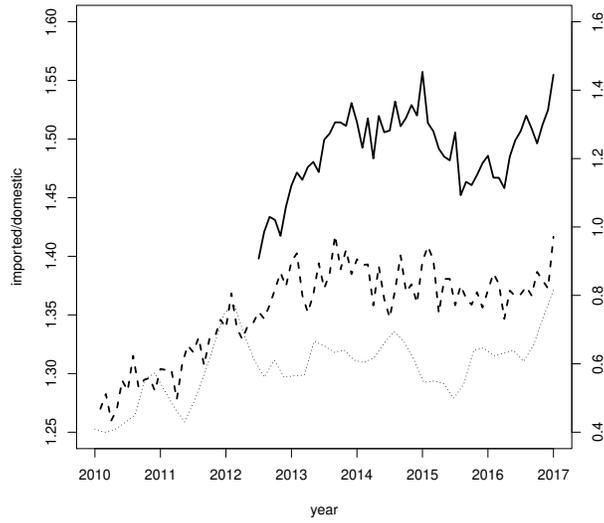
Year	Baseline	Canadian	Irish	Scotch
2012	-0.051 (0.009)	0.072 (0.015)	-0.017 (0.032)	-0.006 (0.016)
2013	-0.016 (0.008)	0.033 (0.012)	-0.005 (0.024)	-0.004 (0.013)
2014	-0.002 (0.007)	0.026 (0.010)	-0.000 (0.021)	-0.028 (0.011)
Product FE	X	X	X	X

Notes: Dependent variables are Oregon whole sale prices. An observation is a product-quarter. All category-year fixed effects are normalized to zero in year 2015-2016.



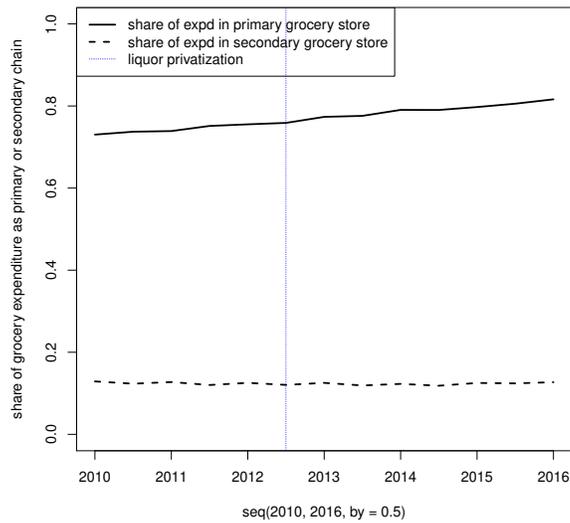
Appendix Figure 10: Liquor expenditure in the household panel

Notes: Liquor expenditure in the household panel, across the focal 6 retailers, and including the state store (or former state store) and other players.



Appendix Figure 11: Price ratio between Scotch/Irish whiskey versus Bourbon/US-made whiskey

Notes: Solid: Washington. Dash: other states. Dotted: GBP/USD exchange rate.



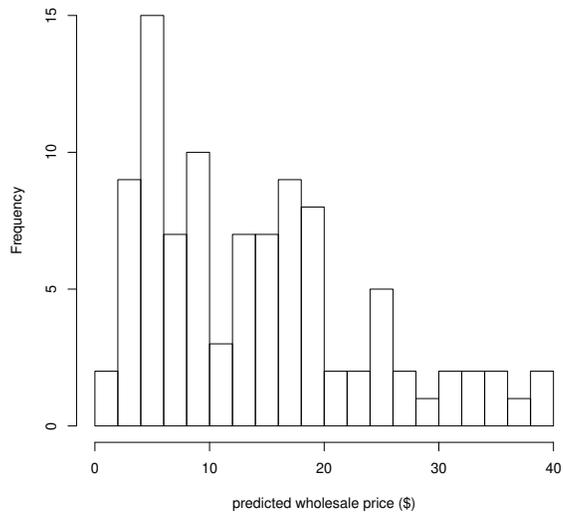
Appendix Figure 12: Shares of grocery expenditure at the primary and secondary chains

Notes: Solid: share of expenditure at the primary chain (i.e. chain with the highest expenditure share for the individual). Dashed: share of expenditure at the secondary chain. Dotted vertical line: liquor privatization.

Appendix Table 8: The response to current price to lagged quantity, Washington and other states

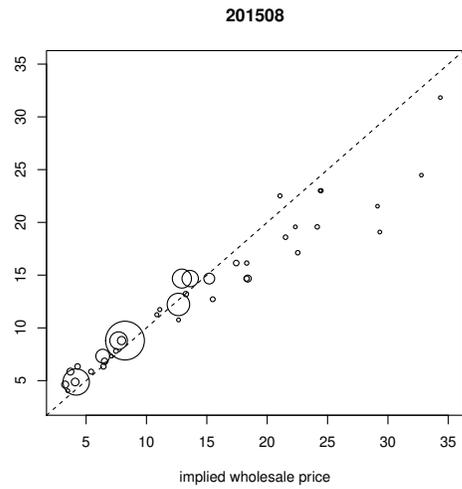
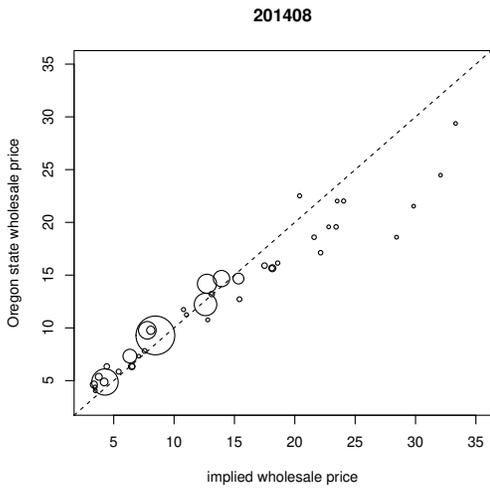
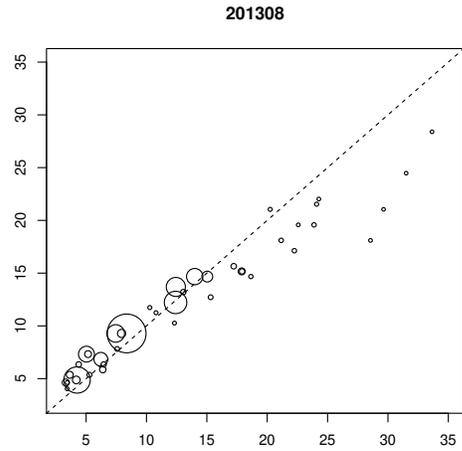
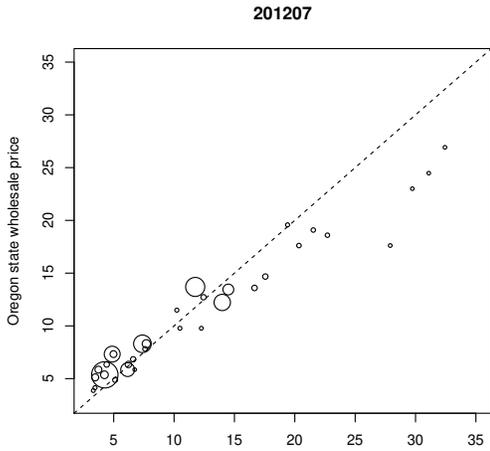
	WA: D.logprice	first stage	Ctrl: D.logprice	first stage
past month D.logprice	0.564*** (0.042)		0.191*** (0.014)	
past month D.logprice		-0.181*** (0.006)		-0.267*** (0.003)
D.log quantity	-0.033*** (0.001)	0.002** (0.001)	-0.028*** (0.000)	0.002*** (0.000)
D.past month log quantity (2012.5)	0.033*** (0.003)	-0.045*** (0.002)	-0.000 (0.002)	-0.021*** (0.002)
D.past month log quantity (2013)	0.030*** (0.003)	-0.038*** (0.002)	-0.001 (0.001)	-0.022*** (0.001)
D.past month log quantity (2013.5)	0.032*** (0.002)	-0.037*** (0.001)	0.001 (0.001)	-0.022*** (0.001)
D.past month log quantity (2014)	0.021*** (0.002)	-0.033*** (0.001)	-0.000 (0.001)	-0.023*** (0.001)
D.past month log quantity (2014.5)	0.018*** (0.002)	-0.028*** (0.001)	0.003** (0.001)	-0.025*** (0.001)
D.past month log quantity (2015)	0.011*** (0.002)	-0.029*** (0.001)	0.003*** (0.001)	-0.026*** (0.001)
D.past month log quantity (2015.5)	0.017*** (0.002)	-0.025*** (0.001)	0.003*** (0.001)	-0.027*** (0.001)
D.past month log quantity (2016)	0.013*** (0.002)	-0.020*** (0.001)	0.002* (0.001)	-0.024*** (0.001)
D.past month log quantity (2016.5)	0.010*** (0.001)	-0.013*** (0.001)	0.001 (0.001)	-0.016*** (0.001)
constant	0.032*** (0.010)	-0.036*** (0.007)	0.003 (0.005)	0.005 (0.004)
yearmonth FE	Yes	Yes	Yes	Yes
Rsq.	.	0.132	.	0.116
obs.	2.9e+04	2.9e+04	9.7e+04	9.7e+04

Notes: Column 1 reports estimates of equation (1), estimated using Washington data (first two columns) and compared against data from other states (last two columns). Lagged price difference is instrumented by the third lagged price difference, and Column 2 reports first stage results (dependent variable is lagged log price difference). Columns 3 and 4 estimates the same specification on other states. All standard errors are robust and clustered on the product-retailer level. *, **, and *** indicate significance at the 90, 95, and 99 percent. F-test for excluded variables in first stage: 392 for Washington, 912 for other states.



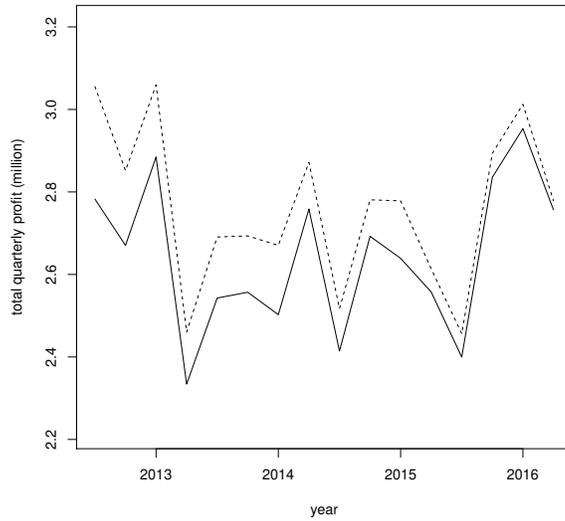
Appendix Figure 13: Distribution of implied wholesale price

Notes: Distribution of the implied product-retailer-specific wholesale price in 2016.



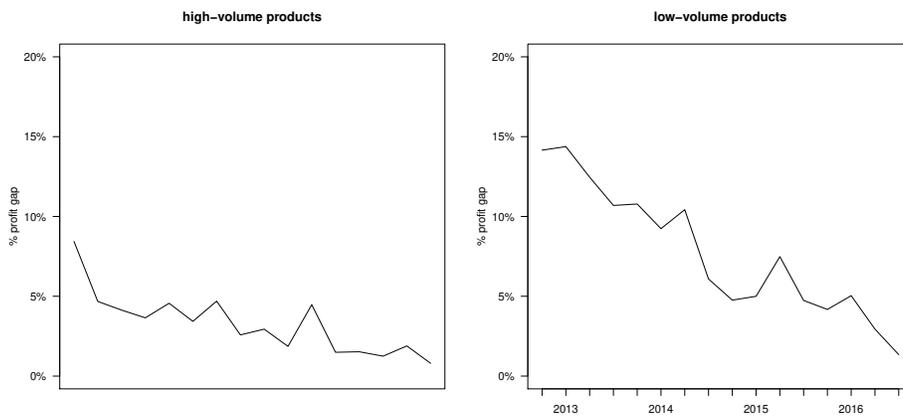
Appendix Figure 14: Comparison between wholesale prices: Washington retailers (estimated) vs Oregon State

Notes: See notes for Figure 7.



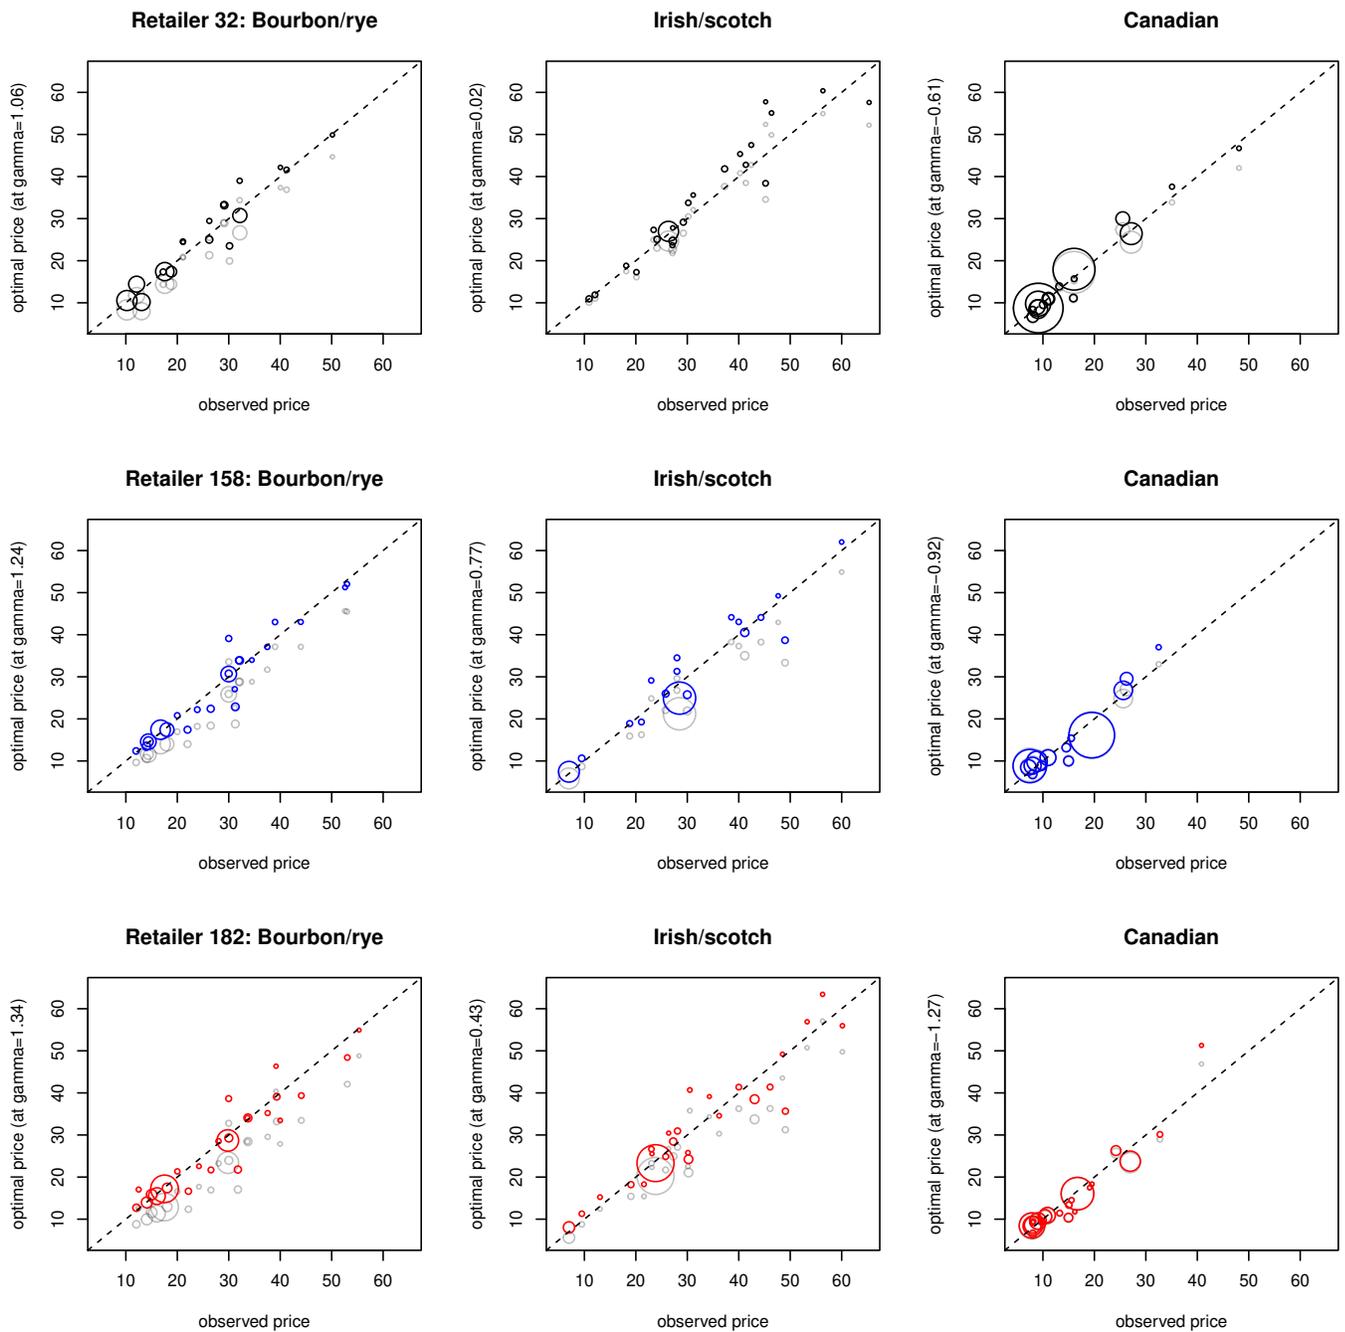
Appendix Figure 15: Total profit under observed and model-implied optimal prices

Notes: The profits are sum of profit from the six grocery retailers in given half-year window.



Appendix Figure 16: Profit gaps by high- and low-volume products

Notes: The percent total profit gap for products of which the average sales volume is above or below median. The average sales volume are measured using data in the last six months of the sample.



Appendix Figure 17: Calibrated initial prices: perceived customer composition differ from reality

Notes: These figures plot observed prices in the first quarter of the sample against model-implied optimal prices under the assumption that retailers maximize profit under perceived customer composition by product category (i.e. with different γ_k^i in their belief). Circle size is proportional to the average sales quantity in the last year of the sample. For comparison, we also show full-information prices (i.e. optimal prices where retailers are fully informed about demand) in the grey circles.