

# Learning to Set Prices

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

November 10, 2021

## Abstract

The authors empirically examine how firms learn to set prices in a new market. The 2012 privatization of off-premise liquor sales in Washington State created a unique opportunity to observe retailers learn to set prices from the point at which their learning process began. Tracking this market as it evolved through time, the authors find that firms indeed learn to set more profitable prices, that these prices increasingly reflect demand fundamentals, and they ultimately converge to levels consistent with (static) profit maximization. The paper further demonstrates that initial pricing mistakes are largest for products whose demand conditions differ the most from those of previously privatized markets, that retailers with previous experience in the category are initially better-informed, and that learning is faster for products with more precise sales information. These findings indicate that firm behavior converges to rational models of firm conduct, but also reveal that such convergence takes time to unfold and play out differently for different firms. These patterns suggest important roles for both firm learning and heterogeneous firm capabilities.

Keywords: Pricing strategy; Structural learning; Entrant behavior; Alcoholic beverage 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, Nitin Mehta, 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 Carnegie Mellon University, Northwestern University, 2018 NYU IO Day, 2017 SHUFE IO Conference, Stanford University, 2018 SICS Conference, 2019 QME Conference, University of Chicago, University of Rochester, and 2021 UTD Bass FORMS Conference. We thank Zhe Hong for excellent research assistance. Our analyses are derived based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of ours and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

# 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!<sup>1</sup>

In classical models of imperfect competition, firms are relentless in their pursuit of profit, infinite in their capacity for computation, and fully-informed of underlying economic primitives. The resulting long-run equilibria are well-defined, but often leave little scope for managerial ability or persistent economic profit. A more behavioralist viewpoint has consistently challenged whether firms in fact reach this efficient ideal (Simon, 1955) and emphasized a more central role for firms’ heterogeneous capabilities and persistent competitive advantage (Teece, Pisano and Shuen, 1997; Goldfarb and Xiao, 2011). Recent empirical evidence reveals both large departures from optimality and heterogeneous performance (Hortaçsu et al., 2019; Arcidiacono et al., 2020). Learning models, which maintain the pursuit of optimality but provide a concrete role for information acquisition and process heterogeneity (see Aguirregabiria and Jeon, 2020 for a survey), represent a promising middle ground between these competing viewpoints, illuminating the pathways by which equilibria arise and highlighting sources of at least transitory advantage. But do firms in fact learn and converge to optimal practices? If so, what market features must they infer to get there? And are there meaningful differences in how, and how quickly, they reach the optimum?

Addressing these questions requires empirically characterizing the learning process without assuming whether or how firms actually learn. To do so, we focus on a setting – new category pricing in an established retail channel – that includes both the necessary variation to detect learning and rich enough observable information to characterize optimal behavior (and identify departures from it). In particular, we examine the newly privatized Washington State retail liquor market, where firm prices, market outcomes and strong proxies for costs can be observed from the market’s inception when

---

<sup>1</sup>Source: Paul Gregutt, “BevMo! Ramps It Up in Washington State”. From [www.paulgregutt.com](http://www.paulgregutt.com).

firms had incomplete information. Tracking this market as it evolved through time, we demonstrate that firms make significant pricing mistakes upon entry, but eventually learn to set profit-maximizing prices; the broad features of the learning process are also consistent with canonical learning models. Our descriptive approach yields new evidence about how firms learn and what they uncover. We find that, though all firms eventually converge to optimal pricing behavior, firm differences in the learning process provide substantial short-term profit advantages over competitors. Notably, these advantages are tied to prior experience, a potential source of dynamic capability. We further find that these information gains accrue from learning quite nuanced features of consumer demand (e.g. a heterogeneous distribution of tastes), a more complex target than those tackled in the extant literature (but consistent with its intent).

Our research objectives require an empirical strategy that avoids placing strong assumptions on the form and structure of the underlying learning process. At a basic level, we seek to characterize the empirical relationship between the underlying market conditions (the object of firms' potential learning) and the individual prices charged (the clearest reflection of their changing beliefs), without specifying *ex-ante* the exact parameter(s) targeted by learning or the mechanism by which such learning takes place. To address this challenge, we exploit institutional features of the Washington liquor market to rule out confounding effects outside of learning. Leveraging both rich demand data and strong proxies of costs, we employ a multi-stage approach in which the least restrictive assumptions are imposed at each step.

We begin our analysis by documenting that prices display sizable and heterogeneous movements in the first two years after privatization and remain stable thereafter. These transitory, nonstationary price movements are consistent with a wide range of learning models, all of which imply systematic changes in firm behavior due to learning. We also find that the prices of different products (and product types), relative to other states, adjust in different directions, suggesting that firms are learning about the particular tastes of Washington consumers.

We next present two novel pieces of evidence consistent with firm learning about demand. Both build upon the notion that learning should manifest through an initial, but eventually vanishing, non-stationarity in the joint distribution between prices and information about demand conditions. First,

we show that prices for a given product respond to lagged demand shocks of that product, and that the rate of response to these shocks declines steadily over time. We further demonstrate via simulation that this pattern is consistent with both canonical Bayesian (Hitsch, 2006; Ching, 2010) and adaptive (Doraszelski, Lewis and Pakes, 2018) learning models, in that firms continue learning *from* the new information contained in sales until they accumulate sufficient experience to adequately infer demand. Second, we show that prices gradually adjust to capture the underlying demand fundamentals – prices increase for popular products in Washington and decrease for less popular products. Simulations establish that this pattern is also consistent with Bayesian or adaptive learning: firms learn *about* demand, identifying products that customers are willing to pay more (less) for, and setting higher (lower) prices correspondingly. These analyses also reveal that most of the learning occurs early in the sample. Later in the sample, firm strategies do not systematically adjust, suggesting that learning has ceased.

We then present more direct evidence that firm learning in fact converges to optimal practice. Aghion et al. (1991) prove that adequate<sup>2</sup> learning *can* occur under relatively mild conditions, although not universally guaranteed. How long learning takes or whether firms, in practice, reach optimality, are empirical questions. Our approach to answering both is to first recover consumer demand (the “target” of learning) and then evaluate whether firms’ eventual pricing behavior converges to the short-run maximum given this target. To do so, we employ a flexible random coefficient demand model (Berry, Levinsohn and Pakes, 1995) that includes aggregate- and micro-level moments to pin down substitution patterns. We find that the resulting parameter estimates and implications regarding preferences accord well with existing evidence on demand in related liquor markets (Miravete, Seim and Thurk, 2018; Conlon and Rao, 2019).

We then impose two conditions regarding observed steady-state behavior, namely that firms have reached full information by the last half-year of our sample (four years into the new market) and, once at full information, prices are set to maximize static profits. These conditions allow us to recover

---

<sup>2</sup>Aghion et al. (1991) define adequate learning as having occurred once “the agent acquires enough information to allow him to obtain the true maximum payoff.” At this point the agent no longer faces a tradeoff between maximizing short-run payoffs and exploiting the information content of current actions. They further note that establishing such adequacy is important because “when adequate learning does not obtain one cannot understand the long-run outcome independently of the priors or the adjustment process by which it was reached.”

firms' marginal costs. The external validity of these conditions is demonstrated by showing that prices stabilize and that the costs implied by these prices (under static profit maximization) closely match our best-available proxies: state-owned retailers' wholesale prices in Washington before the privatization and in the neighboring Oregon over several periods. The finding that prices have converged to the optimum suggests that systematic departures from this optimum are indeed transitory and eventually eliminated by learning.

We next use the structural model estimates to simulate a counterfactual benchmark against which to contrast observed firm behavior. This exercise allows us to compute the gap between full-information and limited-information prices and quantify the economic importance of learning. We hereafter refer to prices' deviations from the optimum as "mistakes," noting that they might be actual pricing mistakes due to limited information or retailers' strategic experimentation to acquire information.<sup>3</sup> We find that, at the median, prices in the first quarter are 9% above the full-information optimum and that these mistakes lead to 9.3% lower profit. However, retailers are quick to learn: they adjust prices and close one-third of the profit gap within a quarter and another one-third within the next six quarters.

We finish with some descriptive evidence on the nature of firms' pricing mistakes and the underlying learning process, framing the discussion within the context of canonical learning models. In particular, Bayesian learning implies an optimal weighting of prior beliefs and new information, with weights tied to each component's information precision (inverse variance). As a result, the Bayesian firm will lean on the prior when it is more precise (e.g., when the firm has more experience) and will lean on the signal when it is more precise. Consistent with this prediction, we demonstrate that the learning rate (the rate at which firms correct pricing mistakes) *decreases* with firms' accumulated experience in Washington, *decreases* with their experience in other markets prior to entering Washington, and *increases* with the precision of new information from realized sales quantities. Together with the finding that retailers achieve optimality, we find that the Bayesian learning framework (Hitsch, 2006; Ching, 2010) presents a credible depiction of retailer price adjustments in this new market.

Our evidence also highlights two important aspects of learning in this market. First, firm dif-

---

<sup>3</sup>Throughout we do not distinguish between mistakes and experimentation. Section 3.4 discusses this distinction in more detail.

ferences in learning can support short-term profit advantages, and such learning appears to transfer between contexts (different markets where demand primitives differ). This provides additional empirical support for the profit impact of dynamic capabilities and firm heterogeneity (Teece, Pisano and Shuen, 1997; Goldfarb and Xiao, 2011). Second, in this context at least, firms appear capable of learning a fairly high-dimensional target, namely the distribution of consumer tastes for different product types. This suggests a high degree of managerial sophistication and capability. In particular, it implies a more sophisticated process than either descriptive pricing research (Noble and Gruca, 1999) or adaptive learning models suggest, and a more complex target than canonical Bayesian learning models presume (likely reflecting model tractability).

Our paper contributes most directly to a recent firm learning literature aimed at relaxing the rational expectations assumption, as well as the Bayesian requirement that firms weight information optimally. Jeon (2020) examines firm investment decisions in the container shipping industry, where firms adaptively learn about the level of demand from historical data. Without assuming firms use optimal information weights, she finds that an adaptive learning model that overweights recent information best explains firm behavior. Doraszelski, Lewis and Pakes (2018) study firm learning about demand and equilibrium competitor behavior. They estimate a model where firms adaptively learn about demand (firms estimate demand with all available information) and learn about competitor actions using fictitious play. Covert (2015) finds that firms in the hydraulic fracturing industry overweight own information, compared to public information from other firms. Li and Ching (2021) study how Prosper.com adapts to changes in the marketplace. They consider a class of data-selection algorithms (where one uses recent data to adaptively learn about the market primitives) and find that one algorithm best explains Prosper.com's actions. Our empirical strategy differs from the above papers, in that we aim to describe how firms learn without assuming the learning mechanism (and as a result, without estimating a learning model). Instead, we measure learning through the evolution of the covariance between prices and demand primitives.<sup>4</sup> Even in the absence of a learning model, we demonstrate that how prices converge to the optimum is most consistent with canonical Bayesian

---

<sup>4</sup>Closely related to this idea, Leisten (2021) finds that chain hotels under-respond to non-salient demand shocks (although his paper does not examine learning).

updating.

Our paper is also closely related to the Bayesian firm learning literature (Hitsch, 2006; Ching, 2010; Huang, Luo and Xia, 2019), which characterizes structurally firms' dynamic entry, exit, and pricing decisions with Bayesian learning. Our finding, that firms' learning about demand is broadly consistent with Bayesian updating, provides support to these models. However, whereas the extant literature characterizes firm learning about product (or product-category) quality, our empirical finding presents a case where firms must learn about nuanced segmentation patterns of the demand function.<sup>5</sup>

Broadly, our paper also speaks to the literature on pricing practices. Several earlier studies use surveys and interviews to learn how managers make pricing decisions (see, e.g., Kaplan, Dirlan and Lanzillotti, 1958 and Noble and Gruca, 1999) and find that managers rely on suboptimal heuristics such as "cost-plus" pricing. Consistent with this view, a collection of recent empirical evidence suggest bounded-rational prices by supermarkets (DellaVigna and Gentzkow, 2019; Arcidiacono et al., 2020) and other firms (Hortaçsu et al., 2019). Our findings demonstrate that (at least in our context) supermarkets can learn to set optimal prices, but the process takes time and is heterogeneous across firms.

Finally, our paper connects to the recent literature on the regulation and privatization of the liquor industry per se. Conlon and Rao (2019) and Miravete, Seim and Thurk (2018, 2020) examine the impact of tax policy and market structure on upstream competition, optimal taxation and consumer welfare, respectively. Aguirregabiria, Ershov and Suzuki (2016) study regulation and tax regimes in the Ontario wine market. Seo (2016) and Illanes and Moshary (2018) leverage Washington State's liquor privatization to study, respectively, the value of one-stop shopping and the impact of market structure on prices and product variety.

The rest of the paper is organized as follows. Section 2 describes the institutional background, data sources, and sample construction. Section 3 provides descriptive statistics about price movements in the privatized market, and key reduced-form evidence of retailer learning. Section 4 employs a structural model to infer demand and recover costs. Section 5 uses the structural estimates to evaluate

---

<sup>5</sup>Also related to our paper, Liu et al. (2020) empirically document sellers' learning to exploit particular platform design features in the peer-to-peer lending market.

the optimality of firm behavior, explore the moderators of the learning process, and discuss what they reveal about the learning mechanism and firm capabilities. Section 6 concludes.

## 2 Context, Data, and Sample Construction

The 2012 privatization of off-premise liquor sales in Washington State provides a unique opportunity to observe how retailers learn to set prices from the point at which that learning process begins. The liquor category *itself* is not new to the market, which both limits the scope of confounding factors such as consumers' learning about their own preferences (as would occur with the introduction of entirely new products) or manufacturers' uncertainty about product quality, and provides a sharper focus for our analysis. In this section, we first provide some institutional background on the Washington liquor market both before and after the privatization, and then describe the data we assemble to examine the learning process.

**Institutional Setting.** Given its strong association with negative health and societal outcomes, the production and distribution of distilled spirits in the United States have been strictly regulated since the repeal of prohibition in 1933. The regulation of retail sales generally takes one of two forms. In “control” states, sales take place directly through a monopoly of state-owned outlets, sometimes referred to as “package stores.” To raise tax revenue and reduce consumption, most control states impose a large and fixed markup over cost in setting shelf prices, levying additional taxes per volume of ethanol. In “license” states, the government grants individual firms the right to sell alcohol in their own retail outlets. While shelf prices are generally left to the discretion of these independent firms, most states impose an ad valorem tax on revenue and a specific excise tax on ethanol to reduce consumption and collect revenue.

Prior to June 2012, Washington State was a liquor control state with a fixed retail markup of 51.9% and a specific tax of \$3.77 per liter. Following the passage of Initiative 1183, the state-owned chain was replaced by privately-owned retailers, licensed by the state. The new policy prioritized the issuance of retail liquor licenses for “existing grocery premises licensed to sell beer and/or wine,” and

further stipulated that retail licenses only be issued to outlets with at least 10,000 square feet of floor space. 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 ad valorem taxes, the specific tax remained at \$3.77 per liter.

As in other states, Washington adheres to a three-tier alcohol distribution system, in which manufacturers may sell to retailers only through licensed distributors. The new policy in Washington affects only retail distribution and does not alter the distributor side of the market. In addition, the new policy only affects liquor (i.e. distilled spirits), as licensed retail wine and beer sales were already permitted under existing law.

The focus on licensing existing grocery premises motivates our use of the Nielsen Retail Measurement Services (RMS) dataset for our empirical analysis, as this database provides good coverage of the grocery, drug, and mass merchandise channels that comprise the grocery market. Moreover, it also partly mitigates concerns regarding changes in market structure, due to the fact that these retail channels are already quite mature and well-established. We further leverage two additional institutional details in constructing our dataset and carrying out our empirical analyses. First, liquor products are visibly distinguished and marketed by category. The main liquor categories include whiskey, gin, rum, tequila, and vodka. Because there is very little substitution across product categories, pricing is generally done on a category-by-category basis (Conlon and Rao, 2019). Therefore, to maintain tractability and allow for flexible substitution across products within a category, we focus on a single liquor category: whiskey products. Whiskey is the largest category by sales volume and includes a rich variety of products that are differentiated both horizontally (e.g. bourbon, Scotch/Irish, and Canadian) and vertically (Johnny Walker Red versus Blue). Second, as has been shown to be the case with most consumer packaged goods (DellaVigna and Gentzkow, 2019; Hitsch, Hortaçsu and Lin, 2021), Washington State retailers set uniform prices across stores, within product, at a given point in time. In our later structural analysis, we will exploit this regularity when inferring costs from observed pricing behavior. Further documentation of these patterns is provided in Web Appendix W.1.

**Data.** As noted above, our primary data source is the Nielsen RMS Dataset. We focus on the whiskey category in Washington State, and primarily the period from June 2012 (the month of privatization) until the last quarter of 2016.

For some analyses, we supplement the Nielsen data from Washington State with auxiliary information from other states or data sources. First, in some descriptive analyses, we compare Washington to other license states in which retailers have up to several decades of experience with private retail grocery liquor sales. The Nielsen data include 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 the District of Columbia. Second, we use posted (control-state) retail prices from Washington and Oregon to help estimate and validate the wholesale prices inferred from our structural analysis. State-run outlets in Washington pre-privatization and in Oregon throughout the entire sample period followed fixed markup policies (with Oregon markup fixed at 104%), allowing us to directly calculate wholesale prices (retailer costs) from their posted shelf 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, as noted above, we narrow our focus to a category within liquor, the broad whiskey category. This category is sizable, has little substitution with other categories (e.g., vodka or tequila), and offers the most product diversity, including distinct whiskey types such as single-malt and blended scotch, Canadian whiskey, American bourbon and rye, and American whiskey. 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 (during which prices may be poorly measured or non-representative, due to stock-outs or close-out sales). 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 the observations from the overall sample.

Third, we filter out 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 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 W.3.2. After these sample selection steps, our sample contains 4,985,621 observations, from 276 products, six retailers, and 561 stores. In Web Appendix W.3 we present further sample adjustments to facilitate our structural analysis. Specifically, there we simplify to focus 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 (see Web Appendix W.3 for further rationale and details).

### **3 Learning about Demand: Descriptive Evidence**

In this section, we present our initial descriptive evidence of firm learning. We begin by outlining a conceptual framework that characterizes firm learning about demand. This conceptual framework yields three testable predictions, focusing on key features of the joint distribution of prices and quantities and how these sample features transition over time. We empirically examine these predictions in the post-privatization Washington and present novel empirical evidence of firm learning about demand.

We first show that prices do in fact change within the first two years after privatization and then stabilize, and that the directions and magnitudes of these changes differ across products. These patterns are consistent with retailers learning demand conditions in the early period, and with this learning

eventually ceasing. Second, we demonstrate that prices immediately begin reacting to realized quantity shocks as retailers start selling liquor in Washington, but that, over time, they adjust less and less to these shocks. Third, we show that the price of a product increasingly reflects its demand fundamentals (e.g., whether the product is especially popular or caters to a price-insensitive customer segment). This evidence suggests that retailers steadily acquire information about demand fundamentals, and the prices they set increasingly reflect their refined information. We also present placebo tests that replicate these descriptive exercises in other states in which retailers have sold liquor for many years prior to 2012, and demonstrate that these comparison states do not exhibit such evidence of learning.

Our case for learning turns on the claim that firms gradually acquire new knowledge about demand, and this new information is reflected in the changing prices we observe. It is therefore important to rule out the possibility that demand (or the environment) itself is changing, so that firms might instead be reacting to these changing conditions. If the environment exhibits ongoing nonstationarities, the price adjustments we observe could reflect firms' rational expectations adjusting to reflect these changes, rather than (initially) biased perceptions adapting through learning. To evaluate this possibility, we provide evidence that consumer tastes are not changing, that the competitive landscape and upstream (distributor) behavior are stable, and that most other retailer strategies remain relatively fixed over this period.

### 3.1 Conceptual model

To help fix ideas, we first present a simple conceptual model that motivates our initial set of testable implications. Note that this is not the structural model we later estimate, but rather a stylized representation intended to highlight the key intuitions. Consider a single firm that produces  $J$  products, with each product  $j$  facing stable demand given by the linear function

$$q_{jt} = \beta_j + \alpha_j p_{jt} + \xi_{jt} \tag{1}$$

in which each product has its own intercept  $\beta_j$  and price sensitivity parameter  $\alpha_j$ . The product-specific intercept captures each product's popularity in the new market, and the product-specific price

sensitivity captures the different customer segments each product caters to. Implicitly, there is no substitution between products (in this stylized abstraction), and prices are set independently based on the firm’s belief about each product’s demand. We further assume that the demand shock  $\xi_{jt}$  is realized after prices are set and that the marginal cost of each product is a constant  $c_j$ .

The firm forms beliefs about demand parameters  $(\beta_j, \alpha_j)$ , which are denoted  $\mathcal{F}_{jt}$ . After prices are set,  $\xi_{jt}$  is realized and the firm observes the total quantity  $q_{jt}$  at the end of the period (but does not observe the parameters of the demand function directly). Throughout the paper, we remain agnostic about how the firm updates its information set  $\mathcal{F}_{jt}$ , other than assuming that all firms eventually learn fully about demand (note that this is in line with Aghion et al., 1991 in that, if the demand function is smooth or analytic, both of which are satisfied here, adequate learning is guaranteed). In Web Appendix W.2, we develop a special case where the firm updates its beliefs  $\mathcal{F}_{jt}$  in a Bayesian manner (where it optimally incorporates all past information) and sets static-optimal prices given its belief. We also demonstrate that model predictions are similar in two alternative cases, where the firm (1) estimates demand adaptively using all available data, without a prior (as in Doraszelski, Lewis and Pakes, 2018), and (2) updates  $\mathcal{F}_{jt}$  in a Bayesian manner while also conducting price experiments.

To be clear, we do not assume that firms employ any particular form of learning in our empirical analysis. However, these numerical examples provide a point of reference to complement and motivate our empirical findings.

### **3.2 Price changes after the privatization of liquor sales**

Our first hypothesis is that, if firms learn about demand in the new market, one should expect the distribution of prices to adjust in a non-stationary manner. It is straightforward to show that our conceptual model would exhibit such a pattern. We now examine how the average Washington whiskey price changes upon privatization and over time. To facilitate comparison to other states, we focus on 65 popular 750ml products that are available both before and after privatization and compute the average sticker (pre-tax) prices using time-invariant quantity weights. The left panel of Figure 1 shows a sharp, 41% increase in the average price upon market privatization, followed by a gradual decrease

in the next two years that eventually stabilizes to a mild time trend.

[Insert Figure 1 about here]

Next, we examine whether such price changes differ across products. To account for possible (mild) cost changes we focus our analysis on the ratio of the average price of each product (product-retailer-half year) in Washington to the average across the other fifteen states,  $\tilde{p}_{jrt} = \bar{p}_{jrt}^{\text{WA}} / \bar{p}_{jrt}^{\text{other}}$ , normalizing this ratio by its initial value,  $(\tilde{p}_{jrt} - \tilde{p}_{jrt_0}) / \tilde{p}_{jrt_0}$ , in order to facilitate comparison across products and retailers. We plot quantiles of these normalized prices over time in the right panel of Figure 1. We find that the bulk of the Washington price changes relative to other states' occur in the first two years after the privatization. During this period, not only does the median price go down by more than 5%, but there is a large dispersion in how prices change for individual products. Consistent with learning about demand, the systematic departure of Washington prices from other states occurs soon after the new market opens, and observed price changes are heterogeneous across products.

### 3.3 Price response to lagged demand shocks

Our second testable hypothesis is that, 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 once they are sufficiently informed. We first illustrate how this prediction plays out in the simulated environment of our conceptual model, where the firm updates its beliefs  $\mathcal{F}_{jt}$  in a Bayesian manner (the specifics of this learning process, and alternative learning processes, are detailed in Web Appendix W.2).

Using this conceptual model, we now show that prices respond to the firm's unexpected lagged quantity  $q_{jt-1} - \mathbb{E}[q_{jt-1} | p_{jt-1}, \mathcal{F}_{jt-1}]$ , with the response diminishing over time. Initially, the demand function is uncertain to the firm. If the realized quantity is higher than expected, the firm will attribute part of this realization to a signal of higher-than-anticipated product quality (higher  $\beta_j$ ) or a less-price-sensitive customer base (higher  $\alpha_j$ ) than previously expected, and thus, update its beliefs, and raise the product's price. If the realized quantity is lower than expected, the firm should adjust its beliefs, and the price, in the opposite direction. However, as demand uncertainty resolves

through learning, the firm will eventually attribute all of the unexpected quantity realization to the idiosyncratic demand shock  $\xi_{jt}$ .

We demonstrate this intuition in Figure 2A. If the firm learns about demand over time, prices (reflective of firm beliefs) respond positively to the unexpected quantity. In addition, as uncertainty is resolved over time, prices eventually stop responding to unexpected quantity. Note that the ideal testable implication would examine realized quantities relative to the firm’s expected quantities (as above), but in reality firm beliefs are unobserved to the researcher (without further strong assumptions). As such, Figure 2B uses the same Bayesian learning example to demonstrate that simply regressing *changes* in price on *changes* in lagged quantity will yield a similar pattern, albeit with attenuated regression coefficients (as lagged quantity is effectively a noisy signal of the ideal measure).<sup>6,7</sup>

[Insert Figure 2 about here]

Turning now to the actual data, we now present evidence of firm learning in line with the intuition in Figure 2B (where we proxy the firm’s unexpected sales quantity with past quantity shocks). We demonstrate that past quantity shocks do indeed have significant initial informational content for Washington retailers and that this information effect is eventually exhausted. To document these patterns, we estimate a linear model of how the current price is correlated with the lagged sales quantity, conditional on various controls and fixed effects. Denoting  $j$  as a product,  $r$  as a retailer and  $t$  as a month, we estimate a linear model of the current price on one-month lagged quantity, controlling for current quantity and lagged prices:

$$\begin{aligned} \log(p_{jrt}) = & \beta_{\tau} \log(q_{jrt-1}) + \\ & \rho \log(p_{jrt-1}) + \alpha^{-1} \log(q_{jrt}) + \psi_{jr} + \phi_t + \eta_{jrt}. \end{aligned} \quad (2)$$

Our key parameter of interest is the sensitivity of the current price to  $\log(q_{jrt-1})$ , the units sold for

---

<sup>6</sup>The particular Bayesian learning model assumes normal prior and independent demand shocks. We also show similar patterns for adaptive (regression-based) learning and learning with epsilon-greedy type experiments. Nevertheless, lagged demand shocks are not guaranteed to be a good proxy for unexpected demand shocks in all data-generating processes.

<sup>7</sup>Note as well that, in a regression with multiple controls, the coefficient on the change in lagged quantity will capture the impact of that difference after partialling out the direct impact of the included controls (hence our use of the term “quantity shock” throughout).

product  $j$  by retailer  $r$  in the previous month,  $t - 1$ . We allow  $\beta_\tau$  to take a different value for each half-year. Our controls capture stable product-retailer characteristics ( $\psi_{jr}$ ), common variations within a month ( $\phi_t$ ), serial correlation in prices ( $\rho$ ), and the inverse demand ( $\alpha^{-1}$ ). For estimation, we take first differences within each product-retailer and employ the standard instrumental variable approach from the dynamic linear model literature (Arellano and Bond, 1991).<sup>8</sup>

Figure 3 plots the estimates of  $\beta_\tau$  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 realized sales quantities. This pattern for Washington State stands in sharp contrast to what we find for the control states where retailers already have extensive experience in the liquor market. In these states, the responsiveness to sales shocks is much smaller, centered close to zero, with most coefficients statistically insignificant. Hence, for retailers with greater experience, prices do not respond to past sales shocks.

[Insert Figure 3 about here]

### 3.4 Prices’ correlation with demand fundamentals

Our third testable hypothesis is that, if firms start with incorrect beliefs about demand but learn over time, these beliefs, and the prices they set based upon them, should increasingly reflect (true) demand fundamentals. Relative to a given prior belief, prices for products that have (surprisingly) high demand will tend to rise to reflect that higher demand, and prices for (surprisingly) low demand products will fall. The end result is that prices charged and quantities sold of products should become increasingly correlated as retailers learn about demand. Hence, learning should be reflected in a clear pattern of increasing covariation between prices and the underlying demand fundamentals.

---

<sup>8</sup>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  $\eta_{jrt}$ , for instance, arising from retailers being unwilling to change price right after a previous price change due to menu costs. In Web Appendix Table 7, we present further details of the estimation results that support the use of the instruments and controls.

In the conceptual model of equation (1), the static optimal price given belief  $\mathcal{F}_{jt}$  is  $p_{jt} = -\frac{\mathbb{E}[\beta_j|\mathcal{F}_{jt}]}{2\mathbb{E}[\alpha_j|\mathcal{F}_{jt}]} + \frac{c_j}{2}$ , a simple function of the firm's belief about demand ( $\mathbb{E}[(\beta_j, \alpha_j)|\mathcal{F}_{jt}]$ ). If retailers initially set prices based on noisy beliefs over the demand primitives, their initial prices will only partly reflect these constructs. However, if their beliefs in fact converge towards the truth via learning,  $\mathbb{E}[(\beta_j, \alpha_j)|\mathcal{F}_{jt}] \rightarrow (\beta_j, \alpha_j)$ , observed prices will increasingly capture these demand primitives. Our second numerical exercise illustrates that, in a case where the retailer learns about the linear demand function via Bayesian updating, observed prices will become increasingly correlated with the optimal markup  $-\frac{\beta_j}{2\alpha_j}$  (eventually reaching the limit of 1). If the researcher knows demand primitives and can compute optimal prices, the way prices converge to optimal prices (or markups) provides a direct test of learning, as demonstrated in Figure 4A.

[Insert Figure 4 about here]

The conventional approach is to estimate the demand primitives and evaluate the extent to which prices capture these primitives. We present a new descriptive test that does not rely on ex ante demand estimates. 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. For example, with the linear demand in equation (1), a cross-sectional regression would group product-specific quality and price sensitivity into the error term:

$$q_{jt} = \bar{\beta}_t + \bar{\alpha}_t p_{jt} + \underbrace{\beta_j - \bar{\beta}_t + (\alpha_j - \bar{\alpha}_t) p_{jt}}_{\text{error term, correlated with price}} + \xi_{jt}. \quad (3)$$

When the firm sets profit-maximizing prices that incorporate some (though not necessarily all) information about the demand primitives, one should expect prices to reflect underlying product quality or positioning (Trajtenberg, 1989), and, as a result, estimates that rely on cross-sectional variation in price and quantity between products will yield a price-sensitivity parameter estimate that is biased upward. In addition, as the firm learns about demand, its beliefs should reflect the true demand fundamentals more strongly, intensifying the price-endogeneity bias in the cross-sectional regression. Therefore, if firms learn about demand over time, the full extent of bias will not show up right away,

but will progressively develop as prices gradually incorporate the demand fundamentals. When learning is complete, this correlation should stabilize, as will the estimates of these (now fully biased) coefficients. Conversely, if learning does not occur, this correlation should instead be stable from the start. In the simulated example where the firm learns via Bayesian updating, we report the estimated  $\hat{\alpha}_t$  over time. Figure 4B shows that these estimates reflect an increasing endogeneity bias (the true average price sensitivity is -0.5) and closely track the pattern in panel A, where the researcher observes the true demand (and optimal markups).

We now implement this test in our empirical context. Using data at the product-retailer-week level, we regress quantity on price, estimating the price coefficient  $\bar{\alpha}_\tau$  for each six-month time window ( $\tau$ ), while controlling for retailer-week fixed effects but not product fixed effects. Hence, following the classic omitted-variable bias formula, the estimated price coefficient  $\hat{\alpha}_\tau$  will be the true slope of demand plus any bias due to the way prices capture unobserved demand primitives.<sup>9</sup> We report the estimated  $\alpha_\tau$ 's and corresponding standard errors for the Washington sample in Figure 5. We also examine retailers in other states as a placebo test, but with one change to account for including multiple states: we include retailer-state-week fixed effects.

In Washington, we find that retail prices show increasing correlation with demand primitives, reflected by the way  $\hat{\alpha}_t$  increases from 2012 to 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,  $\hat{\alpha}_t$  stabilizes in the second half of the sample, suggesting that most of learning occurs in the first two years after the market opens and that, after this initial period, there is no systematic learning that changes how prices reflect demand fundamentals. In contrast to the finding in Washington, we find that other states'  $\hat{\alpha}_t$ 's are stable over the full time period, consistent with experienced retailers in the existing liquor markets being fully informed (i.e. no longer learning about demand in systematic ways).

[Insert Figure 5 about here]

---

<sup>9</sup>The classic bias formula establishes that the focal coefficient in the “short” regression (that omits a key confounder) equals the corresponding coefficient in the “long” regression (that includes it) plus the effect of the omitted times the regression coefficient of the omitted on the included. The pattern documented here arises from the increasing impact of the omitted on the included.

A slightly different explanation of the above finding is that retailers strategically experiment with prices to acquire information. These experiments add noise to prices and result in the initial  $\hat{\alpha}_t$  being less biased. As retailers become more informed about demand, they experimentally vary prices less, resulting in a larger bias in  $\hat{\alpha}_t$  for the later periods. As discussed in the introduction, this paper does not distinguish whether retailers learn by passively observing demand or by strategic experimentation to acquire information. Both types of learning are consistent with this evidence.

Throughout Section 3.2-3.4, we have demonstrated that Washington liquor prices evolve in a non-stationary fashion, are positively correlated with realized sales shocks, and increasingly reflect sales levels for each product. We have also demonstrated that these patterns occur in Washington but not in other states, and mainly in the first two years after the start of the new market. These patterns collectively suggest that retailers learn about demand and adjust their prices accordingly in the new Washington market.

### **3.5 Alternative explanations to learning about demand**

We claimed at the outset that the most probable source of the non-stationary movements in prices we observe in Washington State is retailers' initial uncertainty over *stable* demand primitives. We turn now to justifying this assumption. Of primary concern is the possibility that other features of the environment could be changing at the same time, and thus confounding the main evidence of learning – that the joint distribution of price and quantities exhibit non-stationary movements in the early periods of post-privatization Washington. In this section, we consider four such confounds, namely: (1) changes in consumer demand, (2) changes in the competitive market structure, (3) changes in distributors' behavior, and (4) changes in other retailer strategies. We account for the impact of cost changes later (in Section 4.4). Here, we provide evidence that these alternative mechanisms are not of first-order importance in our setting. This section outlines the main findings, with Web Appendix W.3 providing further detail.

**Stable consumer demand and store traffic.** First, it is important to rule out systematic changes in consumer behavior post-privatization. On the one hand, if consumers gradually learn about the

set of products or search for low-price products, they might *appear* to be increasingly price-sensitive over time. In this case, the observed price adjustments might be responses to *changes* in consumers' price sensitivities, instead of responses to firms' having better information about demand. To address this possibility, we estimate a simple log-log demand curve with a rich set of fixed effects. We find little changes in the slope of demand before and after 2014, suggesting that consumers' price sensitivity remains stable. We also examine whether demand levels are different over time, using our later structural estimates. We find little variation over time both in the demand intercepts and in the variance of demand shocks. We present details of these analyses in Web Appendix W.3.1.

On the other hand, optimal liquor pricing might be dynamic if consumer demand is inertial (Dubé, Hitsch and Rossi, 2010), or if they take advantage of sales to stockpile liquor (Erdem, Imai and Keane, 2003; Hendel and Nevo, 2006), or if liquor prices affect persistent consumer store traffic. Using Homescan panel data, we find no evidence of first-order state dependence (see Web Appendix Table 8), little evidence in support of consumer stockpiling (see Web Appendix Table 1), and no evidence that some retailers' carrying liquor in 2012 affects consumer store traffic (see Web Appendix Table 3). Based on these findings, optimal prices should maximize *static* liquor-category profit for each retailer.

**No retail competition in liquor sales.** Next, it is also important to examine the extent of competition in the Washington retail liquor market, and potential changes in the competitive market structure. If retailers face intense competition in the liquor market, they might be learning about competitor behavior instead of consumer demand. We begin by demonstrating that liquor customers are part of the installed grocery customer-base (instead of a stand-alone market). We further show in Web Appendix W.3.3 that the demand for a given product has negligible substitution between local retailers – consistent with Illanes and Moshary (2018) who find that local retail market structure does not affect Washington liquor prices. This evidence suggests that competition between chains is not a salient aspect of pricing in the liquor market. We proceed by assuming that retailers are aware of this fact from the start and do not use price-setting to learn about competitor behavior (or in response to it).

**Passive distributors.** Third, one might wonder whether liquor distributors change their behavior after the retail privatization. Whereas we do not have data on actual wholesale contracts, we conducted a series of interviews to better understand the degree of coordination and strategic interaction along the vertical channel. The complex nature of the three-tier system and its challenges for upstream information flow, in particular, came up repeatedly in these interviews. A common theme was the passive nature of the middle tier (distributors), which apparently follow simple pass-through strategies that follow the lead of the (more informed) manufacturers. For example, one respondent characterized the distributors as “not very economic or demand driven.”

We further examine the impact of Washington State wholesale liquor market’s privatization in January 2012 (Senate Bill 5942). We explore how this change in the upstream affects the state retailer’s wholesale prices up to the point of retail privatization. If distributors actively set wholesale prices based on market conditions, observed wholesale prices should exhibit a discrete jump after the distribution system is privatized and competing distributors enter the market. We demonstrate in Web Appendix Figure 14 that the average wholesale price ticks up by about 2% in 2012 (compared to 2011), a modest change compared to the 41% retail price jump at the point of retail privatization. We interpret the result as consistent with our interviews with practitioners who describe distributors as passive.

**Other retailer strategies: promotion, feature, display, and assortment.** Finally, other retailer strategies might change as they learn about consumer demand. We document a lack of systematic changes in price-promotion strategies during the sample period (see Web Appendix W.3.2), suggesting that the observed systematic changes in prices mainly come from regular prices. We also show that feature promotion strategies remain stable, although retailers do put fewer products on display in the second half of the sample. Nonetheless, our estimated demand suggests that feature and display play a minor role in explaining consumer demand. 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

and profitability.

## 4 Model: Demand and Costs

Having presented several pieces of model-free evidence that prices are updated in a manner that increasingly reflects stable underlying demand primitives, but eventually settles down to a stable process, we now propose a structural model of consumer and firm behavior. The goal of this modeling framework is to both provide additional evidence of firm learning and quantify its economic importance, as well as to further illuminate the mechanisms by which learning occurs and the moderators of its success. We start by proposing a model of consumer demand, which we use to recover the primitive components of consumer preferences that firms are purportedly learning about. Importantly, estimating the demand system characterizes the target of learning, and does not require a model of firm behavior.

We then propose a benchmark supply model of how firms *should* behave under full-information. We leverage this model and the fact that prices appear to converge by the end of the sample period to infer the implied costs that those firms must be facing. In doing so, we *assume* that firms are fully-informed rational maximizers at the end of our observation period (but not necessarily earlier). We later *validate* this assumption by comparing the implied costs to the best available proxies. We find that these measures accord closely. This close correspondence indicates that firms eventually learn to set prices (near) optimally. Having concluded that firms indeed learn to price, we then turn, in the two remaining sections, to quantifying the economic importance of learning, exploring how the initial stock and later flow of information shapes the learning process, and identifying the key features of demand about which firms appeared to be most mis-calibrated.

### 4.1 The demand for liquor

We assume that demand for liquor is characterized by a standard random coefficient logit model, and estimate its parameters via established nested fixed point methods that incorporate micro-moments to identify heterogeneous demand parameters (Berry, Levinsohn and Pakes, 1995; Nevo, 2001; Petrin,

2002). We allow preferences to be heterogeneous along two dimensions. First, to accommodate consumers' self-selecting into a broad range of vertically-differentiated whiskey products, we allow for heterogeneity in the price coefficient (operating partly through consumer income). Second, since the whiskey category is comprised of several distinct whiskey *types* (scotch whiskey, American bourbon and American whiskey, and Canadian whiskey) that are clearly targeted to different demographic groups, we allow the tastes for these whiskey types to also vary by income.

We estimate liquor demand at the level of retailer  $r$ , three-digit zipcode (market)  $m$ , and month  $t$ . Consumer  $i$  comes to the retailer to buy groceries, and derives utility from purchasing whiskey product  $j$ :

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

In the above,  $\gamma_{ki}$  represents whiskey type  $k$  and household  $i$  specific intercept, where  $k$  takes 1 (American bourbon and rye), 2 (Canadian whiskey), or 3 (Irish whiskey and scotch). The baseline type is American whiskey (e.g. Jack Daniels).<sup>10</sup> The parameter  $\alpha_i$  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.  $\delta_{jrm}$  are product-retailer and retailer-market fixed effects capturing tastes for different products, which could differ across shoppers frequenting different retailers or residing in different markets.  $\lambda_{rt} = \lambda_{ry}^1 + \lambda_t^2$  are retailer-year and month fixed effects, which capture changes over time in the demand for liquor in grocery stores.  $\xi_{jrmt}$  captures unobserved characteristics or demand shocks.  $\varepsilon_{ijrmt}$  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, in line with our supporting evidence (section 3.5) that consumers do not substitute between local retailers (and hence, retailers effectively act as monopolists over their own store traffic). Given this structure, the market share within a retailer-

---

<sup>10</sup>American whiskey is a small category. As such, we do not model household-specific tastes ( $\gamma_{ki}$ ) for this category. Note that we do control for product-retailer fixed effects for all products.

market is an integral of logit choice probabilities over the included random coefficients,<sup>11</sup>

$$s_{jrm t} = \int s_{ijrm t} dF(\alpha_i, \gamma_i) = \int \frac{\exp(\gamma_{ki} + \alpha_i p_{jrm t} + x_{jrm t} \beta + \delta_{jrm} + \lambda_{rt} + \xi_{jrm t})}{1 + \sum_{j' \in J_{rmt}} \exp(\gamma_{k'i} + \alpha_i p_{j'rm t} + x_{j'rm t} \beta + \delta_{j'rm} + \lambda_{rt} + \xi_{j'rm t})} dF(\alpha_i, \gamma_i). \quad (5)$$

As noted above, a key feature of the whiskey category is sharp (horizontal and vertical) differentiation by spirit type. To capture this aspect in the model, we parameterize the household's preferences towards whiskey type  $k$  as a function of log household income,<sup>12</sup>

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

noting that the average whiskey-type-level intercept is absorbed by the included 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  $\xi_{jrm t}$  are conditional mean zero, given instruments  $z_{jrm t}$  (including non-price covariates, fixed effects, and excluded instruments for price and random coefficients):

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

Price endogeneity might arise for a number of reasons including measurement error (which arises from the aggregation of prices over weeks and stores) and price-setting based on the unobserved characteristics,  $\xi_{jrm t}$ . In section 3.4, we argued and demonstrated that learning creates a pattern increasing

---

<sup>11</sup>Market share  $s_{jrm t}$  is observed quantity divided by market size. We assume that the market size is the total population multiplied by the retailer's grocery trip share (among focal retailers, computed from Homescan data) multiplied by two. Web Appendix Table 9 presents robustness checks of the Berry (1994) logit model estimates with various alternative market size assumptions.

<sup>12</sup>This income-type interaction term follows, but enriches, the specifications in Conlon and Rao (2019) and Miravete, Seim and Thurk (2018).

covariations between prices and the unobserved characteristics as retailers learn the demand primitives. This implies we need instruments that exclude both potential correlated measurement errors, as well as the shocks and unobserved qualities that firms learn about. We construct price instruments similar to Conlon and Rao (2019) and Miravete, Seim and Thurk (2018) to address these concerns. 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 common wholesale price variations across states but do not correlate with either demand shocks or the learning occurring in Washington (after controlling for the above fixed effects). One example of wholesale price co-movement is that prices of Scotch whiskey move with the USD-GBP exchange rate, as illustrated in Web Appendix Figure 13. Also, to the extent that measurement errors from aggregating prices across weeks and stores, the instruments constructed from other retailers' prices should have independent measurement errors.

**Random coefficients.** We identify the random coefficients by combining the instruments suggested by Berry, Levinsohn and Pakes (1995) with additional micro-moments that inform substitution patterns (Petrin, 2002) that are inspired by the constraints used in Conlon and Rao (2019). 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 disproportionate substitution to other products versus to the outside option captured by type-intercepts  $\gamma_{ki}$ .

Second, we construct five sets of micro-moments, using the Nielsen Homescan panel data, to help identify how the type 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

whiskey categories among purchases. For each set of parameters  $\theta$ , we match the model 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 type, and  $j \in k$  represents product  $j$  falling into type  $k$ . As in Petrin (2002), we match the observed overall and type-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 do not report their coefficients in the table. For comparison, 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 the magnitude of the latter. Part of this heterogeneity is driven by income, indicating that high-income consumers are less price-sensitive. Consistent with Conlon and Rao (2019), 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 via the Berry, Levinsohn and Pakes (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.

**Implied elasticities.** 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 the 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 similar to those reported by Miravete, Seim and Thurk (2018), who find elasticities in the Pennsylvania liquor market to be -2.9 for “cheap” products and -4.9 for “expensive” products. Intuitively, this pattern arises 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.

[Insert Table 2 about here]

Thus far, we have not assumed anything about what firms know or how they behave. We have simply recovered the parameters that govern the demand system (about which the firms may be incompletely informed). We next leverage a model of firm pricing behavior in order to further explore the role and nature of learning.

#### **4.4 Wholesale prices (retailers’ marginal costs)**

We now present a supply-side model to recover wholesale prices (retailer marginal costs). The goal of this exercise is two-fold. The first goal is to provide an indirect test of long-run optimality – by

comparing the marginal costs implied by the model to our best available proxies. The second goal is to measure the economic impact of firm learning, the moderators of learning, and to evaluate what demand fundamentals retailers learn about.

This process has two distinct steps. In the first step, we assume the optimality of firm pricing behavior in the *last six months* of our observation window, when we previously identified learning as appearing to be complete.<sup>13</sup> From this step, we obtain cross-sectional variation in the implied retailer-product wholesale prices. The second step imposes a transition process for the wholesale prices and leverages data from Oregon to model the time path of wholesale price changes back to the beginning of the privatization period. We now discuss each step in turn.

**Recovering post-learning wholesale prices under optimality.** To recover wholesale prices, we assume that retailers are fully informed about demand in the *last six months* of the sample. During this period, the fully informed retailer  $r$  sets prices for its products based on demand primitives,<sup>14</sup> with the restriction that the price must be set uniformly for each product across all markets in Washington. Specifically, the retailer, as a multi-product and multi-market monopolist, chooses the vector of prices,  $p_{rt}$  to maximize the total profit in month  $t$ ,

$$\pi_{rt}(p_{rt}) = \sum_{j \in J_{rmt}} \sum_{m \in M_r} ((1-f)p_{jrt} - c_{jrt}) \cdot \tilde{s}_{jrmt}(p_{rt}) \cdot h_{rmt}, \quad (12)$$

where  $\tilde{s}_{jrmt}$  denotes the model-predicted market shares (using an AR(1) process for the demand shock, see footnote 14),  $M_r$  is the set of markets (3-digit zipcode) it operates,  $h_{rmt}$  is the size of market  $m$  in month  $t$  for the retailer (which is local population times the retailer's share of grocery revenue in the market),  $c_{jrt}$  is the marginal cost for product  $j$  in  $t$ , and  $f = 0.17$  is the share of gross revenue

<sup>13</sup>In their analysis of learning in the UK electricity market, Doraszelski, Lewis and Pakes (2018) similarly find that prices converge to a rest point consistent with equilibrium behavior. In a stable environment, such rest points are the result of adaptive or optimizing behavior that converges to the solution of a constrained maximization (for a single agent) or system of first order conditions (for equilibrium play).

<sup>14</sup>The retailer (who has completed learning) also projects demand shocks  $\xi_{jrmt}$  using an AR(1) process,  $\xi_{jrmt} = \rho \xi_{jrmt-1} + \iota_{jrmt}$ , after observing the realized  $\xi_{jrmt-1}$ . The full projection problem requires a computationally complex integration over many random variables (the vector of innovations  $\iota_{rmt}$  for all products). Instead, we assume the retailer projects demand shocks based on  $\hat{\xi}_{jrmt} = \hat{\rho} \xi_{jrmt-1}$ . Empirically, we find that  $\hat{\rho} = 0.651$  (standard error = 0.002). Hitsch (2006) sets the entire  $\tilde{\xi}_{jrmt} = 0$  and finds it does not affect the results meaningfully. We relax his assumption, setting  $\tilde{\iota}_{jrmt} = 0$ .

levied by the state, which is included in the list price. This profit maximization problem leads to the first-order condition such that for all  $j$ ,

$$p_{rt} = \frac{c_{rt}}{1-f} - (\Delta_{rmt})^{-1} \sum_m \tilde{s}_{rmt}(p_{rt}) \cdot h_{rmt} \quad (13)$$

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

We impose this optimality condition only for the last six months of the sample and calculate the implied markups  $(\Delta_{rmt})^{-1} \tilde{s}_{rmt}$  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 margins as  $(1-f)(\Delta_{rmt})^{-1} \tilde{s}_{rmt}$  divided by price (see Table 2 for examples). For example, for product 3, the retailer sets price at \$14.53 and gains 30.7% of the price as gross margin. These estimated margins indicate that retailers set lower percentage margins for high-priced products. This particular, but well-known, feature of liquor demand arises from the increasing elasticities in price.

**Inferring wholesale prices over the learning period (pre-optimality).** We now infer wholesale prices while the retailer is still learning (i.e., from before the six-month window that we assume is optimal). Our approach has three key elements: we recover and employ the final (cross-sectional) average costs, we assume a cost transition process, and we then leverage Oregon prices to estimate the *changes* in wholesale prices. The underlying cost function we assume takes the following form

$$\log(c_{jrt}) = \bar{c}_{jr} + \tau_{k(j)y(t)} + \omega_{jrt}. \quad (14)$$

The log wholesale price is composed of a component that is constant over time, but varies across products and retailers ( $\bar{c}_{jr}$ ), a component that captures common type-year variation ( $\tau_{k(j)y(t)}$ ), and an unobserved mean-zero cost shock,  $\omega_{jrt}$ .<sup>15</sup>

Allowing for type-year cost variation captures differences in exchange-rate movements. For exam-

---

<sup>15</sup>We estimate  $\bar{c}_{jr}$  as the average log wholesale price using 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, normalizing  $\tau_{k(j)y(t)} = 0$  for the period of 2015-2016. Thus, we use the Oregon wholesale prices to construct the average percentage changes in wholesale prices. Our estimate of the log wholesale prices for each product and time period is  $\bar{c}_{jrt} = \exp(\bar{c}_{jr} + \tau_{k(j)y(t)})$ .

ple, Canadian dollars have depreciated during the sample period (whereas British pound and the euro remained stable), potentially leading to different import-price changes for Canadian whiskey. Also, equation (14) assumes that wholesale-price changes are common between Washington and Oregon. To provide support for this assumption, Figure 6 takes a fixed set of products and compares their average prices (de-measured) between the Washington state retailer (before June 2012), the Oregon state retailer, and other liquor-privatized states in the RMS data. The average price series in different markets track each other closely.

[Insert Figure 6 about here]

The recovered wholesale prices have large dispersion across different products; however, such dispersion is expected given the large degree of vertical segmentation 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.

## **5 Empirical characterization of the learning process**

In this section, we use our estimates to empirically characterize firms' pricing strategies during the learning phase. We first examine whether learning is completed during the sample period, and if so, *how much prices and profits differ* between the initial, limited-information state and the eventual, full-information one. The goal of this exercise is to examine the extent of deviation from the full-information optimum and whether this deviation is transitory or persistent. In section 5.1, we conclude that prices are in line with full-information behavior by the end of the sample by verifying that model-implied marginal costs (assuming optimal behavior by the end of the sample) are consistent with our best-available proxy. Section 5.2 then shows that initial prices are systematically different from the optimum, leading to an average 9.3% gap between realized and optimal profit at the outset. Systematic, suboptimal pricing is present and economically important, but also transitory, due to learning.

Having established that prices eventually converge to full-information behavior, section 5.3 examines *how prices evolve* from the initial, limited information state to the eventual, full-information one.

We present descriptive evidence that firms’ initial behavior and learning rates are consistent with the predictions of classical learning models: learning occurs more quickly for products for which sales are more informative and for firms that are less informed about the market. In both cases, the marginal impact of additional learning is higher, in line with canonical (Bayesian or adaptive) learning models.

Finally, we attempt to pin down the object of firm learning and further unpack how firms leverage prior experience and earlier market conditions to shape the learning process. In particular, we demonstrate that learning in this context focuses on Washington consumers’ taste distribution, which differs systematically from other, previously privatized states. We further show that all retailers make similar initial mistakes consistent with their using other states’ consumer behavior as exemplars for the Washington market. However, the firm with direct experience in other markets initially outperformed those that did not. This evidence suggests that, instead of simply copying other states’ prices (which would not have suited the Washington market), the retailer was able to more quickly infer the unique aspects of Washington demand. In other words, this retailer seems to be solving a fairly nuanced (and counterfactual) inference problem.

Throughout this section, we focus on the three grocery retailers that remained stable throughout our observation period – Retailer 158, 182, and 32 – and that account for 80.7% of revenues. Retailers 152 and 182 operate only in states where liquor cannot be sold in grocery stores, meaning that they have no direct experience with this category. Retailer 32, in contrast, had extensive experience selling liquor in the other states in which it operates, before starting the sale of liquor in Washington.

## **5.1 Evidence that retailers have learned to set optimal prices**

Our wholesale price recovery imposes that retailers are fully informed and set static-optimal prices in the final six months of the sample (four years after their entry into the market). In theory, a monopolist retailer will *eventually* learn about demand if the inference problem is bounded and sufficiently “well-behaved” (Aghion et al., 1991).<sup>16</sup> In the empirical literature, this assumption is often either imposed explicitly (Doraszelski, Lewis and Pakes, 2018) or implied by the specific learning process invoked

---

<sup>16</sup>Specific to this context, the requirements for adequate learning to obtain are that the profit function is quasi-concave, and the firm use local experiments to figure out the direction in which it should adjust prices.

(Hitsch, 2006; Ching, 2010). Thus far, we have demonstrated that learning reaches a stable rest point after a couple of years, consistent with the conclusion that learning is complete by the end of the sample. We now empirically *validate* this claim, establishing a central component of our conjectures regarding learning: Upon reaching this rest point, retailers have learned to price optimally.

We compare our wholesale price estimates  $\bar{c}_{jrt}$  to two sets of direct measures,  $\tilde{c}_{jt}$ , recovered from publicly available posted prices under state regulated fixed-markup regimes. Washington State imposed a fixed 51.9% markup pre-privatization and Oregon a 104% markup, which we use to compute  $\tilde{c}_{jt}$  (for each state, over different time periods) from posted retail prices. If Washington private retailers face similar wholesale prices as the (pre-privitization) state chains, full-information optimal pricing implies that  $\bar{c}_{jrt}$  – estimated under the assumption that learning is complete by the end of the sample – should be similar to both proxies  $\tilde{c}_{jt}$ . However, because manufacturers should respond to the change in retail markups from the pre-privatization fixed-markup rule to the post-privatization flexible rule,  $\tilde{c}_{jt}$  should not be the same as  $\bar{c}_{jrt}$ . In particular, Miravete, Seim and Thurk (2020) characterize manufacturer and retailer pricing decisions and show that retail and (manufacturer-set) wholesale markups are strategic substitutes, and thus one should expect lower (higher) wholesale prices for products where retail markups increase (decrease) the most. Hence, in comparing the wholesale prices, we seek to find consistency between pre-privatization and implied wholesale prices, up to a rotation.

We evaluate the relationship between  $\bar{c}_{jrt}$  and  $\tilde{c}_{jt}$  for two distinct sample frames. The first sample is from Washington state in the pre-privatization period, and the second is from Oregon for the contemporaneous period.<sup>17</sup> We show in Figure 7 that the estimated wholesale prices are indeed close to those backed out from the Washington state data. The median wholesale price is \$12.1 in the privatized market, almost identical to the median wholesale price observed in the pre-privatization era at \$12.4. In addition, across products, the model-implied and state-observed wholesale prices have a correlation coefficient of 0.99. Further, visual inspections suggest that the pre-privatization wholesale prices and the corresponding model-implied ones are closely related, with the expected slight rota-

---

<sup>17</sup>Only a subset of products are available in both locales and time periods. Our hand-matching process identified 46 products that were available at the Washington state-owned retailer (pre-privatization), Washington private retailers (post privatization), and Oregon state-owned retailer (throughout the sample). This set of product is held fixed throughout all comparison periods.

tion. The alternative comparison of implied wholesale prices to observed wholesale prices in Oregon provides similar results (see Web Appendix Figure 17). The main difference is the degree of rotation, which one would expect to differ given the large difference in fixed-markups imposed in each of the two states.

[Insert Figure 7 about here]

An alternative approach is to use the state-observed wholesale prices to proxy retailers' costs. For the subset of 46 products available both before and after the privatization and in Oregon, we take their state wholesale prices as the baseline and add Oregon's whiskey type-year trend (as in equation (14)) to project wholesale prices to the entire sample period. We then compute optimal prices for these products and contrast them with observed prices. Web Appendix Figure 18 shows a "mirror-image" result to Figure 7: The observed and model-implied prices track each other closely except for a similar rotation, consistent with manufacturers lowering wholesale prices for the high-end products due to their lower percent-retail-markup.

These results provide external validity that the model effectively recovers the wholesale prices that represent the key marginal costs for these retailers. More importantly, we have demonstrated that, after operating in the market for a few years, retailers behave in ways consistent with full-information optimal pricing. The fact that retailers reach optimal pricing outcomes by the end of the observation period places a boundary on the types and extent of behavioral departures from rationality that are at play in this market. In particular, they are at most transitory in nature and generally self-correcting. Having shown that firms learn to price, we turn now to quantifying the economic importance of their learning task.

## 5.2 The economic importance of learning

**Measure of limited-information behavior and outcomes.** We now quantify the economic importance of learning, contrasting observed pricing decisions with the full-information prices implied by the model. Specifically, we use the structural model to simulate a counterfactual in which retailers have full information about demand and set prices optimally (at all points in time). This counterfactual

serves as a benchmark to evaluate how close observed retailer behavior is to optimal decision-making under perfect information (at any given point in time). To solve for model-implied optimal prices, we use our supply-side estimates to compute the implied marginal costs,  $\bar{c}_{jrt}$ . We solve for product-specific optimal prices for each of the three retailers in each month, assuming each product's price is set uniformly across the state.<sup>18</sup> 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 (12)).

We contrast the fully-informed, optimal prices with the observed ones, using the following measure of the “percent price gap”:

$$\% \text{price gap}_{jrt} = \frac{p_{jrt}^* - p_{jrt}}{p_{jrt}^*}. \quad (15)$$

Similarly, 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.

**Quantifying the economic importance of learning.** Figure 8 presents, across all retailers and for each quarter,<sup>19</sup> the gap between observed and optimal prices, along with the corresponding gap for total profits. Observed prices in the first quarter differ markedly from the full-information optimal prices, resulting in considerably lower profit for this first period. At the median, prices are 9.3% higher than the optimal level, with large dispersions (quartiles are [0%, 21%]). This departure in prices leads to a 9.3% lower profit compared to the optimum, suggesting that retailers' limited information about the new market is consequential for setting prices – the scope for learning about demand is high.

The gap between observed and optimal prices decreases steadily over time, with most of the adjustment occurring in the early periods. By the second quarter, the overall price level is much closer to the optimum: 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 foregone profits are recovered from learning in quarter one.

<sup>18</sup>Specifically, 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 (13). The demand shocks are set to  $\hat{\rho} \hat{\xi}_{jrmu-1}$ , the predictable part of the demand shock, noting that the innovation term  $\iota_{jrmu}$  is very small. To find the optimal prices, we iterate on equation (13) 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.

<sup>19</sup>Here, we define quarters as the periods of June to August, September to November, December to February, and March to May. We shift the standard definition forward a month to keep June 2012 in the first quarter of the sample.

By early 2016 (before we assume optimal pricing), the profit gap closes by 7.4 percentage points, which is 80% of the initial gap. This further improvement results less from average price changes and more from shrinking the variation in the price gap distribution, with the price gap interquartile range shrinking to [-2%, 3%] by 2016. This finding suggests that retailers learn about demand and improve pricing, and do so at a quicker rate at the beginning than the end. The result is a significant improvement in variable profits that, together with the fact that demand is stable over this period, suggests that firm beliefs are initially miscalibrated but gradually refined to reflect the true market conditions.

[Insert Figure 8 about here]

### 5.3 The pattern of learning

Having established that firms are initially mis-informed about demand in this new market, but gradually learn from experience, a natural follow-up question is whether the observed learning process follows patterns consistent with predictions from the extant learning literature (Hitsch 2006; Ching 2010). Our approach is to descriptively characterize the firms' initial mistakes and learning rates (the initial absolute %price gap and the rate at which it declines over time) and examine whether the basic patterns follow those implied by classic learning models.

In particular, we present evidence that is in line with two key prediction of Bayesian learning: Learning occurs at a higher rate a) when the firm is less informed (in the initial period or with less prior experience) and b) when the firm obtains stronger information (i.e., observes stronger signals). Further suggestive evidence points to the nature of initial mistakes, and the role knowledge transfer from prior operations in other states to the new Washington market.

**Observation 1: decreasing rate of learning.** We first document that learning occurs at the highest rate when firms initially enter the market, and this rate declines as firms become more informed. Figure 8 has already revealed that prices start far away from the optimum, quickly become much

closer to it in the first few quarters, and continue moving toward the optimum at a slower pace.<sup>20</sup> This pattern is consistent with either a canonical Bayesian learning model or an adaptive learning model where firms estimate demand using all available data (e.g. Doraszelski, Lewis and Pakes, 2018): as the firm accumulates information, the marginal contribution of additional signals declines.

**Observation 2: learning rate increases with precision of signals.** Our second observation is that the learning rate increases with the precision of demand signals. We define the learning rate as the rate at which prices converge to optimal price levels, using the price gap metric from equation (15). We measure the extent of initial mistakes and the learning rate for product group  $g$  (here,  $g$  refers to high- and low-volume products) by estimating the following linear model over the first year of the sample,

$$|\%price\ gap|_{jrt}|_{j,t \in g} = \gamma_{0g} + \gamma_{1g} \cdot t + v_{jrt}, \quad (16)$$

where  $\gamma_{0g}$  is the initial absolute percent-price gap and  $\gamma_{1g}$  measures how much this gap shrinks per unit of time.

We examine whether learning occurs at a higher pace for products whose average sales quantity is above the median.<sup>21</sup> Average sales volume is a natural proxy for the informativeness of quantity signals because (as we demonstrate in Web Appendix Table 6) high-volume products are carried by more stores and their demand shocks  $\xi_{jmrt}$  exhibit lower variance.

Table 3 demonstrate that the absolute %price gap shrinks at almost double the rate for high-volume products than for low-volume products. This evidence suggests that the learning rate is increasing in the “flow” of new information (which retailers obtain from sales-quantity data), a standard implication of both Bayesian and adaptive learning models. Web Appendix Figure 20 provides a more detailed picture, documenting how the whole distribution of price gaps changes over time.

[Insert Table 3 about here]

---

<sup>20</sup>Formally, initial prices are 12.6% away from the optimal price (by measuring the absolute %price gap). This price gap declines by 5.1% in the first year, 3.4% in the second year, 2.3% in the third year, and 0.7% in the fourth year.

<sup>21</sup>To avoid creating a confound between the definition of these groups and 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).

**Observation 3: initial pricing mistakes reflect Washington’s distinct taste distribution.** Before diving deeper into the learning process, it is helpful to discuss the nature of the initial pricing mistakes we observe. We now show suggestive evidence that these initial pricing mistakes are most consistent with retailers having limited information about customer preferences over whiskey types. We arrive at this conclusion in two steps.

First, using the Homescan panel data, we find that Washington’s high-income consumers have different preferences over product types, relative to the “mainstream” preferences observed in other states. Figure 9 demonstrates that, most notably, Washington high-income consumers (top-third in the income distribution) are less enamored with bourbon (and rye), as well as scotch (and Irish whiskey), but far more enthusiastic about Canadian whiskey.<sup>22</sup> We reckon this pattern could be related to the proximity of Washington to Canada, and the distance from the key domestic producing regions (e.g. Kentucky) and their European counterparts. This difference changes the composition of customers who purchase each whiskey type: As a result, bourbon and scotch/Irish whiskey draw more low-income consumers, and Canadian whiskey attract more high-income consumers.

[Insert Figure 9 about here]

Second, we contrast this difference in customer preferences with retail prices in Washington. If retailers are Bayesian learners with initial beliefs reflecting information from other states, one might expect initial prices to *not* account for Washington’s distinct preference distribution. To validate this conjecture, we provide a simple calibration exercise that measures retailers’ *perceived* demand and examines how these demand perceptions differ from Washington’s actual demand. Firm beliefs are difficult to pin down without further assumptions on pricing behavior. As such, for this calibration exercise, we assume that retailers are static profit maximizers who set prices to maximize profit conditional on *perceived* demand (ignoring demand uncertainty). Specifically, we assume that the perceived demand follows the demand model in Section 4.1 up to differences in  $\gamma_k$ ’s – which govern the sorting of customer demographics into different product types. Although the assumption of naive and myopic pricing behavior is strong, it allows us to *flexibly* recover each retailer’s initial perceptions, without

---

<sup>22</sup>We use data during the sample period where Washington prices have stabilized. As such, the quantity differences are not driven by Washington’s suboptimal initial prices.

specifying how these initial beliefs are formed (e.g., without assuming rational expectation). In Web Appendix W.5, we document how we calibrate the initial (first quarter) pricing perceptions to capture how well the initially inaccurate beliefs explain the initial pricing mistakes.<sup>23</sup>

The first three columns of Table 4 summarize the difference between perceived and estimated  $\gamma_k$ 's, by retailer and whiskey-type. We find that retailers make sizable and systematic mistakes for bourbon and rye, and for Canadian whiskey. Our structural demand model estimate for  $\gamma_1$  (for bourbon and rye) is -0.30, implying that Washington high-income consumers do not favor bourbon. In contrast, retailer-perceived  $\gamma_1$ 's are between 1.36 and 1.65 for all retailers, suggesting that retailers set initial bourbon prices anticipating that a significantly higher fraction of high-income customers favor this whiskey-type. Conversely, the perceived  $\gamma_2$  parameters (corresponding to Canadian whiskey) are all negative compared to the structural demand estimates, suggesting that *all* retailers believe that Canadian whiskeys should be targeted to the lower-income segment.

[Insert Table 4 about here]

These systematic mistakes are consistent with retailer beliefs initially being substantially biased: away from Washington's true consumer preferences and toward the "typical" preferences of an average US state, although these biases are later corrected through market feedback. Therefore, the initial prior plays an important (and persistent) role in retailers' pricing strategies during the learning phase, consistent with the predictions from canonical Bayesian learning models (but distinct from adaptive learning, where prior beliefs do not play a role).

**Observation 4: prior experience helps initial pricing but lowers the initial learning rate.** Our final observation is that retailers' prior experiences in other markets can impart knowledge about conditions in the Washington market. We contrast initial pricing decisions by Retailer 32, which has experience selling liquor in other states, with those made by Retailers 158 and 182, which are entirely new to the retail liquor market. We demonstrate that prior experience in other states allows retailer 32 to set better initial prices, while slowing down its initial learning.

---

<sup>23</sup>We also compare with an alternative assumption on the learning goal: that retailers are initially incorrect about the overall price sensitivity. We show in Appendix W.5 that the learning about customer composition would rationalize about 53%-62% of the initial profit gap – a much better explanation of the data than the alternative.

Because Washington’s consumer preference distribution is distinct from other states’, it is possible that direct experiences from selling liquor in other states gives the retailer strong, but biased, priors that reflect consumer behavior in other states. It is also possible, however, that the experienced retailer accumulates more nuanced knowledge about consumer preferences by operating in different states, allowing it to better predict Washington’s distinct consumer preferences. A third possibility is that all retailers have access to some information about the retail liquor market (e.g., through hiring experienced managers or through market research), with or without having sold liquor in other states.

The calibration exercise above has demonstrated that all retailers’ perceived consumer preferences are biased in the same direction, yet the experienced Retailer 32 has (slightly) smaller biases across all whiskey types, compared to the two inexperienced retailers. This result suggests that all retailers leverage information from other states (despite this information being somewhat misleading about Washington tastes *per se*), but the experienced retailer exhibits a better initial understanding of the new market itself.

To explore this further, we estimate equation (16) separately for experienced versus inexperienced retailers to examine how the *initial* average absolute %price gap, as well as its *rate of change*, differ between the two groups. Table 3 shows that the experienced retailer is better-endowed with information about demand, and thus sets initial prices closer to the optimum (consistent with their more calibrated perceived demand). The average product is priced 10% away from the optimum, and the total profit is 6.3% below the optimum. In contrast, the average product is priced 15% away from the optimum for the inexperienced retailers, with corresponding total profit 13.9% below the optimum. Moreover, despite setting better initial prices, the experience retailer learns at a lower rate. Absolute %price gaps decreases by 2.8 percentage-points after the first year, less than half of the inexperienced retailers’ learning rate.

This evidence suggests that retailer prior beliefs have a persistent influence on their decision-making during the learning phase. Consistent with the prediction of a canonical Bayesian learning model, retailers’ prior beliefs about demand – shaped by previous experiences in other markets – plays a significant role in their initial pricing decisions and the rate at which they absorb subsequent information. Nevertheless, it is noteworthy that such prior experience does not come from operating

in a similar market per se, but rather from operating in markets with different primitives (yet still useful in predicting the Washington market).

**Discussion.** Without imposing a structural learning model, we presented four empirical observations that characterizes the learning process: (1) the learning rate declines over time; (2) the learning rate increases with the precision of new information; (3) the firms' main learning objective appears to be Washington's distinct customer preference distribution; and (4) the experienced retailer, who has sold liquor in other states, forms priors that are closer to the true demand and are more precise, compared to inexperienced retailers.

These observations describe a learning process that is broadly consistent with a canonical Bayesian learning model (Hitsch, 2006; Ching, 2010): the implied learning rate is governed by the strength of the firm's prior (which is affected by experience from various sources) and the precision of its sales signals. Notably, we document these patterns without assuming a Bayesian model structure. As such, these observations provide valuable field evidence that supports the relevance of the Bayesian learning framework to firm learning.

However, two aspects of our empirical evidence go beyond the canonical framework. First, firms are apparently learning complex features of the customer preference distribution – a high-dimensional object. Canonical models characterize firm learning about product quality (Hitsch, 2006) or category-level quality (Ching, 2010), unidimensional objects that greatly facilitate model tractability. Our evidence suggests that the focus of firm learning can be more complex than what has been characterized in existing models. If one were to fully characterize this learning problem by a forward-looking firm, the model would involve the firm strategically setting prices to create data that facilitate future learning. The closest work to this model in marketing is Misra, Schwartz and Abernethy (2019), who present a scalable algorithm that approximates this learning problem. The full solution to the (structural) learning model is a promising area for future research.

Second, the role of firms' previous experience in other markets suggests that the transfer of firm knowledge is important for understanding learning, namely how priors are formed. In particular, firms demonstrate a capacity to predict outcomes in a new (ex ante counterfactual) environment, consistent

with the types of dynamic capabilities emphasized by Teece, Pisano and Shuen (1997). Precisely how prior beliefs are formed and how knowledge is transferred across markets is outside the scope of current learning models (presumably due to the added model complexity). This is another open area for future research.

## **6 Conclusion**

We have shown that, post-privatization, firms in the Washington State liquor market learned to set prices very much in line with classic models of imperfect competition. However, this process took some time to unfold and played out differently for different firms, suggesting important roles for both learning and heterogeneous firm capabilities. In our setting, we find that this convergence of pricing strategies to optimal decision-making does not require the pressure of direct competition. Instead, firm practices alone appear to create sufficient impetus to develop the sophisticated understanding required to reach optimal pricing policies. Future research could examine the sources and limitations of this drive, and whether competition, by adding the complexity of interpreting and anticipating competitive actions, helps or hinders the learning process.

## References

- Aghion, Philippe, Patrick Bolton, Christopher Harris and Bruno Jullien (1991), 'Optimal learning by experimentation', *The review of economic studies* **58**(4), 621–654.
- Aguirregabiria, Victor, Daniel Ershov and Junichi Suzuki (2016), 'Estimating the effects of deregulation in the Ontario wine retail market'.
- Aguirregabiria, Victor and Jihye Jeon (2020), 'Firms' beliefs and learning: Models, identification, and empirical evidence', *Review of Industrial Organization* **56**(2), 203–235.
- Arcidiacono, Peter, Paul B Ellickson, Carl F Mela and John D Singleton (2020), 'The competitive effects of entry: Evidence from Supercenter expansion', *American Economic Journal: Applied Economics* **12**(3).
- Arellano, Manuel and Stephen Bond (1991), 'Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations', *Review of Economic Studies* **58**(2), 277–297.
- Berry, Steven, James Levinsohn and Ariel Pakes (1995), 'Automobile prices in market equilibrium', *Econometrica* **63**(4), 841–890.
- Berry, Steven T (1994), 'Estimating discrete-choice models of product differentiation', *RAND Journal of Economics* **25**(2), 242–262.
- Ching, Andrew T (2010), 'A dynamic oligopoly structural model for the prescription drug market after patent expiration', *International Economic Review* **51**(4), 1175–1207.
- Conlon, Christopher T and Nirupama Rao (2019), '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'.
- DellaVigna, Stefano and Matthew Gentzkow (2019), 'Uniform pricing in US retail chains', *Quarterly Journal of Economics* **134**(4), 2011–2084.
- 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.
- Dubé, Jean-Pierre, Günter J Hitsch and Peter E Rossi (2010), 'State dependence and alternative explanations for consumer inertia', *RAND Journal of Economics* **41**(3), 417–445.
- Erdem, Tülin, Susumu Imai and Michael P Keane (2003), 'Brand and quantity choice dynamics under price uncertainty', *Quantitative Marketing and Economics* **1**(1), 5–64.
- Goldfarb, Avi and Mo Xiao (2011), 'Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets', *American Economic Review* **101**(7), 3130–3161.
- Hendel, Igal and Aviv Nevo (2006), 'Measuring the implications of sales and consumer inventory behavior', *Econometrica* **74**(6), 1637–1673.

- 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 Hortaçsu and Xiliang Lin (2021), ‘Prices and promotions in us retail markets’, *Quantitative Marketing and Economics* pp. 1–80.
- 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.
- Huang, Guofang, Hong Luo and Jing Xia (2019), ‘Invest in information or wing it? a model of dynamic pricing with seller learning’, *Management Science* **65**(12), 5556–5583.
- Illanes, Gastón and Sarah Moshary (2018), ‘Market structure and product variety: Evidence from a natural experiment in liquor licensure’.
- Jeon, Jihye (2020), ‘Learning and investment under demand uncertainty in container shipping’.
- Kaplan, Abraham, Joel B Dirlan and Robert F Lanzillotti (1958), *Pricing in big business: A case approach*.
- Leisten, Matthew (2021), ‘Information, managerial incentives, and scale: Evidence from hotel pricing’.
- Li, Xinlong and Andrew T Ching (2021), ‘How does a firm adapt in a changing world? the case of prosper marketplace’, Available at SSRN: <https://ssrn.com/abstract=3403404> or <http://dx.doi.org/10.2139/ssrn.3403404>.
- Liu, Xinyuan, Zaiyan Wei and Mo Xiao (2020), ‘Platform mispricing and lender learning in peer-to-peer lending’, *Review of Industrial Organization* **56**(2), 281–314.
- Miravete, Eugenio J, Katja Seim and Jeff Thurk (2018), ‘Market power and the Laffer curve’, *Econometrica* **86**(5), 1651–1687.
- Miravete, Eugenio J, Katja Seim and Jeff Thurk (2020), ‘One markup to rule them all: Taxation by liquor pricing regulation’, *American Economic Journal: Microeconomics* **12**(1), 1–41.
- Misra, Kanishka, Eric M Schwartz and Jacob Abernethy (2019), ‘Dynamic online pricing with incomplete information using multiarmed bandit experiments’, *Marketing Science* **38**(2), 226–252.
- 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.
- 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', *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.
- Trajtenberg, Manuel (1989), 'The welfare analysis of product innovations, with an application to computed tomography scanners', *Journal of Political Economy* **97**(2), 444–479.

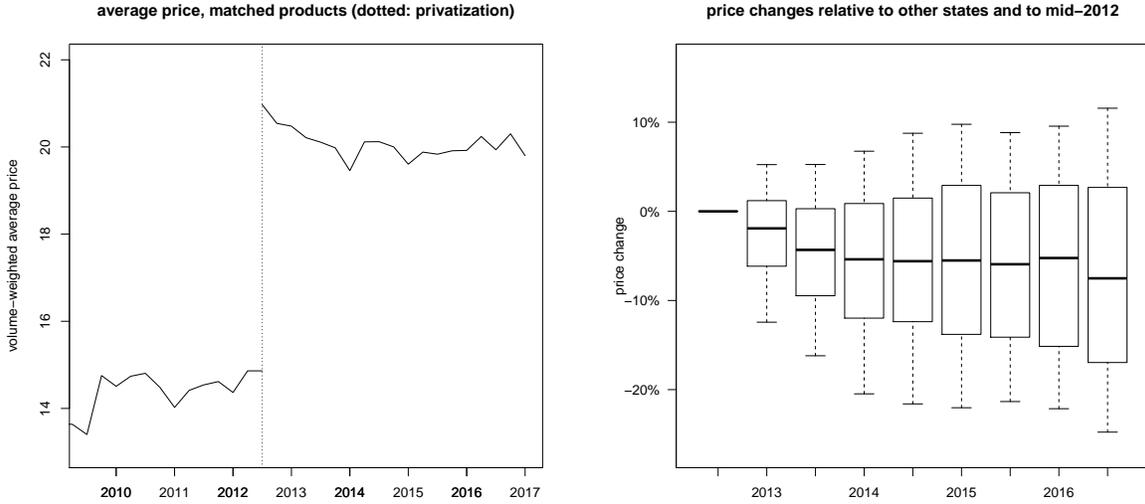


Figure 1: Price changes over time and between Washington and other states

**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 as volume weight, and fix the weight over time. The right-panel shows the distribution of price changes relative to mid-2012 and to other states. The y-axis represents the relative differences between price and initial price for a given product. The boxes are 25th and 75th percentile while the whiskers are 10th and 90th percentile.

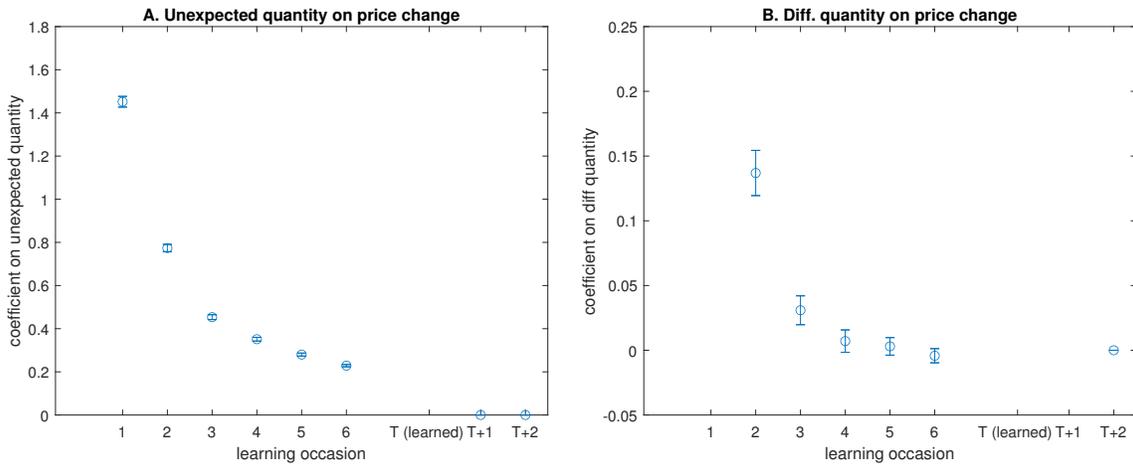


Figure 2: Price adjustments and lagged unexpected quantity: simulation under Bayesian learning

**Notes:** The left panel presents the regression coefficients of price changes,  $p_{jt} - p_{jt-1}$  on unexpected quantity, defined as  $q_{jt-1} - \mathbb{E}[q_{jt-1} | p_{jt-1}, \mathcal{F}_{jt-1}]$ , where  $\mathcal{F}_{jt-1}$  is the retailer's prior belief at  $t - 1$  about the demand for product  $j$ . Each of the learning occasions 1-6 refers to retailer setting prices given posterior belief after a set number of observed quantity signals. "T (learned)," T+1, and T+2 refers to when the retailer has perfect information about demand (and still Bayesian updates, as a form of placebo check). Finally, the right panel presents the regression coefficients of price changes on lagged quantity changes,  $q_{jt-1} - q_{jt-2}$ .

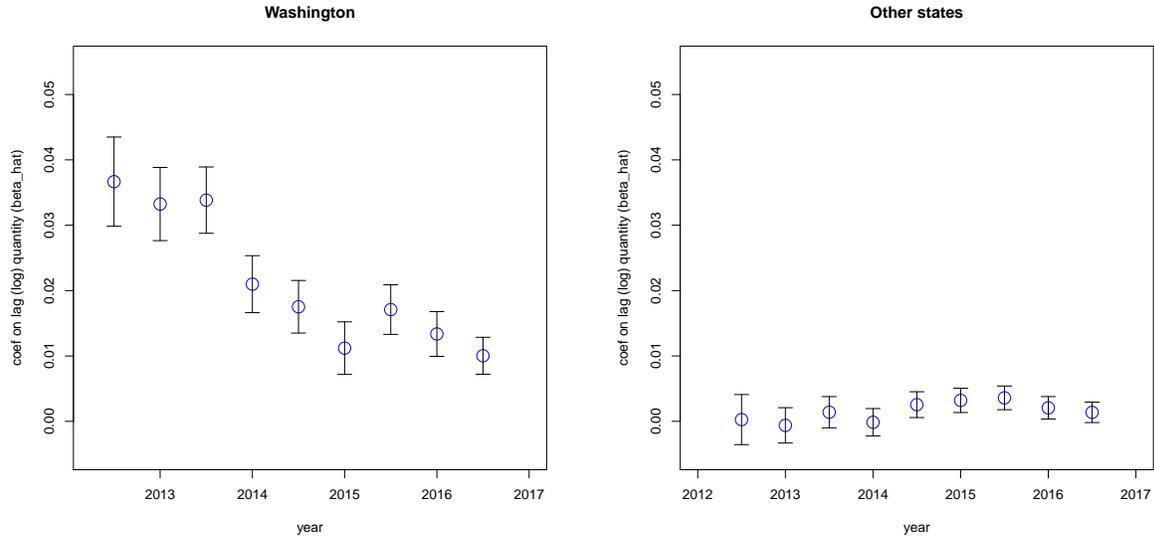


Figure 3: The response of current price to lagged quantity, Washington (left) and other states (right)

Notes: Estimates  $\beta_\tau$  from equation (2). See Web Appendix Table 7 for more detail. Confidence intervals are two standard error.

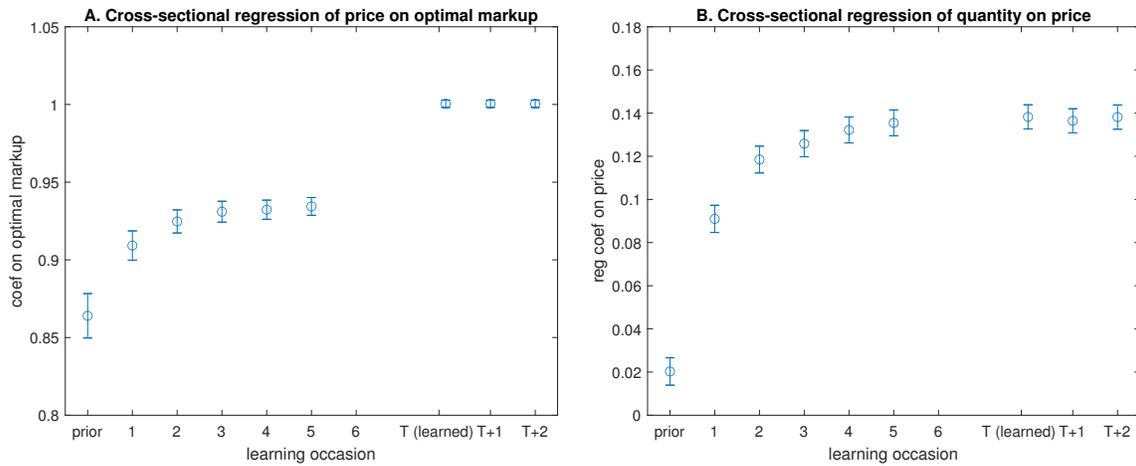


Figure 4: Price and demand fundamentals: simulation under Bayesian learning

Notes: The left panel presents cross-sectional regression coefficients of price on optimal markup. “Prior” refers to the period when the retailer draws a normal prior about demand primitives. Each of the learning occasions 1 to 6 and T to T+2 is defined in the same way as the above. Finally, the right panel presents cross-sectional regression coefficients of quantity on price.

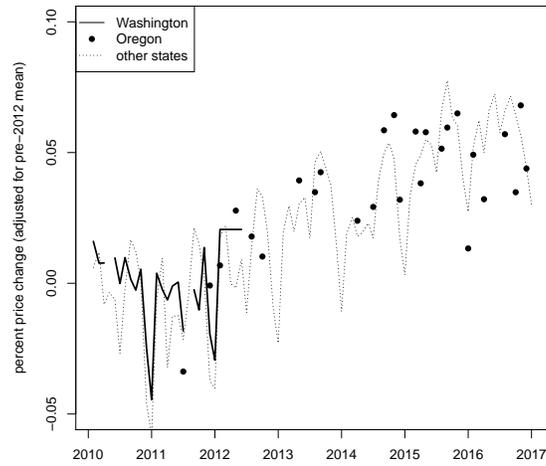


Figure 6: Percent retail price change: comparison across states

**Notes:** This figure shows the percent average retail price change relative to the period of 2010-June 2012. The bold, solid line shows Washington state-chain's price, the points are Oregon state chain's, and the dotted line is the average between other privatized states.

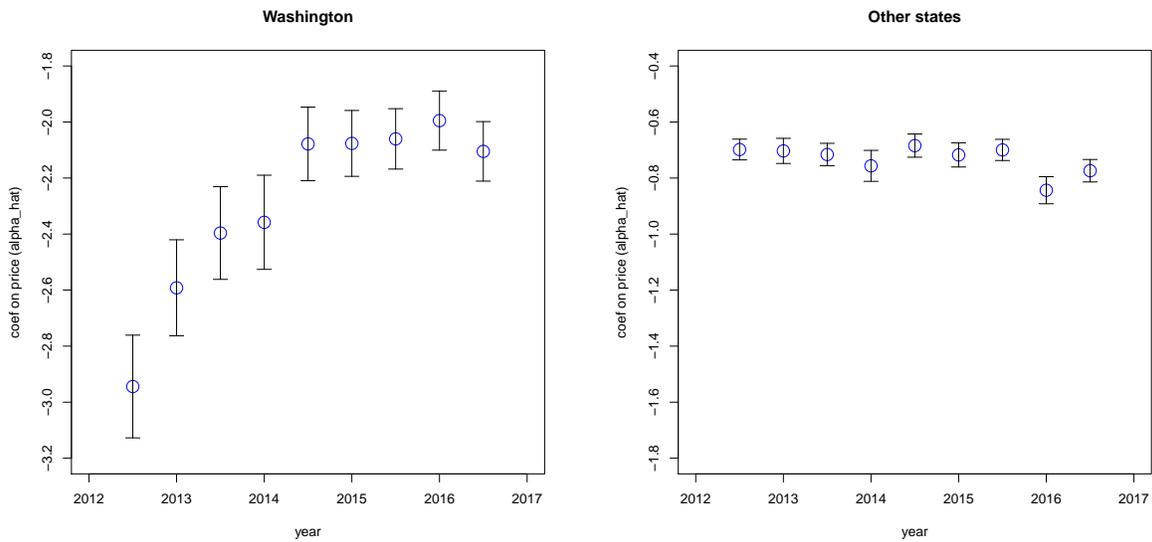


Figure 5: Coefficient estimates of  $\alpha_\tau$  from equation (3), Washington (left) and other states (right)

**Notes:** We report coefficient estimates of  $\alpha_\tau$  from equation (3) where  $\tau$  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.

Table 1: Demand parameter estimates

	Main spec.	Berry (1994)	price (first stage)
price ( $\alpha_0$ )	-0.833 (0.002)	-0.171 (0.015)	
price $\times$ log (income) ( $\alpha_1$ )	0.121 (0.002)		
std. dev. of price coef. ( $\sigma_v$ )	0.053 (0.014)		
bourbon/rye $\times$ log (income) ( $\gamma_1$ )	-0.301 (0.116)		
Canadian $\times$ log (income) ( $\gamma_2$ )	0.773 (0.100)		
Irish/scotch $\times$ log (income) ( $\gamma_3$ )	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 ( $\delta_{jrm}$ )	X	X	X
retailer-year and month FE ( $\lambda_{rt}$ )	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.

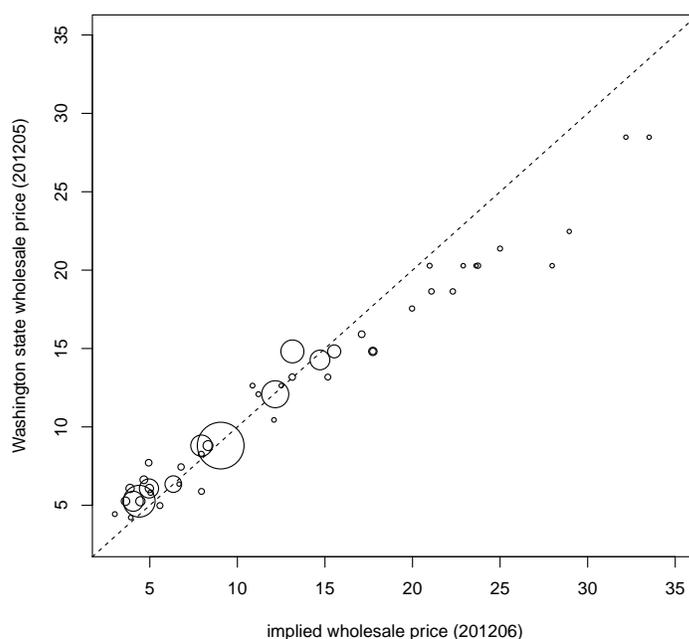


Figure 7: Comparison between estimated and observed wholesale prices

**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.

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: Initial mistakes and learning rate between time periods, retailers, and product groups

	Average volume		Retailer experience	
	High	Low	Inexpr.	Expr.
Initial absolute %price gap	0.122 (0.001)	0.131 (0.002)	0.145 (0.001)	0.099 (0.001)
Change in price gap per year	-0.064 (0.002)	-0.035 (0.003)	-0.067 (0.003)	-0.028 (0.003)
Observations	12,199	9,774	13,181	8,792

**Notes:** This table reports the initial absolute percentage price gap (which is the absolute value of the %price gap defined in Section 5.2) and how this price gap decreases per year in the first year of the privatization, across products with high and low sales volume and across retailers with different prior experience. “Inexperienced” retailers include retailer 158 and 182, the two local retailers who started selling liquor since the privatization in Washington. “Experienced” retailer refers to retailer 32, which has sold liquor for a long time in other states. Standard errors in parenthesis.

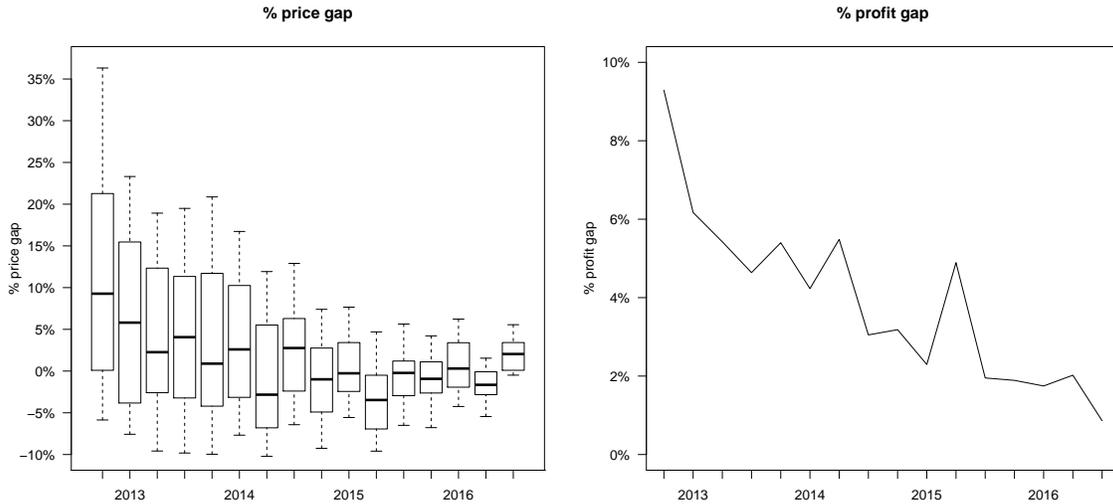


Figure 8: Percentage price and profit gap: pooled across retailers

**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.

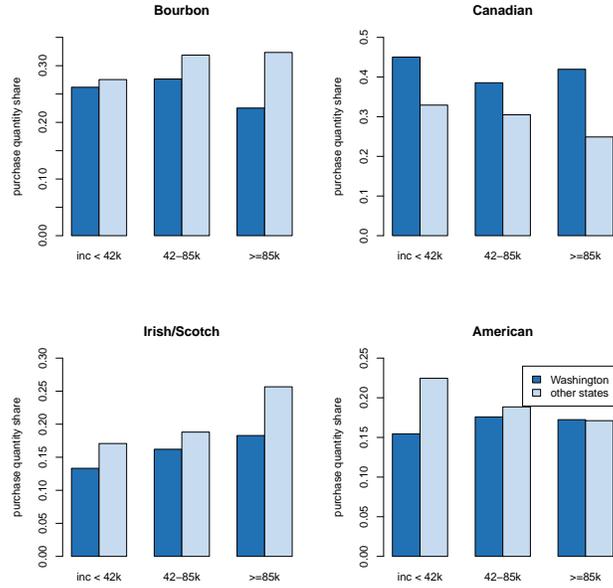


Figure 9: Share of purchases by product type and income group

**Notes:** Probability of purchasing each product type, 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.

Table 4: Retailer-perceived product type preferences and observed consumer type choice

	perceived $\tilde{\gamma}_k$ – estimated $\gamma_k$			$\Delta$ 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 product-type shares between other states and Washington. Column 1-3 present the gap between perceived and estimated type-income interaction term, by type and retailer. For example, the first column reads: Retailer 32's initial prices can be rationalized as optimal prices in a market where  $\gamma_1$  (for bourbon and rye) is 1.063 higher than the true demand,  $\gamma_2$  (for Canadian whiskey) is 1.689 lower, and  $\gamma_3$  (for Irish whiskey and scotch) is 0.297 lower. Column 4-5 are differences in type shares conditional on purchase, calculated from the Homescan data (Figure 9). 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.