Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board's Entry Decisions*

Katja Seim
The Wharton School
University of Pennsylvania
kseim@wharton.upenn.edu

Joel Waldfogel
The Carlson School of Management, Department of Economics, and Law School
University of Minnesota and NBER
jwaldfog@umn.edu

May 14, 2012

Abstract

We estimate a spatial model of liquor demand to analyze the impact of government controlled retailing on entry patterns. In the absence of the Pennsylvania Liquor Control Board, the state would have roughly 2.5 times the current number of stores, higher consumer surplus, and lower payments to liquor store employees. With just over half the number of stores that would maximize welfare, the government system is instead best rationalized as profit maximization with profit sharing. Government operation mitigates, but does not eliminate, free entry's bias against rural consumers. We find only limited evidence of political influence on entry.

* We thank seminar participants at Kellogg, MIT, Ohio State, Texas, Toronto, Virginia, Wisconsin, the Tenth CEPR Conference on Applied Industrial Organization, and the 2010 NBER Winter IO Meetings at Stanford. We are grateful to Tim Bresnahan for comments on an earlier draft. We thank Thomas Krantz at the Pennsylvania Liquor Control Board and Mike Ehtesham at the National Alcohol Beverage Control State Association, Inc., for helping us to get access to the data. John Coglianese provided valuable research assistance. All errors are our own.

1. Introduction

An economic system can leave entry decisions to markets or to government. Markets have many well-known advantages, along with some well-understood challenges. For example, private action can result in insufficient entry when benefits cover costs but revenue does not; and private entry can lead to excessive entry when revenue covers the cost of an additional outlet even though the incremental social benefit does not. Moreover, even if the number of outlets is fixed at the correct level, private entry can result in the choice of sub-optimal locations, as in the well-known Hotelling 2-firm result (Hotelling 1929). A planner can, in principle, avoid these problems if he internalizes business stealing effects while also attaching a benefit to consumer surplus. But even such a planner's entry decisions may face challenges. Government-controlled entities can be captured politically and may allocate resources to serve political ends rather than to promote economic efficiency. For example, labor costs may be higher if union labor is favored; and store location decisions might be subject to political pressure.

It is difficult to evaluate the efficiency and apparent motives of centralized entry decisions because few contexts allow for direct comparison of government and market entry patterns. One exception is liquor retailing in the United States. Since Prohibition, liquor distribution has been heavily regulated by state and local governments, each of which has chosen its own regulatory path. The 50 US states are divided broadly into two allocative camps: 32 "private" or "open" states where the number of stores operating is generally regulated but operators are free to choose particular locations and 18 "control" states, where the government has a monopoly on liquor retailing, wholesaling, or both. In Pennsylvania, all stores are both controlled and operated by the state's Liquor Control Board (PLCB), with unionized government employees.

This paper studies entry decisions made by the PLCB with the goal of addressing two questions, one positive and one normative. First, how does allocation by a government monopoly affect outcomes? That is, how do store configurations and resulting welfare under the PLCB compare with plausible private alternatives? Second, what implicit motives underlie the government-operated system? This second question has three parts: (i) how closely does its operation conform to the theoretical benchmarks of free entry, or profit or welfare maximization? (ii) What do PLCB entry patterns reveal about the government's attitude toward different types of consumers? and (iii) Is there evidence of political influence?[1]

---

[1] Our work has similarities with recent studies of store entry decisions by big-box retail chains (see, e.g. Jia 2008 and Holmes 2011). In contrast to these settings, where static or dynamic profit maximization appears a natural objective for the firms, this

We explore these questions in seven sections. Section 2 presents background on liquor retailing and, in particular, a comparison of the systems in private and control states. Two facts about the PLCB are clear from this comparison: relative to private states, Pennsylvania has higher store operation costs and operates far fewer stores per capita. Section 3 describes the data used for estimation. Section 4 presents a model of spatial demand that we use to calculate the quantities sold at each store location, as well as consumer and producer surplus in each location, for any configuration of stores. Section 4 also describes how we use the model to calculate various counterfactual store configurations, including free entry as well as efficient configurations that maximize welfare or profit.

We then use the modeling to answer two sets of substantive questions. Section 5 presents a comparison of the current system with free entry simulations to describe the private system Pennsylvania would have absent the PLCB. We find that the welfare impact of the PLCB is to reduce consumer surplus and raise producer surplus, much of which is shared with labor under the current system. Section 6 provides evidence on motives underlying the PLCB's store configuration. We use our model to characterize a continuum of "efficient" store configurations that maximize a weighted sum of profit and consumer surplus ($\pi + \gamma CS$). Viewed against theoretical benchmarks of profit maximization ($\gamma = 0$) and welfare maximization ($\gamma = 1$), we find that the current system has just over half the number of stores that would maximize welfare if the state faced competitive labor costs. Instead, the PLCB system resembles profit maximization with labor profit sharing, or welfare maximization, given a constraint of paying super-competitive labor costs. The PLCB system mitigates a bias against rural consumers that would prevail under free entry. We see only limited evidence of political influence on store location decisions. Section 7 concludes with a discussion of the likely effects of the PLCB on aggregate welfare.

2. Liquor Distribution in Pennsylvania versus Other States

Pennsylvania is at an extreme among control states, acting as a state monopolist in the wholesale and retail distribution of wine and liquor through a system of state-run stores staffed by unionized government employees. Pennsylvania has a private system for the sale of beer, which is sold by the case in licensed private "beer distributors" and by the six-pack at bars and restaurants. By contrast, some of the control states, like Ohio and Maine, contract with private firms to operate retail stores on the state's behalf. In others, such as Utah and Washington, the

is less apparent in the context of a public enterprise. See, for example, Boardman and Vining (1989) for a prominent study comparing the efficiency of private and public enterprises.

state operates some stores, while private licensees operate others. Private states, on the other hand, employ regulated private entry, allowing fully private retailing operations but limiting the supply of licenses, generally within each municipality. This section compares Pennsylvania to other control and private states along several dimensions, including number of stores and their workforce, liquor taxes, pricing and selection at stores, and consumption per capita.

a. *Entry in control and private states*

Since private states do not have unregulated free entry, but typically award a limited number of liquor store licenses, it is not clear *a priori* that a private system in Pennsylvania would have more or fewer stores than the government system. In principle, a comparison of the number of liquor stores in control and private states for any given level of demand is easy. The 2007 Economic Census provides data on the number of stores selling beer, wine, and liquor in each state, and the 2000 Census provides data on population, which provides a reasonable proxy for demand. There are a few complications, however. First, many states allow the sale of alcoholic beverages in grocery stores; and such states will have fewer standalone liquor stores per capita. Using the 2007 Economic Census data on sales by line of business, we can calculate the share of packaged alcoholic beverage sales occurring in dedicated liquor stores (α). If N is the number of liquor stores, N/α is an approximation of the number of liquor stores if all packaged liquor demand were satisfied by dedicated liquor stores. Second, unlike liquor stores elsewhere, PLCB stores sell only wine and spirits (and not beer), depressing the number of PLCB outlets relative to population. The Economic Census product line data indicate that 35 percent of packaged liquor sales in Pennsylvania are beer, so we adjust the number of PLCB stores by scaling by (1/0.65). Figure 1 plots the resulting adjusted number of liquor stores against population in log terms.

As of the first week of 2005, Pennsylvania operated 621 wine and spirits stores, each serving an average of 14,562 residents over the age of 21. In contrast, stores in private states serve an average of 7,944 residents, while stores in other control states serve 11,184. Even with both of our adjustments, Pennsylvania – and a good number of the other control states – thus has fewer stores per capita than do private states. The relative paucity of Pennsylvania stores may depress drinking: in 2005 wine and spirit consumption averaged 3.61 gallons per capita in Pennsylvania, significantly short of the 5.15 gallons consumed in the average private state. In contrast, other control states that are active in alcohol retailing typically focus on spirits products only; their per capita consumption of spirits averaged 1.94 gallons in 2005, while Pennsylvania's was 1.53 gallons and private states averaged 2.19 gallons.[2]

How much more entry would we expect to see absent the PLCB? Based on the fitted relationship between log adjusted entry and log population for the private states, the PLCB operates 59 percent fewer stores than one might otherwise expect in Pennsylvania (2,355 stores). Hence, we can roughly estimate that a private system for selling only wine and spirits in Pennsylvania would have (0.65)×(2,355), or 1,531 stores, roughly 2.5 times the current number.

b. Pricing and Selection

The PLCB charges an identical retail price for a particular product in all of its stores using a simple markup rule to determine the price. The pricing rule is set in the Pennsylvania Liquor Code by the State Legislature. Accordingly, the PLCB applies a 30% markup and an 18% liquor tax to the wholesale price.[3] In effect, Pennsylvania's liquor tax is 2.3 times higher than the average for other states: for the average bottle in our data, Pennsylvania's liquor tax is \$1.89 per bottle, compared with \$0.81 in other states.[4]

As we document in related work (Miravete, Seim, Thurk, and Waldfogel (2012)), we have no conclusive evidence that retail prices vary systematically between Pennsylvania, other control, and private states: Pennsylvania's prices are in line with, and frequently below, those in other control states. Small-sample comparisons of prices for specific products in Pennsylvania and retail stores in neighboring private states similarly do not suggest significant differences.

Another possible difference between Pennsylvania's liquor retailing system and what might prevail in a private system is the product selection carried by each liquor store. According to the store level data that we use in this paper (described in detail in Section 3), the mean (median) PLCB store sold a total of 1,371 (1,254) different wine and spirits products, with a standard deviation of 709. While we lack similarly detailed product availability data for stores in

---

[2] Calculated as the state's total apparent consumption by type of beverage divided by its population over the age of 21. Source: LaVallee, R.A. and Yi, H., 2011. Surveillance Report #92: Apparent Per Capita Alcohol Consumption: National, State, and Regional Trends, 1977-2009. Bethesda, MD: NIAAA, Alcohol Epidemiologic Data System.

[3] The specific pricing rule is: $retail\ price = (wholesale\ price(1.3) + bottle\ fee)(1.18)$, where the bottle fee amounts typically to \$1 and the PLCB rounds the resulting retail price to end in the nearest nine cents. In addition, the consumer pays a 6% Pennsylvania sales tax.

[4] See the American Wine Institute (2011) http://www.wineinstitute.org/resources/statistics/article86, the Distilled Spirits Council (2011) http://www.discus.org/pdf/Spirits_Category_Tables_2010.pdf, and Commonwealth of Pennsylvania (2011) http://www.portal.state.pa.us/portal/server.pt/community/liquor_privatization_analysis__final_report/4575, accessed January 13, 2011. We convert wine and spirits-based gallonage taxes from other states into a single, value-based, Pennsylvania tax rate by calculating a weighted average gallonage rate using the break-down of sales into wines and spirits and expressing the resulting tax as a percentage of the mean marked-up Pennsylvania wholesale price.

other states, we can compare the square footage of dedicated liquor stores in a random sample of zip codes in states bordering Pennsylvania (NJ, NY, OH, WV)[5] to the size of the PLCB stores. For the PLCB, the correlation between a store's average product selection and its square footage is 0.56, suggesting that store size proxies reasonably well for variety. Sample stores in the adjacent states are significantly smaller than Pennsylvania stores, with a median store size of 55 to 64% of the median PLCB store's size. Pennsylvania thus operates fewer, but larger, stores than alternative systems. These statistics suggest that the typical PLCB store does not carry fewer products than do stores in other states.

c. Labor Costs

The PLCB employs unionized store clerks and pays them according to a single, state-wide pay-scale. The 2007 Economic Census reports that the average pay per Pennsylvania employee in beer, wine, and liquor retailing was $26,000, or $43,680 including benefits.[6] The PLCB employed 4,896 workers in 2009, and total operating expenses ("Store, Warehouse, and Transportation Costs") were $299.7 million that year.[7] Hence, labor costs were 5/7th of total operating expenses.[8]

How do these labor costs compare with those in private states? According to the 2007 Economic Census, pay at stores selling beer, wine, and liquor (NAICS 4453) averaged $16,000 per worker, or $21,000 with benefits, in private states, less than half the rate at the PLCB.[9, 10] In addition to paying more per worker, PLCB stores employ more workers per store. PLCB stores have an average of 7.9 workers per store, while, according to the 2007 Economic Census, liquor stores outside Pennsylvania had an average of 4.6 employees per store.

---

5 The data were obtained from Dunn & Bradstreet and contain information on 64 stores in New Jersey; 136 stores in New York; 49 stores in Ohio; and 84 stores in West Virginia.

6 The Bureau of Labor Statistics reports that state and local government service employees received $0.68 in benefits per dollar of pay. See http://www.bls.gov/news.release/ecec.t04.htm. We derive total labor costs inclusive of benefits by scaling wage payments by 1.68.

7 See the PLCB Fiscal Year 09-10 Summary.

8 Luciew reports in a 2009 article in the Patriot News that the PLCB paid $224 million in total labor costs in 2007, when the agency had 4,439 employees, implying total labor cost of $50,000 per employee. For the sake of conservatism, we adopt the estimate in the text.

9 That is, excluding Pennsylvania as well as other states with at least some direct government involvement in retailing: Alabama, Idaho, New Hampshire, North Carolina, South Carolina, Utah, and Virginia, as well as Maryland, Minnesota, and Washington.

10 The Bureau of Labor Statistics reports that employees in retail trade earned $0.33 in benefits per dollar of pay. See http://www.bls.gov/news.release/ecec.t10.htm.

The PLCB currently spends $1,110 per day to operate a store. How much would it cost to run a store absent the current system? We do not have information on rental expenses and distribution costs (each of which account for 1/7th of PLCB operating expenses) in private states. Holding these fixed at current levels, we obtain one answer from assuming that PA stores would have their current levels of employment but half the current rate of pay. We refer to this as the "competitive wage" alternative, and it results in $713 per store per day. We obtain a second estimate by assuming the Pennsylvania stores would otherwise have both typical rates of pay and the more common levels of employment per store. We term this the "competitive cost" scenario, and it results in $549 per store per day. Thus, current store operation costs appear to be twice those in states with private systems.

Before moving on, two descriptive facts uncovered above bear emphasis: relative to private states, 1) the PLCB faces high store operation costs, and 2) the PLCB operates far fewer stores per capita than would likely prevail in a private system. Our goal in the remainder of the paper is to use these facts, along with a model of demand and a method for describing entry, to evaluate the welfare consequences of PLCB operation, along with its implicit motives.

3. Data

The basic data set for the study is a store-level panel obtained from the PLCB under the Pennsylvania Right-to-Know Law.[11] It contains daily information on quantities sold and gross receipts at the product and store level during 2005. In addition, we received information on the wholesale cost of each product that is constant across stores and varies over time according to reporting periods described below. We geocode the stores' street addresses to assign them to a geographic location, which we link to data on population and demographic characteristics for nearby consumers based on information from the 2000 Census and Reference USA. Because stores open and close during the year, the characteristics of stores' ambient consumers change over time.

We aggregate our data across products to the level of either the day or the week. This periodicity accounts for the strong seasonality inherent in liquor sales, which are disguised in more aggregate definitions. Averaging across 32,509 store weeks in 2005, stores sell an average of 2,674 bottles per week. Figure 2 exhibits the strong seasonal pattern to sales, with a trough after New Year's (week 1) and peaks at July Fourth (week 26), Thanksgiving (week 47), and Christmas through New Year's Eve (weeks 50-52).

---

[11] 65 P.S. §§ 66.1 *et seq.*, as amended.

Because we treat liquor as a single quantity in our analysis below, we also need a single price. PLCB stores carry thousands of products, and we calculate a state-wide price index that is a weighted average of the system-wide product prices in each week. We use fixed weights – the products' shares of 2005 sales – to avoid contaminating the price index with quantity responses.

As discussed in Section 2, the PLCB uses a markup formula to calculate prices. The PLCB is further able to pass on temporary wholesale price reductions to the consumer in the form of system-wide monthly sale periods (28-day period beginning on the Monday closest to the end of the month). As a result, we observe changes in prices for two reasons: (1) an adjustment in wholesale prices, or (2) temporary sale prices on a subset of products. The PLCB negotiates wholesale prices directly with its suppliers. A new product's wholesale price remains fixed for one year after introduction. For established products, the PLCB re-negotiates over cost increases on a quarterly basis rotating through product categories over the course of its four-week-long reporting periods. Each reporting period, the wholesale price of a subset of products is adjusted, translating into changes in the retail price. In contrast to sales periods, which typically begin on the last Monday of a month, reporting periods begin on a Thursday, usually in the middle of the month. Prices can therefore change at two discrete times per month, and our price series resembles a step function.[12]

While stores differ in the mix of products sold, these differences reflect heterogeneity in consumer preferences more than differences in availability. Of the 100 best-selling products statewide in 2005, the median store carried 98.0% in its median week, while a store at the fifth percentile carried 72.0% of the products. Similarly, of the 1000 best selling products statewide in 2005, the median store carried 82.03% in its median week, while a store at the fifth percentile carried 44.2% of the products. The PLCB operates 65 larger stores that are designated "premium-collection" stores.[13] The product availability at premium stores is somewhat better than the average, with the median premium store carrying all of the top 100 products and 95.1% of the top 1000 products. But most stores carry most products, supporting our assumption below that differences in product availability do not drive customers' store choices to a significant

[12] In our data, 90.26% of price changes occur within one week from the beginning of a new reporting or sales period, reflecting that not all products have daily sales in at least one PLCB store.

[13] The PLCB also operates seven "outlet" stores near the borders with neighboring states. In addition to the typical selection, the PLCB sells certain products – typically larger-sized bottles or multi-packs – at these stores that are unavailable in the remaining stores.

degree. In our empirical exercises below we employ a single statewide price index reflecting our model's implicit assumption of a single identical product available at each store.[14]

a. Descriptive Evidence

Our model of demand links purchase behavior to demographic characteristics, the configuration of stores, and price. In this section we explore these relationships as a step toward more formal estimation. We first examine the relationship between prices and demand, via regressions of log quantities on measures of log prices.

It is possible that prices move endogenously with anticipated changes in demand. We address this potential endogeneity of the price series in a number of ways. First, we control for unusual spikes or declines in demand around holidays by including time dummies for holiday weeks or days, or a more flexible quadratic seasonality specification; since prices vary only across time and not place, we cannot include fully flexible time dummies. These time terms address endogeneity concerns to the extent that they control for the relevant temporary changes in demand that manufacturers anticipate when choosing their wholesale price discounts. The price elasticity is identified from the co-variation in quantity and the price index after accounting for common contemporaneous changes in sales experienced at the same time. Second, we employ a price index that removes price declines due to the potentially endogenous discounts. We call this the list price and build a statewide, fixed weight, price index based on it.

Across specifications that differ in seasonality controls, periodicity of the data, the inclusion of store fixed effect, and the selection of the sample, we find that demand is inelastic with a price elasticity ranging from –0.7 to –1.9 (see Table 1). This is in line with estimates from the large empirical literature estimating the elasticity of demand for liquor. Cook and Moore (1999) review the literature on demand for alcohol, most of which use state-level time series data. According to Chaloupka, Grossman, and Safer (2002), "An extensive review of the economic literature on alcohol demand concluded that based on studies using aggregate data (i.e., data that report the amount of alcohol consumed by large groups of people), the price elasticities of demand for beer, wine, and distilled spirits are –0.3, –1.0, and –1.5, respectively (Leung and Phelps 1993)."

The second relationship of interest is between ambient population and quantity demanded. Table 2 explores this relationship systematically with multiple regression using the

---

[14] We performed various descriptive exercises (like those in Table 1 below) with store-specific price indices, and their use results in demand elasticities similar to those implied by the statewide index.

population-weighted average great-circle distance[15] to the store as a proxy for travel cost. In aggregate, assuming that all population resides at Census tract centroids, the average (median) great-circle distance to the nearest store is 3.2 (2.4) kilometers, with an interquartile range of 1.0 to 3.6 km. The descriptive results suggest that population increases demand while demand declines with distance to the nearest store; the estimated price elasticity is robust to the inclusion of demographics at –0.9.

Table 2, as well as our estimates below, employs daily price and quantity data. With this level of aggregation, there is variation in catchment areas over time since different stores are open on different days and at different times of the year. We observe several permanent changes to the store configuration during the year: twelve new stores opened in 2005, while six existing stores closed. Three other stores relocated. There is also regular variation in catchment areas over the course of the week. While most stores are open six days per week, 10% of PLCB stores are open on Sundays as of the beginning of 2005. Following authorization by the state legislature to increase this set of stores, we observe an additional 90 stores recording Sunday sales by the end of 2005. The PLCB phased in the conversion of these stores to seven days a week gradually over the course of the year. Twelve stores have limited hours and are consistently closed on one or two of the six regular business days. There are also temporary closings, which we identify in the data as regular sales days where no sales were recorded for a given store. Two stores were closed for an extended period of several weeks, while 61 stores recorded no sales for a subset of their regular sales days for at least one, and frequently for several, weekdays. These openings and closings help identify the effect of distance to the store on demand beyond purely cross-sectional variation.[16]

We also explore descriptively how sensitive the results in Table 2 are to some of the salient features of the Pennsylvania liquor market. First, we re-estimate specifications (3) and (4) excluding holiday weeks (Thanksgiving week and last three weeks of December) from the sample, to test whether the base results are driven by differences in willingness to pay for liquor or travel to the store in these unusual weeks. We obtain very similar results with this limited

[15] Great-circle distances are calculated according to the Haversine formula and measure the shortest distance along the surface of a sphere between any two locations.

[16] Note that despite the panel nature of our data, store fixed effects do not address a possible concern about unobserved spatial heterogeneity. We would ideally like to control for unobserved preference shifters of consumers that may be correlated with the distance such consumers travel to the store. However, we do not observe the demand associated with particular consumers. Instead, we observe store-level demand. Because the group of consumers nearest each store varies across days, a store fixed effect does not control for the same consumers' unobserved demand. While we report a fixed-effect estimate of the distance coefficient in Table 2 nevertheless, we address a concern over spatial heterogeneity in demand by investigating the robustness of the estimates of our full demand model to the inclusion of a host of potential observable demand shifters below.

sample. Second, we explore whether systematic differences in demand in areas close to Pennsylvania's borders, including in Philadelphia and Pittsburgh, drive the relationships in Table 2. Demand in these areas may be more elastic than in the interior of the state due to the easier access to alternative shopping sources. The descriptive regressions do not yield conclusive evidence to that effect.

Table 1 and Table 2 provide clear evidence for the mechanisms that underlie our story: having more potential customers nearby raises demand, as does their proximity to their nearest store. Higher prices reduce demand, via the demand curve. We now turn to a simple model to estimate these effects, allowing us to predict sales under alternative store configurations.

4. A Simple Model of Demand with Travel Cost

We seek a model that, for any set of store locations, can indicate both the demand and producer and consumer surplus from consumption by individuals in each piece of geography. The key behavioral relationships that the model must describe are a) the sensitivity of demand to consumers' distance to stores and b) the price elasticity of demand, which allows us to attach a dollar value to proximity. We could directly relate quantities sold at a store to, say, population in its area and other demand shifters, such as median income in the area. Table 2 reports such regressions, but they cannot be used to predict sales under a counterfactual set of stores or locations and, in turn, to calculate the change in consumer surplus from an additional store or a change in store configuration. This goal, instead, requires a model of demand at the level of geography where consumers reside. We use a discrete-choice framework to model the consumer's decision to purchase liquor and estimate its parameters based on the PLCB's current stores to address these requirements.

a. Demand and Distance

There are S stores located around the state. We assume that a consumer i patronizes the store s nearest his residence. This assumption, which would arise endogenously if stores were identical in selection, given that pricing is identical across stores, divides the state into S catchment areas containing all of the population nearest to each store. We make this assumption, as well as several others, to facilitate the determination of optimal store configurations, discussed below.

We denote each store's catchment area by C_{st}. Formally, C_{st} is the set of consumer locations r such that store location s is the closest to location r on day t, or $C_{st}\colon \{d_{rst} = \min_{s'} d_{rs'} \forall s' \in S_{open,t}, \forall r = 1, \dots, R\}$, where d_{rs} denotes the distance, measured in an appropriate metric, of consumer i in location r from the store's location s, for all stores open on day t and contained in set $S_{open,t}$. We discretize consumer locations in the state by modeling demand at the level of the Census tract and place all residents at each tract's centroid. We then assign Census tracts to store catchment areas by finding the store s whose street address is closest in distance to each tract centroid. The use of Census tracts – relative to finer divisions of the state such as Census block groups – introduces some measurement error into the distances consumers travel. It yields, however, a more manageable set of 3,125 consumer locations, which we also use as potential store locations in the simulations that follow.

Our lack of data on individual purchases prevents us from distinguishing between the decision to visit a store and the decision of how many bottles to purchase.[17] Instead, we assume that consumers purchase a single bottle of liquor during a shopping occasion and model consumer i's conditional indirect utility from traveling to store s on day t to purchase a bottle of liquor as:

$$V_{ijrst} = X_{jrt}'\beta_x - \beta_d d_{rs} - \beta_p p_t + \varepsilon_{ijrst}. \tag{1}$$

We aggregate consumers to demographic groups j. In equation (1), X_{jrt} is a vector of attributes for consumers of type j in location r and seasonal effects. The term ε_{ijrst} denotes an unobserved utility shifter that we assume to be distributed extreme value. Prior studies of alcohol demand suggest that demand varies with age, income, and the racial composition of households (see e.g., Heien and Pompelli 1989 and Wang, Gao, Wailes, Cramer 1996). Consequently, we differentiate between black (B) and other-race residents (O) and include among the X_{jrt} the group's per-capita income and median age using data from the Census 2000.

A consumer chooses to purchase from his store provided that his utility exceeds the utility of the outside option of not purchasing. We normalize its value to zero. Our assumption of extreme-value distributed ε_{ijrst} yields Logit purchase probabilities for consumers of each demographic group j in a particular location r:

$$s_{jrst} = \frac{\exp\left(X_{jrt}'\beta_x - \beta_d d_{rs} - \beta_p p_t\right)}{1+\exp\left(X_{jrt}'\beta_x - \beta_d d_{rs} - \beta_p p_t\right)}. \tag{2}$$

To derive a store's predicted demand, we aggregate over the decisions of potential consumers across demographic groups j within a tract location and across all of the locations that make up a store catchment area, C_{st}. We consider as potential consumers the population of each

[17] A further downside to observing store, rather than consumer, level data is that we cannot explore the extent to which people who live further from a store choose to make fewer, but larger, shopping trips and store the product more.

Census tract over the age of 21. Aggregate demand for liquor in tract r, $\hat{Q}_{rst}$, and at store s, $\hat{Q}_{st}$, is thus the weighted average probability of purchase across demographic types and, for the store, across tracts, using as weights each tract's mass of consumers of a particular type, scaled up by the total potential consumers:

$$\begin{aligned} \hat{Q}_{rst} &= \sum_{j=\{B,O\}} \mathcal{s}_{jrst}(X_{jrt}, d_{rs}, p_t|\beta) M_{jrt} \\ \hat{Q}_{st} &= \left(\sum_{r\in C_{st}} \frac{\hat{Q}_{rst}}{M_{st}}\right) M_{st} = \mathcal{s}_{st} M_{st}, \end{aligned} \tag{3}$$

where M_{jrt} denotes the number of potential consumers of type j in tract location r and $M_{st} = \sum_{r\in C_{st}} \sum_{j=\{B,O\}} M_{jrt}$ the potential consumers in the store's aggregate catchment area.

Estimation proceeds via maximum likelihood. The parameter estimates maximize the likelihood of observing actual sales in store s on day t, Q_{st}, given data on the price of liquor and on the demographics and distance from the store of the locations making up the store catchment area. The log-likelihood function is given by:

$$\ln\mathcal{L} = -\sum_{t=1}^{T} \sum_{s=1}^{S} I(open_{st})(Q_{st} \ln(\mathcal{s}_{st}) + (M_{st} - Q_{st}) \ln(1 - \mathcal{s}_{st})), \tag{4}$$

where $I(open_{st})$ is an indicator of whether store s is open on day t. We identify the parameters from observing variation in the price of liquor over time (β_p) and cross-sectional and time-series variation in the composition of catchment areas, resulting in variation in distances traveled (β_d) and demographic attributes (β_x).

b. Demand Model Estimates

To keep the estimation manageable, we rely on a randomly drawn 10% subset of the daily data.[18] Beyond age, race, and income, we proxy for variation in local attitudes toward liquor consumption by including in utility each tract's number of churches per-capita, derived from a state-wide listing of religious organizations from Reference USA.

We capture travel costs by considering three different distance-based measures. First, we use distance traveled from the centroid of each tract to the store along the existing road network. We compute the distances based on the shortest route between two locations, using the program

---

[18] For the descriptive regressions in Table 2, the estimated parameters using the subsample do not differ significantly from the results obtained using the full sample of data.

MPMileage. We do so for the distances between all Pennsylvania tracts and the existing stores, as well as – for the purposes of computing demand under counterfactual store configurations below – between all tracts themselves.[19] Second, we employ the great-circle distance in kilometers between locations. MPMileage further generates the average travel time in minutes between any two locations, which we use as our last travel cost proxy.

We allow for systematic variation in the travel cost depending on features of the consumer's place of residence by interacting the distance to the store with the percent of tract households that lack a car. This specification reflects that the mode of transportation to the store may differ between residents of cities and those in less urban areas.

As in our descriptive regressions, we address the fact that the PLCB may time sales and thus price reductions to coincide with unobserved temporal variation in liquor demand by employing the list-price prior to sales as our price index for the composite liquor product. We also control for seasonal effects by including day of the week effects, week dummies for holiday weeks (the week after New Year's (week 1), July Fourth (week 26), Thanksgiving (week 47), and Christmas through New Year's Eve (weeks 50-52)), and additional holiday dummies for Memorial Day (May 28, 2005), days close to July 4 (June 30, July 1 – July 3, 2005), Labor Day (September 3, 2005), and days around Thanksgiving (November 23 – 26, 2005). The price elasticity is thus identified from a response in sales to price changes in otherwise similar days.

Driving distance is, not surprisingly, systematically larger than, but closely related to, great-circle distance. A regression of driving distance to the closest store on great-circle distance to the closest store for each of Pennsylvania's 3,125 tracts indicates that each additional kilometer of great-circle distance adds 1.4 km of driving distance, with an R^2 of 0.94. The regression also indicates that, on top of the aspect of driving distance that is proportional to great-circle distance, driving distance contains an additional 0.2 km, or systematic deviations from proportionality. These deviations from proportionality leave room for possible differences in the estimated demand models using the alternative distance measures.

The coefficients of the estimated demand function appear in Table 3. Column (1) reports results based on driving distance in km as our distance metric. We rely on these results in the remainder of the paper. Columns (2) and (3) report the results based on great-circle distance and

[19] Due to the computational burden of computing driving distances for 3,125×3,124 tract combinations, we calculate exact driving distances only for the 25 tracts nearest each consumer tract location based on straight-line distance. We use an approximation based on an estimated linear relationship between driving and straight-line distance for more distant tracts. In our simulations of alternative store networks below, consumers in all tract locations are typically assigned to a store in one of their neighboring ten tracts for store configurations of plausible size.

driving time in minutes, respectively. Most of the parameters are stable across specifications. The estimated price coefficients translate into an average price elasticity of –1.18 to –1.48, similar to the estimates in Table 1 and Table 2.

In specification (1), the estimated parameters on distance and distance interacted with the percent of the population without access to a car imply that demand declines by 61 (98) cents for every kilometer (mile) driven to the store for a tract with the median share of households without car access of 8.18%. Based on straight-line distance in column (2), we estimate a travel cost of 84 cents per kilometer of straight-line distance to the store. The increase in the estimated effect relative to the driving distance model reflects that driving distance is typically larger than great-circle distance. The estimated travel cost is similar to the implied travel cost under driving distance when scaled down by the factor of proportionality of 1.4 above, resulting in an equivalent travel cost of 60 cents per kilometer of driving distance. In the driving-time model in specification (3), we estimate an implied travel cost of 50 cents per minute added to each leg of a round trip to the store. Based on customers traveling between 35 and 50 km per hour, this translates into a cost per kilometer of driving distance of 50 to 86 cents. Our alternative distance specifications thus result in relatively similar travel costs.

The results suggest further that travel costs increase in the percentage of households without a car; based on the 10^{th} and 90^{th} percentiles, travel costs per kilometer range from 39 cents (when virtually all households have access to a car) to 157 cents (when 35% of households do not have access to a car). The decline of travel cost with greater car access reflects the time difference between driving and its alternatives and lends credence to a travel-cost interpretation of our distance coefficient.

Our travel cost estimates are consistent with the existing work, although the literature contains a relatively wide range of travel cost estimates. Davis (2006) estimates that a consumer who travels 3.2 km in total incurs a travel cost of approximately 20 cents per km of great-circle distance, while Thomadsen (2005) finds travel costs of $1.86 per kilometer of driving distance. McManus (2007) finds that consumers are willing to pay $4 to avoid walking an additional mile from their location to the retail outlet reflecting the increase in time spent to cover one mile walking relative to driving. Houde (forthcoming) estimates that an average consumer is willing to add 1 minute in travel time to save 67 cents on a purchase of 20 liters of gasoline.

Across specifications, areas with higher median income have higher demand; demand does not vary significantly with age. While the point estimates suggest that demand is lower in areas with a larger number of churches per capita and a lower percent of black households, the effects are statistically significant at conventional levels for specification (2) only.

While we rely primarily on specification (1) in the simulations that follow, we investigate several alternative specifications of our travel demand model. Across specifications, the price and distance coefficients are similar to the ones in Table 3. First, we estimated a variant of specification (1) based on both daytime and evening / weekend population, allowing consumers a choice of consuming either from their place of residence or from their place of work. The estimates are similar to the main results with a slightly lower demand elasticity of –1.59.

Second, we test the role of various alternative determinants of demand to ensure that their effect does not get absorbed by our main demand drivers, most notably distance. We allow demand to vary between rural and urban tracts and with the population density of the county. We investigate whether the presence of fundamentalist churches (as classified in Smith 1986), whose congregants might place a higher value on limited alcohol consumption than church congregants in general, is a stronger proxy for demand than aggregate church density. In both cases, the additional regressors were not statistically significant in affecting demand and travel cost remained stable, ranging from 53 to 64 cents per km of driving distance.

Third, we consider various, more flexible specifications for travel costs and the price coefficient. We approximate the cost of travel with a quadratic distance specification. The price elasticity under this alternative specification is –1.34 and the travel cost implied by the quadratic specification increases slightly in distance. At the mean distance of consumer to store locations, it amounts to 66 cents per km of driving distance, similar to the estimates above, and ranges from 60 cents to 67 cents for the 25th and 75th percentile of distances traveled, respectively.

We consider whether consumers are less sensitive to distance traveled when they are able to combine the trip to the liquor store with other shopping occasions. Results including interactions of distance with the number of grocery stores or the number of discounters in the liquor store's tract do not suggest, however, that consumers are willing to travel a larger distance to liquor stores in close proximity to other similar retailers. For the median tract, travel costs remain at 61 cents per km of driving distance. We further do not find significant evidence that the distance coefficient varies significantly with tract income. Lastly, we allow the price coefficient to vary with tract income. Our results suggest that demand is less responsive in higher-income areas with an interquartile range for the price elasticity of −1.65 to −1.36 for consumers in tracts with the 25th and 75th percentiles of income. Appendix A provides details on the specifications, data sources, and, in Table A-1, the results of these alternative demand specifications.

c. Welfare Measures

To evaluate openings or closures of stores and changes in store locations, we need to compute the welfare benefit of alternative store configurations. Our model shows how much the demand by persons in each location (and, by extension, the quantity sold at each store) changes with the distance to the closest store. Opening a store near location r has two effects on consumer welfare: First, current purchasers in and close to location r face a lower effective price (inclusive of travel). Second, a larger share of consumers in location r purchase under the lower effective price. This generates additional consumer surplus.

For the chosen specification, daily consumer surplus (CS) for consumers in location r is given by:

$$CS_{rst} = -\frac{1}{\beta_p}\left(\sum_{j=\{B,O\}} \ln\left(1 + \mathrm{e}^{X_{jrt}^{'}\beta_x - \beta_d d_{rs} - \beta_p p_t}\right) M_{jrt}\right) \tag{5}$$

if store s serves tract location r (see Small and Rosen 1981). The consumers in location r generate daily producer surplus (PS) to the store, based on the markup of the retail price p_t over the wholesale price c_t:

$$PS_{rst} = (p_t - c_t)\hat{Q}_{rst}. \tag{6}$$

The daily total surplus (TS) generated by store s is therefore:

$$TS_{st} = \sum_{r \in C_{st}} (CS_{rst} + PS_{rst}) - K, \tag{7}$$

where K denotes the daily fixed cost of operating a store.

d. Comparing Alternative Entry Patterns

To assess the goals underlying the PLCB's store configuration, we derive several benchmark configurations, including the store layout chosen by a profit-maximizing monopolist and a benevolent monopolist. These rely on the $R \times R$ matrix Y of consumer-location-to-store-location matches. We define Y_{rs} to be one if consumers in location r are served by a store in location s, and zero otherwise. The Y matrix also indicates S, the total number of stores operating, as trace(Y)=$\sum_{s=1}^{R} Y_{ss}$. We continue to assume in our simulations that locations are Census tracts. Since we do not observe a store in every tract in the data and do not model where within a tract the store would locate, we use each tract's centroid as a potential store's location.

For a given store configuration, the average daily profits of the system are then:

$$\Pi = \sum_{t=1}^{T} \frac{1}{T} \sum_{s=1}^{R} \sum_{r=1}^{R} PS_{rst}\, Y_{rs} - K \sum_{s=1}^{R} Y_{ss}. \tag{8}$$

The profit in Equation 8 includes two parts. The first, $\sum_{t=1}^{T} 1/T \sum_{s=1}^{R} \sum_{r=1}^{R} PS_{rst}\, Y_{rs}$, is the producer surplus that results from a particular configuration of stores and the rule that demand is assigned to its closest locations. The second part of the maximand is simply the fixed cost of operating the chosen number of stores. The profit-maximizing monopolist's problem is to find the store configuration that maximizes profit, while a benevolent monopolist's problem is to find the configuration that maximizes welfare (replacing PS_{rst} with $PS_{rst} + CS_{rst}$).

Solving this optimization problem is difficult because of the sheer number of possible store configurations. There are 2^R possible configurations of stores to evaluate. Even with a small set of possible locations, e.g. 25, there are over 33 million configurations. Operations researchers have developed efficient integer programming algorithms, such as "branch and bound," for solving problems of this sort.[20] We are able to rely on these sophisticated algorithms to solve problems of moderately large size.[21] Here we state the problem as an integer program; Appendix B provides an overview of the branch-and-bound algorithm we employ in finding optimal store configurations.

Expressed as an integer programming problem, the profit-maximizing planner's maximand is:

$$\max_Y \Pi = \sum_{t=1}^{T} \frac{1}{T} \sum_{s=1}^{R} \sum_{r=1}^{R} PS_{rst}\, Y_{rs} - K \sum_{s=1}^{R} Y_{ss} \tag{9}$$

subject to

$$\sum_{s=1}^{R} Y_{rs} = 1 \quad \forall r, \tag{10}$$

$$Y_{ss} \geq Y_{rs} \quad \forall r, s, r \neq s, \tag{11}$$

$$Y_{rs} = \{0,1\} \quad \forall r, s. \tag{12}$$

[20] We employ LINGO 13.0 to solve these problems.

[21] Our problem is closely related to the facilities location problem analyzed in Perl and Ho (1990). Chan, Padmanabhan, and Seetharaman (2007) employ the same integer programming techniques we use in their study of the optimal location choices of retail gas stations in Singapore where the regulator determines outlet locations, but then licenses the outlet operations to private firms. They illustrate how to estimate a reduced-form demand distribution across consumer locations from realized outlet locations under the maintained assumption that the government's objective is the minimization of the sum of consumer distances from their closest gas station and that actual location choices are optimal given this objective.

Constraint (10) indicates that each demand location must be assigned to a single store location. Constraint (11) prevents the assignment of demand to locations without a store. Constraint (12) makes the assignment of demand to supply binary: each demand location is either served by a particular supply location, or not. The alternative problem where the monopolist maximizes welfare less fixed costs can be expressed analogously.

Finding a solution via integer programming requires fixed coefficients on the binary store-location variables. Here, these fixed coefficients are the values of PS_{rst} and CS_{rst}. That is, we need to know the amount of demand or welfare that each demand location would contribute to each store in each possible configuration. We are able to calculate these coefficients in advance of the optimization because our demand model assigns each demand location to its nearest store. This would not be the case if we allowed consumers to choose not only whether to purchase liquor, but also from which store to purchase in a multinomial choice model of demand. Then a store's demand from any location would depend not simply on the distance between the store and demand locations but rather on the entire configuration of stores. That is, each PS_{rst} and CS_{rst} would depend on the entire 2^R configuration.

Integer programming approaches are strained by the problem of locating stores throughout the state's 3,125 tracts. We consider two alternatives. First, we find the optimal configurations on a county-by-county basis for each of Pennsylvania's 67 counties. We then aggregate across counties to derive profit, consumer welfare, and total welfare across the state. This procedure likely differs from the statewide optimal configuration in counties where a significant share of the population resides close to the county borders and might choose to consume out-of-county, which we preclude. As a second alternative, we find a statewide store configuration by turning to "greedy" algorithms, which provide intuitive and less computationally burdensome approaches (Daskin 1995). We implement such an algorithm, which we term "sequential myopic entry" (SME), as follows. Beginning from a first location that maximizes its standalone profits (or welfare) among the state's full set of tracts, we keep adding stores that maximize incremental profit (or welfare), holding the previous stores' locations fixed, until the marginal profit or welfare of the incremental location falls below the fixed cost of an additional store.

The SME configuration is not in general the same as the configuration that simultaneously maximizes the profit available from n stores. Sequential myopic entry overstates the benefit of each inframarginal entrant because its marginal benefit is – myopically – predicated on the $(n-1)$ stores already operating, rather than the total number that will ultimately operate. When the last store has been added, the marginal benefits of the infra-marginal stores are smaller than they were when the stores were marginal. To assess the magnitude of such biases, we compare results under sequential myopic entry with the simultaneous-move optima for small areas where these can be calculated.

e. Private Entry

In addition to examining profit and welfare maximizing store configurations, we would also like to explore configurations that would arise under atomistic private entry, either unconstrained or regulated to a constrained number of entrants. The usual condition for equilibrium with free entry by symmetric firms is that the S firms operating are each profitable while $(S+1)$ would not be.[22] Here, because of the vagaries of geography, equilibrium is more complicated. Each firm (store) must be profitable; there must be no room for further entry; and no firm may wish to switch its location.

A challenge in employing our estimates to assess a private alternative to the current regime is that in our empirical context, prices and markups are fixed and set by the State Legislature. This undermines our ability to predict the extent of spatial price competition in a free-entry alternative, and we continue to assume that firms charge the regulated price in the private entry context. Because the price-reducing mechanism usually present with free entry is absent, the model likely generates more stores than would actually operate if entry were truly unregulated. Hence, the number of firms under unregulated free entry from the model should be viewed either as an upper bound or as a simulation of a fixed-price regime, as might operate if the state regulated prices with an optimal Pigouvian tax.

Due to the computational burden of identifying the equilibrium in a simultaneous-move game of the size we consider, we employ a sequential myopic algorithm similar to those introduced above, although some adaptation is needed for free entry. First, we find the location that maximizes a lone store's revenue.[23] If this location is profitable, it remains. The second store locates at the location that generates the most profit, given the location of the first store. That is, the second store locates at its best response, evaluated given the first store's location. If either store is unprofitable, it is withdrawn. Then another store locates at the most profitable available location, and so on. The process ends when there is no location for profitable entry,

[22] This is the condition for equilibrium in homogeneous goods entry models such as Bresnahan and Reiss (1991) and Berry (1992). Entry models dealing with product positioning include Mazzeo (2002) and Seim (2006).

[23] We investigated the sensitivity of the resulting configuration to our choice of the initial store's location using one Pennsylvania County as a case study. Configurations that result from starting the SME algorithm in each of the County's tracts in turn result in an identical final configuration in all but one instance that differs in the location of a single store.

and each existing store is profitable.[24] This deviates from Nash equilibrium because the stores, while profitable, might be more profitable if they switched locations. Only the last entrant is necessarily on its best response function. Still, the algorithm shows – approximately – how many stores free entry could support.

This algorithm is clearly neither fully rational nor – as a result – fully optimal. When stores enter, they find the location that is currently most profitable, given existing entry. Entrepreneurs do not anticipate, however, how subsequent entry will affect the profitability of the locations they choose and continue operating until they are rendered unprofitable by other – unforeseen – entry. Still, it seems reasonable to expect, if S simultaneously operating stores are profitable, that the free entry equilibrium has at least S stores. Even this simple algorithm is somewhat computationally challenging since in each iteration, we must check the profitability of each store (rather than just the entire system).

5. Effect of State Control on Liquor Retailing

We have already seen, in Section 2, that the PLCB operates fewer stores than would likely exist under a private system. Our goal here is to quantify the welfare and distributional consequences of the PLCB using our demand model along with our characterization of private entry. To this end, we compare a model simulation of the current PLCB configuration against one of two plausible alternatives: first, privatization of liquor retailing in Pennsylvania holding the liquor tax at current levels and second, free entry under a reduced liquor tax that is typical for private states, using the national average tax rate.[25] Given, as discussed in Section 4.e, that we hold prices fixed in our free-entry simulations, the lower liquor tax is equivalent to a higher variable profit per bottle.

We calculate each store's variable profit using the demand estimates from our main specification in column (1) in Table 3. We set the retail and wholesale prices to their mean values in 2005 with p = \$12.38 and marginal cost c = \$7.31. For the privatization simulations we initially presume that the current tax structure would remain in the absence of the PLCB system. Of the \$5.07 difference between average retail and wholesale prices for a bottle, \$1.89 is

---

[24] Our free entry simulations do not always converge to a single configuration. Instead, they generally cycle among a small number of possible configurations. For example, the free entry simulation with a fixed cost threshold of \$618 eventually cycles among eight possible configuration sizes: 1,523, 1,524, …, and 1,530. Once the cycling begins, 95 percent of iterations produce configurations of between 1,525 and 1,529.

[25] These differ from current private systems in that we consider the issuance of a statewide pool of liquor licenses, while in practice governments commonly allocate licenses at the level of the municipality or county. For an overview of state policies, see Toma (1988).

liquor tax, while the remaining $3.18 is variable profit. Then we reduce taxes to private-system rates, where we assume a liquor tax of $0.81 per bottle, leaving $4.26 as variable profit.

The actual PLCB system has 621 stores in 603 distinct locations. The model simulation of the actual system predicts that the sale of 256,502 bottles per typical day generates $10.50 million in consumer surplus. Each day the system generates $0.13 million in profit, along with $0.48 million in liquor tax and $0.35 million in labor surplus. The total of these three components, which we collectively term "total producer surplus/rents," is $0.96 million per day. See Table 4.

The comparison of Pennsylvania liquor retailing with other states in Section 2 suggests that if it had a private system selling wine and spirits, Pennsylvania would have substantially more stores: private states have on average 1,531 stores to serve markets with the size of Pennsylvania's wine and spirits market. Therefore, we would like to compare the welfare properties of the actual configuration with the properties of a private entry configuration of the predicted size.

Privatized free entry with competitive daily fixed costs of $549 – and retaining the current liquor tax – gives rise to a system with 1,290 stores. Consumption is 301,172 bottles per day, nearly a fifth above its current level. Consumer surplus is $10.78 million per day, while private profit is $0.25 million per day, positive only because of integer constraints. Daily liquor tax revenue is $0.57 million, and there is no labor surplus. Total producer surplus is thus $0.82 million per day. That is, privatization that retains the current liquor tax would increase overall surplus relative to its PLCB level by 4.6% of consumer expenditure: consumers would gain by having more stores, while workers would lose their above-competitive payments. The free entry configuration has significant duplication: the 1,290 stores operate in only 1,112 distinct locations (tracts). Locations with sufficient equilibrium demand to cover the costs of multiple stores get more than one.

We cannot directly choose the number of stores operating for our free entry algorithm. Instead, to use our model to generate a Pennsylvania more closely resembling a private state, we adjust the fixed-cost threshold that determines entry and use the algorithm to calculate the number of stores that can be sustained at that cost. The fixed-cost threshold can also be expressed in terms of number of bottles sold per day, with entering firms selling daily quantities in excess of the ratio of fixed costs to variable profit per bottle (excluding liquor taxes). After experimenting, we find that a bottle threshold of 145 produces a private entry configuration with

1,527 stores in 1,177 distinct tracts.[26] If Pennsylvania's liquor tax fell to the average level of other states, variable profit per bottle would rise to \$4.26; hence, store operation costs of \$618 would give this threshold (618/4.26=145). We interpret the excess of this \$618 over competitive costs of \$549 as the cost of having a liquor license, and the payment for the license is part of the fixed cost from operating a store.[27] In the resulting configuration, bottle consumption is 303,192 per day, again about a fifth above its current level; and consumer surplus is \$10.80 million. The system generates \$0.34 million in private profit, \$0.25 million in daily liquor tax, \$0.11 million in daily license rents, and no labor surplus. Total producer rents are \$0.70 million per day.

Relative to either Pennsylvania privatization retaining the current liquor tax or reducing it to typical private state levels, the PLCB system has three major effects. First, the PLCB substantially limits the number of stores, to 621 rather than 1,500 or more. This limitation on the number of stores reduces consumer surplus by about \$0.3 million per day, but it also raises total producer surplus. Second, the PLCB reduces consumption by about 15 percent. Third, the PLCB delivers a substantial labor surplus that would not exist with a private system. Aggregate welfare is lower by about 5% of expenditure under the PLCB to its value under the two forms of free entry considered here.

6. Comparison with Optimal Configurations and Implicit Motives

Given an objective for the planner and an assumption about store operation costs, we can also use our model to calculate the optimal Pennsylvania liquor store configuration. We are interested primarily in statewide estimates. However, as discussed in Section 4.d, we are able to calculate exact solutions only for smaller pieces of geography (individual counties) and aggregate them to the whole state or employ a simplified algorithm for the whole state. To compare the performance of these two algorithms, we first derive profit and welfare maximizing benchmark configurations for counties, calculated both exactly and using the simplified statewide algorithm. We then re-do this exercise at the state level, before turning to alternative planner objectives and assessing the PLCB relative to these objectives.

[26] Performing a grid search over values of the bottle threshold to find the exact threshold that entails the predicted size of the Pennsylvania liquor market from Section 2 is computationally taxing. We therefore rely on the free entry configuration resulting from the threshold of 145 bottles as an approximation.

[27] Under the competitive cost assumption, the additional payment for the liquor license is \$69 per day. On an annual basis, this implies a payment of roughly \$20,000. Discounting at 5-10 percent, this implies that the value of a liquor license is between \$200,000 and \$400,000. To get a sense of whether this implied license value is reasonable, we analyzed the listings of 51 liquor stores for sale (outside of Pennsylvania) at http://www.bizbuysell.com/liquor-stores-for-sale/ as of December 19, 2011. Removing the stated value of included fixtures, inventory, and real estate, the mean (median) asking price was \$473,294 (\$240,000).

a. *Exact County Estimates*

At the county level we can implement the integer programming approach to find efficient configurations. We derive optimal configurations under profit and welfare maximization, assuming the fixed store operating costs stay at current levels. We do this for five counties, and the leftmost columns of Table 5 summarize the exact welfare maximizing solution. The rightmost columns in Table 5 repeat the exercise using the sequential myopic entry algorithm for each county. The results are similar: the maximum welfare under the myopic algorithm is within 0.5% for all five counties. The comparison of the profit maximizing configurations yields comparable results.

b. *Statewide estimates*

We now apply the two solution methods to calculating statewide efficient configurations, aggregating across counties in the case of the county-by-county efficient ("exact county") configurations. We begin by assuming that the true store operation cost is the competitive cost of $549 per day and that, from the planner's perspective, the entire $5.07 in gross variable profit, including taxes, contributes to its profit. We contrast profit and welfare-maximizing configurations under the exact county and the SME approaches. Table 6 reports the profit and welfare maximizing configurations from these respective approaches, and the welfare properties of the results are similar. While the respective welfare and profit maximizing configurations from the two approaches differ in size by 0.4 and 4 percent, the associated sales and welfare measures are within 0.5 percent of each other. In both cases, welfare maximization is achieved with a store network of approximately 1,120 stores. Profit maximization is accomplished with around 480 stores. In what follows we focus on the aggregation of the less computationally costly county-by-county exact results.

In the analyses so far, we derived the benchmark configurations that maximize profit and welfare. A range of efficient configurations results, however, if we consider that the state maximizes a weighted sum of profit (variable profit less store operation costs) and consumer surplus. That is, the state's objective function $W = PS + \gamma CS$, where γ is the weight that the planner attaches to consumer surplus relative producer surplus.[28] When $\gamma = 0$, this is simply profit maximization; when $\gamma = 1$, this yields welfare maximization (equal weights on profit and consumer surplus).

[28] Our tradeoff between consumer surplus and profit is reminiscent of the framework employed in Armstrong, Cowan, and Vickers (1994).

The term γ is the planner's willingness to trade off CS for PS, and it has a natural interpretation. By choosing different store configurations, the planner can generate a range of consumer and producer surplus. Initially, for small networks, both CS and PS rise when comparing a network with n stores to a network with $(n-1)$ stores until the network size reaches the profit maximizing monopoly configuration. As stores continue to be added, consumer surplus rises and producer surplus falls. The ensuing relationship between CS and PS is a welfare possibilities frontier. When we observe a chosen store configuration, we can use this Pareto frontier of profit and consumer surplus combinations to infer the planner's tradeoff between the two.

Calculating the frontier requires an assumption about the store operation cost facing the planner. One interpretation of the PLCBs' current super-competitive store operation cost is that the PLCB actually faces this cost as a constraint. A second interpretation is that the planner faces competitive costs but chooses to make higher store operation payments as a means of sharing profit with labor. These contrasting assumptions give rise to different welfare frontiers and therefore different interpretations of the system's current size.

We derive the Pareto frontier under competitive costs by calculating optimal store configurations and their welfare properties for a range of γ's between 0 and 3.5. Figure 3 depicts the resulting Pareto frontier, starting with the profit-maximizing network size; Table 7 details the welfare and profit maximizing configurations contained in the frontier. It is interesting to note that the welfare maximizing configuration, at 1,124 stores, is substantially smaller than the configurations that would likely obtain absent the current PLCB system. It also seems clear that welfare maximization with competitive costs – and treating the gross variable profits as profit – is a poor positive description of the current system.

We can also create a Pareto frontier based on current store operation costs (FC=$1,110). With this higher cost, pure welfare-maximization is achieved with 566 stores, while profit maximization is accomplished with 249 stores. Competitive wages, holding current employment levels constant, imply that profit maximization is achieved with 370 stores, while 883 are required for welfare maximization.

The Pareto frontiers provide a lens for viewing the current system size of 621. In the case of competitive costs, the point on the frontier corresponding to an efficient configuration with N=621 is achieved by maximizing $\pi + 0.18 \times CS$. Thus, the current system size would result from maximization by a planner who values profits 5.6 times more than consumer surplus and shares some of the gross profits with labor. In other words, pure profit maximization with competitive costs provides a rough characterization of the current system size. By contrast, using a Pareto frontier based on current costs, an efficient system of current size is achieved by maximizing $\pi + 1.17 \times CS$. That is, a planner facing current costs and seeking to maximize welfare would choose a system of roughly the current size.

Yet, both of these characterizations, based entirely on the number of stores operating, are incomplete. The profit or welfare maximizing configurations in Table 7 are those that lie on the Pareto frontier of efficient configurations. If the actual system were on the frontier, we could infer state motives from the frontier's slope. For example, a system maximizing producer surplus would indicate a disregard for consumers. But the actual system is well inside the efficient frontiers, foregoing 8.7% and 7.1% in welfare relative to the welfare maximizing configuration under current costs and the profit maximizing configuration under competitive costs.

A second informative comparison to the efficient frontier is thus one where we compare the actual system to an efficient system of equal size. Based on both the myopic and the exact county algorithms, the efficient system with 621 stores generates CS of \$10.32 million per day and profit of \$1.09 million. Relative to the optimal system of equal size, the actual system forgoes \$0.18 million, or 5.3% of expenditure, in daily CS and \$0.13 in daily profit. Consumption declines by 10%, as Table 7 shows.[29] The fact that the actual system is interior to the Pareto frontier suggests that the system's store configuration is not simply maximizing a weighted sum of producer and consumer surplus.[30] We turn to this question below, with an attempt to infer system motives.

Because the simulations in this section rest on a number of inputs, we explored the sensitivity of our results to some of our assumptions. First, we consider a setup where the rental expense contribution to store operating costs is allowed to vary with tract-level residential rents. The resulting configuration generates a welfare improvement over the actual configuration whose magnitude is within 0.1 percentage points from the welfare differences under the constant-cost specification. Second, to investigate whether dynamic adjustment costs to changing the store configuration, such as long-term leases, can explain the apparent locational inefficiencies, we use 1990 demographic data to predict the optimal store configuration at that time and compare it to the current configuration. The analysis provides little support for this explanation. Third, a closely related demand model to our main specification finds that our

[29] The amounts of CS and profit foregone are similar when we compare the actual configuration to one where we constrain the size to be the number of distinct locations the PLCB serves, or 603 locations.

[30] We say "suggests" rather than "indicates" because the actual system's distance to the Pareto frontier may also arise from model mis-specification.

results are robust to the use of the different distance metrics depicted in columns (1)-(3) of Table 3. Last, we investigated the sensitivity of our results to the chosen demand specification. We re-derived the welfare- and profit-maximizing configurations using an alternative demand specification that entailed an economically low travel cost of only 20 cents. While the optimal configurations under this demand system are 20 to 35% smaller in size than the ones in Table 7, the majority of welfare losses continue to stem from locational inefficiencies. Appendix C provides additional detail.

The model allows us one more exercise of interest, quantification of the welfare loss associated with free entry and a division of this loss into two parts: the overall loss from having too many stores in the wrong locations and the loss from having simply the wrong locations, for a given number of stores. We do this by comparing a welfare-maximizing configuration to a free entry configuration with an equal number of stores. One complication is that we cannot easily target a particular configuration size with free entry; but we can compare the free entry configurations in Table 4 with equal-sized efficient configurations.

A "fair" comparison of welfare maximization and free entry requires us to calculate profits analogously under both entry regimes. For welfare maximization we treat the entire gross variable profit per bottle as profit. Hence, we need to do the same for free entry. The ensuing free entry configuration, without any liquor tax, is thus useful as an evaluation of free entry; but it is not meant as a plausible characterization of a private liquor retailing system. The resulting free entry configuration has 2,230 stores. Daily consumption is 322,197 bottles, consumer surplus is high at $10.92 million per day, and profit is low: $0.41 million per day. Relative to the welfare-maximizing configuration (with 1,124) stores, free entry raises CS by $0.05 million and reduces profits by $0.57 million. Overall, free entry dissipates $0.52 million per day.

The theoretically familiar welfare loss from free entry with homogeneous goods (Mankiw and Whinston 1986 and Dixit and Stiglitz 1977) arises entirely from too many outlets. Here, where goods are distinguished by location, we are able to ask how much of this lost welfare is due purely to wrong locations as opposed to too many locations. To answer this we compare a free entry configuration with a given number of stores against a frontier configuration of equal size (see the bottom panel of Figure 3). We perform this comparison for both N= 1,124 (the size of our welfare maximizing configuration) and N=2,230 (the result of unconstrained free entry). We calculate the efficient 2,230-store configuration via our SME algorithm (see Table 7), and by experimentation with different bottle thresholds we determine a 1,130-store free entry configuration, reported in Table 4. At N=2,230, the aggregate welfare loss from wrong locations is $0.27 million, while the loss in the neighborhood of N=1,124 is $0.3 million. These losses are between 52 percent and 57 percent of the overall welfare loss from free entry in this context.

Hence, half of the loss from free entry in this context would arise from wrong locations; the other half would arise from too many stores. Our modeling setup is unusual in that we fix prices despite free entry. This feature will tend to increase the overall and locational welfare losses from free entry as incentives to enter remain artificially as entry occurs.

c. Deviations from Efficiency Implicit in Free Entry and the Actual System

We saw above that the actual system is interior to the Pareto frontier. This sub-optimality can arise because the system's store configuration favors some types of consumers over others. We can explore the nature of this favoritism using the distance between each tract and the nearest store. We would not expect these distances to be equal across tracts in an efficient system; rather, the distance to the nearest store in an efficient (on-the-frontier) configuration provides a benchmark measure of the efficient distance for consumers to their nearest liquor store. Define d_t^* as the distance between tract *t* and its nearest store in an efficient configuration, d_t^{FE} as the distance to the nearest store in a free entry configuration, and d_t^{PLCB} as the distance to the tract containing the nearest actual store. We can compare $(d_t^{FE} - d_t^*)$ across tracts with different characteristics to infer how atomistic free entrants regard different types of consumers. Similar analysis of $(d_t^{PLCB} - d_t^*)$ reveals the goals of the implicit PLCB planner. We compare configurations of equal size – the size of the PLCB system of 621 stores – to isolate the pure impact of entry rationale.

Free entry is well understood to foster potentially excessive entry in high-demand areas and to effect inefficiently insufficient entry in low-demand areas (Spence 1976). By contrast, a major ostensible PLCB goal is to offer service to Pennsylvania consumers located throughout the state, even if they live in remote locations. We would therefore expect free entry to deviate from an efficient configuration by favoring urban consumers and for the PLCB's chosen locations to reverse this market bias.

As Table 8 shows, a regression of $(d_t^{FE} - d_t^*)$ on the tract's rural population share along with tract median income produces a coefficient of 8.12 (s.e. =0.32) on the rural share, indicating that in a 100-percent rural tract, the nearest liquor store is 8.12 kilometers more distant under free entry than in an equal-sized efficient configuration.[31] The coefficient on median income is negative, indicating high-income tracts are closer to liquor stores under free entry, compared with the optimum. This confirms the free entry bias against low-demand areas. An analogous regression of $(d_t^{PLCB} - d_t^*)$ on the rural share produces a coefficient of 1.87 (s.e.=0.20), while the

[31] The analogous regression based on the 1,588-store free entry configuration produces a rural share coefficient of 2.17 (s.e. =0.17).

median income coefficient goes from -0.05 to 0.01. As under free entry, the rural coefficient indicates a bias against rural consumers relative to the efficient configuration of equal size. However, the coefficient is less than a quarter as large, indicating that the PLCB's configuration substantially mitigates the bias against rural consumers implicit in the free entry configuration equal in size to the actual configuration.

d. Direct Evidence of Politics

The PLCB is ultimately controlled by the Pennsylvania General Assembly, and there is speculation in the press that political considerations and lobbying play a significant role in store closings, countering the stated profit motives of the board.[32] Here we thus ask whether the PLCB's entry patterns reflect its political control. Oversight of the PLCB rests with the House Committee on Liquor Control, whose membership numbers 28 among the 203 General Assembly overall. In the 2005-2006 Session, the Liquor Control committee was in Republican control.[33]

We can locate each liquor store in its House and Senate district.[34] Of 203 districts, 198 contained a liquor store as of the end (start) of 2005 (198/203). Districts represented by Democrats have slightly more stores, although this difference is not significant. Of the 175 districts whose representatives did not serve on the Liquor Control committee, 97 percent had a store. All 28 of the Liquor Committee member-represented districts had a store, although this difference is not statistically significant (p-val=0.37).

We explore this more systematically by regressing the number of liquor stores in a House district on population, median income, the percent rural and percent black, using the years 2000-2005. When we include all years but do not include district fixed effects, then after accounting for demographic characteristics of districts, those represented by a Democrat have 0.86 additional stores, while those represented by a Liquor Control committee member do not have

[32] According to the January 1, 2008 Pittsburgh Post Gazette article, "LCB works in curious ways" (http://www.post-gazette.com/pg/08028/852743-85.stm, accessed October 17, 2008), then PLCB Chairman Stapleton "did allow that the board hears from legislators 'all the time' when a store closing or store transfer is in the works. 'A lot of times there is a strong belief by legislators that certain downtown areas should be served by a store,' he said." The article cites the example of Representative C. George who "has been an outspoken advocate for state stores in his district…[including] the store in Houtzdale, Mr. George's hometown, [that] has lost from $11,000 to $20,000 in each of the past three years, but, he vowed, 'I would fight tooth and nail against any plan that took that store out of our town.'"

[33] The committee had the following structure: a chair from each of the majority (Republican) and minority (Democratic) parties, four chairs of two subcommittees drawn from the two parties, a secretary drawn from each party, as well as twelve rank-and-file Republicans and 8 rank-and-file Democrats. We located stores in districts using the Find Your Legislator feature of the Commonwealth of Pennsylvania website See "Standing Committees of the House of Representatives, 2005-2006 Session."

[34] See http://www.legis.state.pa.us/cfdocs/legis/home/findyourlegislator/index.cfm?CFID=25192954&CFTOKEN=16631361 last accessed August 29, 2011.

more stores. Distinguishing Liquor Control committee members by party reveals a different pattern: Districts represented by a Democrat on the Liquor Control committee have 0.6 fewer stores, while those represented by a Republican Liquor Control member have 0.5 more stores. When we include district fixed effects, however, member party becomes insignificant, while committee membership now has a significant coefficient of 0.2. When committee membership's effect is allowed to vary by party, the minority party effect disappears, while the majority (Republican) party impact remains significant (0.3 additional stores). Appendix table A2 contains detailed results. Overall, there is only modest evidence of a political impact on store location decisions.[35]

7. Conclusion

The PLCB's retailing system provides a rare glimpse into government decisions about entry. Comparisons with other states indicate that states with private liquor retailing have lower labor costs and substantially more stores per capita. How does government operation affect outcomes? Under a private system, Pennsylvania would likely have 2.5 times as many stores. Using a simple spatial demand model we are able to compare the current system to plausible free entry configurations. The plausible counterfactual configuration would raise consumer surplus by 9% of current consumer expenditure, simply because more consumers would have a closer store. Privatization would have two distinct effects on the rents enjoyed by producers. First, with more stores operating, overall producer surplus would fall. But paying Pennsylvania liquor store employees at private state rates would eliminate the rents currently experienced by PLCB employees. A significant welfare aspect of privatization is thus pure redistribution as aggregate welfare increases by only 4.6%.

Based on the number of stores it operates, what is the PLCB currently doing in relation to theoretical benchmarks of welfare and profit maximization? If the planner faced competitive store operation costs, it would maximize welfare with nearly double the current number of stores. One can roughly rationalize the current configuration as welfare maximizing, if one takes the current super-competitive store operation costs as given. Alternatively, the current system is similar in size to a system that would maximize profits for a planner facing competitive costs and sharing some of the profits with employees. But the PLCB configuration is well below the consumer surplus-profit Pareto frontier, indicating that the implicit planner cares about something other than simply a weighted sum of profits and consumer surplus. While we cannot

[35] We also investigated possible political influence on the 2005 choice of which stores to operate on Sundays. Sunday store presence is systematically more likely in higher-income House districts, but political variables are not systematically related.

uniquely identify the motives of the planner, we find that the PLCB's choices serve to mitigate – but not eliminate – the bias of free entry against rural consumers. Satisfying political goals could be a further motive of the system, but we find little evidence of explicit political influence on store locations.

According to our estimates, the PLCB's choice to reduce the number of stores operating also reduces consumption by 15 percent. Because the consumption of alcohol creates substantial social costs for third parties, consumer surplus alone is an inadequate measure of consumption's impact. For example, Young and Bielinska-Kwapisz (2006) document an elasticity of 1.13 of traffic fatalities with respect to aggregate statewide alcohol consumption. Cook, Osterman, and Sloan (2005) find a 0.23 all-cause elasticity of mortality with respect to statewide alcohol consumption. While alcohol consumption could be controlled without state operation of liquor retailing – for example with strict entry regulation or high taxes – the PLCB's effective discouragement of alcohol consumption reduces social costs, and these effects may represent additional motives of the PLCB.

We have one other novel finding on the welfare loss from free entry. Usual estimates of the welfare loss from free entry are driven by the number of outlets. We are able to estimate the welfare losses from free entry arising from both the wrong number of stores and the wrong locations for stores. In our context, wrong locations alone produce half of this loss.

Our analysis has focused on the store location and network size considerations of a public versus a private system. In doing so, we abstract from other strategic choices. It should be noted again that the simulations in the paper take prices as given. We further treat stores as identical in selection and abstract from systematic differences between stores that would encourage consumers to patronize stores further afield than their closest. These choices are motivated in part by lack of systematic data on prices and variety choices by retailers in private states; we leave the comparison of product selection under the government system to what might result in a private Pennsylvania system to future research.

Figures and Tables

Figure 1: Scaled-Up Number of Liquor Stores vs State Population

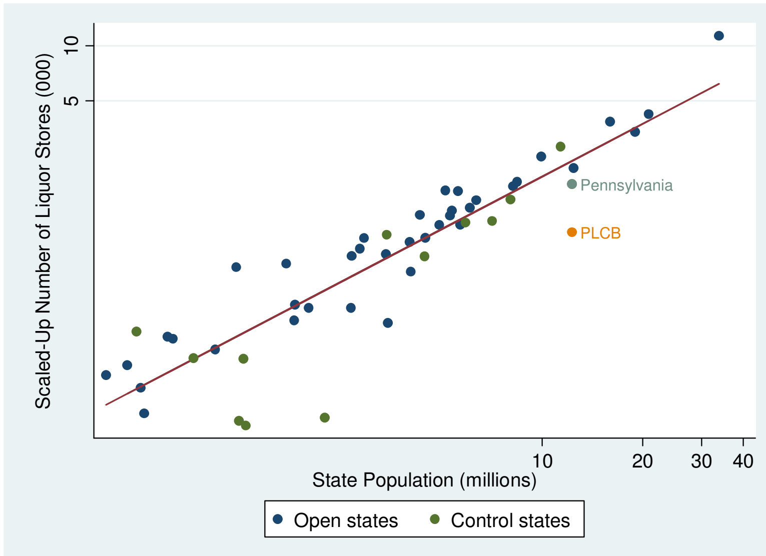

Figure 2: Weekly Average Number of Bottles Sold per Store, 2005

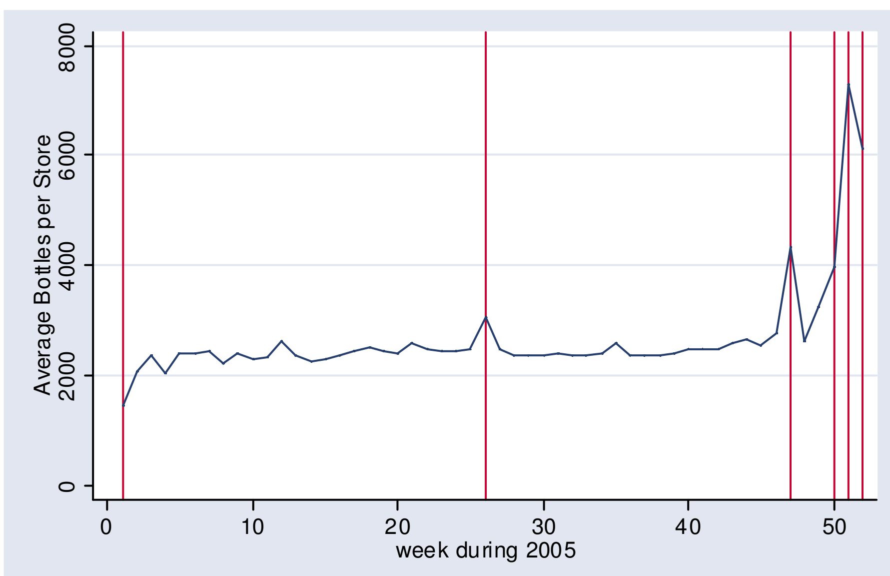

Figure 3: Consumer Surplus-Producer Surplus Frontier, Alternative Configuration Sizes

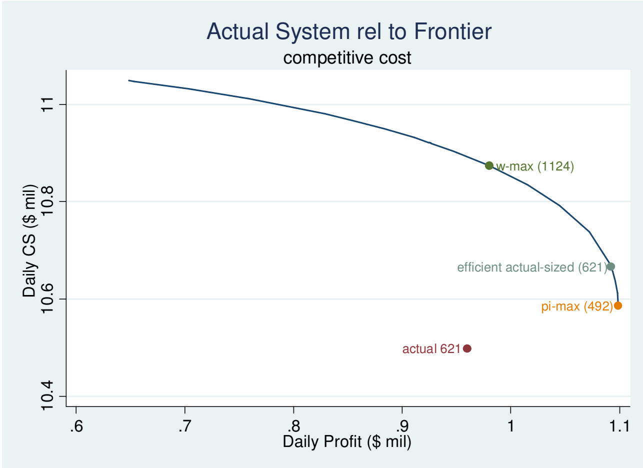

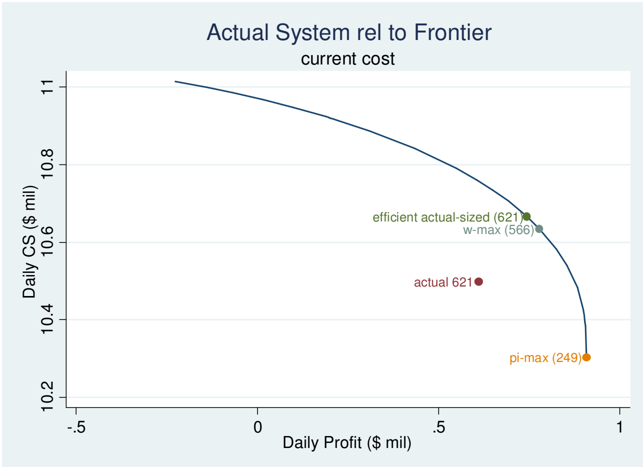

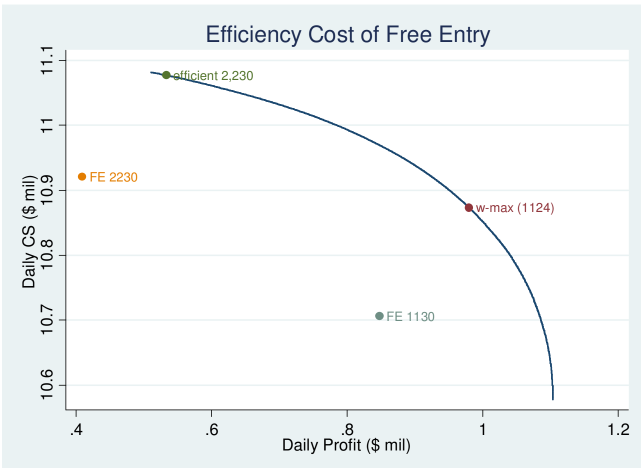

Table 1: Price Elasticity Evidence

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Log State Bundle List Price	-0.9627	-1.0900	-1.3184	-1.3396								
	(1.230)	(0.176)**	(1.356)	(0.181)**								
Log State Bundle Price					-1.8926	-1.9404	-0.8591	-0.8587	-0.7325	-0.7113	-1.4872	-1.0541
					(0.738)*	(0.105)**	(0.822)	(0.110)**	(0.387)	(0.110)**	(0.947)	(0.257)**
Constant	9.8875	10.2075	10.6897	10.7466	12.2274	12.3476	9.5368	9.5396	7.7585	6.9034	9.7677	7.882
	(3.095)**	(0.444)**	(3.410)**	(0.454)**	(1.857)**	(0.265)**	(2.071)**	(0.276)**	(0.973)**	(0.276)**	(2.379)**	(0.645)**
Observations	32,509	32,509	32,509	32,509	32,509	32,509	32,509	32,509	191,994	191,994	23,587	23,587
R-squared	0.06	0.77	0.07	0.80	0.06	0.77	0.07	0.80	0.19	0.77	0.07	0.38
Sample	weekly	weekly	weekly	weekly	weekly	weekly	weekly	weekly	daily	daily	daily, price changes	daily, price changes
Holiday weeks	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Quadratic time trend	no	no	yes	yes	no	no	yes	yes	yes	yes	no	no
Holidays									yes	yes	no	no
Day of the Week									yes	yes	yes	yes
Store FE	no	yes	no	yes	no	yes	no	yes	no	yes	no	yes

Notes: Dependent variable is log bottles per time period per store. Regressions of log bottles sold on various measures of the price. Holiday weeks include weeks 1, 26, 47, 50, 51, and 52. We include separate time trends for the period January - October and the holiday period of November - December. Standard errors in parentheses. * significant at 5%; ** significant at 1%. State-bundle prices use a constant bundle for computing the price and vary only by time and not across stores.

Table 2: Demand, Population, and Distance to the Nearest Store

	(1)	(2)	(3)	(4)	(5)	(6)
	Average Daily Sales per Store	**Average Log Daily Sales per Store**	**Log Daily Sales per Store**	**Log Daily Sales per Store**	**Log Daily Sales per Store**	**Log Daily Sales per Store**
Catchment Area Pop (0000)	173.7613	0.4485	0.1873	0.0359	0.1872	0.0359
	(12.4713)**	(0.0302)**	(0.0194)**	(0.0007)**	(0.0194)**	(0.0007)**
Average Distance to Nearest Store	-23.8609	-0.0812	-0.0768	-0.0163	-0.0768	-0.0163
	(2.8530)**	(0.0069)**	(0.0066)**	(0.0004)**	(0.0066)**	(0.0004)**
Log State Bundle List Price			-0.8458	-0.6573		
			(0.1759)**	(0.1284)**		
Log State Bundle Price					-0.6230	-0.6963
					(0.1386)**	(0.1092)**
Constant	273.7057	5.3648	7.5861	6.7124	7.0251	6.8094
	(25.1748)**	(0.0609)**	(0.4515)**	(0.3228)**	(0.3930)**	(0.2744)**
Observations	635	635	191,921	191,921	191,921	191,921
R-squared	0.27	0.33	0.33	0.77	0.33	0.77
Store FE	No	No	No	Yes	No	Yes

Table 3: Demand Model Estimates

	(1)	(2)	(3)
State-bundle list price	-0.1641	-0.1534	-0.1237
	(0.0488)***	(0.0768)**	(0.0984)
Driving distance	-0.0502		
	(0.0160)***		
(Driving distance)×(% w/o car)	-0.0060		
	(0.0025)**		
Straight-line distance		-0.0608	
		(0.0192)***	
(Straight-line distance)×(% w/o car)		-0.0083	
		(0.0028)***	
Driving time			-0.0615
			(0.0133)***
(Driving time)×(% w/o car)			-0.0007
			(0.0012)
Black	0.2143	0.2085	0.1652
	(0.1880)	(0.1723)	(0.1370)
Median Income	0.0341	0.0332	0.0325
	(0.0047)***	(0.0039)***	(0.0045)***
Median Age	-0.0002	0.0001	0.0128
	(0.0002)	(0.0002)	(0.0127)
No Churches per capita	-0.1053	-0.1544	-0.2260
	(0.0892)	(0.0601)**	(0.4277)
Monday	0.5009	0.5165	0.5932
	(0.0569)***	(0.0581)***	(0.0672)***
Tuesday	0.5462	0.5631	0.6269
	(0.0607)***	(0.0614)***	(0.0728)***
Wednesday	0.6629	0.6807	0.7480
	(0.0587)***	(0.0597)***	(0.0728)***
Thursday	0.8189	0.8337	0.9190
	(0.0578)***	(0.0600)***	(0.0692)***
Friday	1.3684	1.3839	1.4550
	(0.0599)***	(0.0582)***	(0.0669)***
Saturday	1.3925	1.4139	1.4874
	(0.0563)***	(0.0562)***	(0.0649)***
Implied elasticity of demand	-1.2668	-1.1836	-1.4779
Implied travel cost ($) per unit	0.6064	0.8405	0.4976

Note: Results based on daily store-level data for a 10% subset (19,255 obs) of all store-day observations. Bootstrapped standard errors (50 replications). We include separate holiday effects for May 28, June 30-July 3, Sept 3, and Nov. 23-26. Specification (1) uses the shortest distance in km along the road network; specification (2) the straight-line distance in km; and specification (3) the travel time in minutes associated with the shortest travel distance.

Table 4: Statewide Comparisons: Actual and Free Entry Configurations

Configuration	stores	bottles	consumer surplus	total producer rents	total surplus
Actual (cost=$1,110)	621	256,502	10,498	960	11,458
Private Pennsylvania (Liquor Tax = Current Liquor Tax)					
Free Entry (FC = $1,110)	527	254,885	10,433	1,003	11,436
Free Entry (FC = $713)	906	279,633	10,638	920	11,558
Free Entry (FC = $549)	1,290	301,172	10,784	819	11,603
Welfare Maximizing Planner Configurations of the Size of Privatized Configurations					
N^{FE} under FC = $1,110	527	274,507	10,615	1,102	11,717
N^{FE} under FC = $713	906	303,454	10,799	1,041	11,840
N^{FE} under FC = $549	1,290	322,287	10,920	926	11,846
Private Pennsylvania (Liquor Tax = Average Tax of Private States)					
Free Entry (FC = $618)	1,527	303,192	10,797	699	11,495
Free Entry Targeting W-Max	1,130	289,533	10,706	848	11,554
No liquor tax					
Free Entry (FC = $549)	2,230	322,197	10,921	409	11,330

Notes: Consumer surplus, total producer rents, and total surplus reported in 000s. Producer rents calculated as the sum of variable profit under the given tax structure, labor surplus ((FC−$549)×number of stores), and tax revenue, less the total store operating costs. Welfare-maximizing planner configurations of the size of the free-entry configurations derived using the SME algorithm.

Table 5: Performance of Myopic Algorithm:
Comparison of Welfare Maximizing Configuration for Five Counties

	Exact Equilibrium Configuration				**Sequential Myopic Entry Configuration**				
County	**No stores**	**Net Welf**	**Q**	**Profit**	**No stores**	**Net Welf**	**Q**	**Profit**	**Net Welf Dev**†
Berks	22	341,055	8,772	20,055	21	340,977	8,668	20,637	-0.02%
Blair	5	112,859	2,167	5,435	5	112,791	2,161	5,405	-0.06%
Lancaster	20	406,433	9,130	24,090	21	405,455	9,141	23,033	-0.24%
Lycoming	6	102,885	2,072	3,843	6	102,529	2,040	3,682	-0.35%
Schuylkill	10	142,324	2,932	3,766	10	142,324	2,932	3,766	0%

†The welfare deviations are calculated as the percentage change in welfare in going from the welfare under the welfare-maximizing configuration to the welfare under the configuration predicted by the myopic algorithm to maximize welfare.

Table 6: Comparison of Exact County and SME Algorithms: Statewide Estimates

Configuration	solution algorithm	stores	bottles	consumer surplus	total producer rents	total surplus
Welfare max N	Exact county	1,124	315,017	10,873	980	11,853
	SME	1,120	314,806	10,872	981	11,853
Profit maximizing N	Exact County	492	269,985	10,586	1,099	11,685
	SME	473	268,924	10,579	1,104	11,682

Notes: Consumer surplus, total producer rents, and total surplus reported in 000s. Fixed costs set to competitive level of $549. Producer rents calculated as the sum of variable profit including taxes, less the total store operating costs.

Table 7: Statewide Comparisons: Actual and Efficient Configurations

Configuration	stores	Bottles	consumer surplus	total producer rents	total surplus
Actual	621	256,502	10,498	960	11,458
Efficient Configurations					
N = locations with at least 1 store	603	278,941	10,643	1,083	11,726
N = N^{actual}	621	282,098	10,667	1,092	11,759
Welfare max N (FC=$1,110)	566	277,521	10,634	1,096	11,730
Welfare max N (FC=$713)	883	301,954	10,790	1,046	11,836
Welfare max N (FC=$549)	1,124	315,017	10,873	980	11,853
Profit maximizing N (FC=$1,110)	249	233,858	10,303	1,049	11,352
Profit maximizing N (FC=$713)	370	254,897	10,490	1,089	11,579
Profit maximizing N (FC=$549)	492	269,985	10,586	1,099	11,685
SME Welfare Max Targeting FE	2,230	346,654	11,077	533	11,610

Notes: Consumer surplus, total producer rents, and total surplus in 000s. Producer rents equal the sum of variable profit, taxes, and labor surplus ((FC−$549)×N), less fixed costs. All efficient configurations calculated using exact county algorithm.

Table 8: Distance Traveled: Actual and Free Entry vs Optimal Configurations

	(1) $d_t^{PLCB} - d_t^*$	(2) $d_t^{FE} - d_t^*$
Rural Share	1.8728	8.1224
	(0.1962)**	(0.3202)**
Median income (000)	0.0116	-0.0451
	(0.0037)**	(0.0061)**
Constant	-0.6304	2.2248
	(0.1773)**	(0.2894)**
Observations	3,123	3,123
R-squared	0.03	0.18

Notes: Standard errors in parentheses.* significant at 5%; ** significant at 1%. Dependent variables defined in text.

References

Armstrong, Mark, Simon Cowan, and John Vickers, Regulatory Reform: Economic Analysis and British Experience. MIT Press, Cambridge, MA. 1994.

Berry, Steven T., "Estimation of a Model of Entry in the Airline Industry." *Econometrica* 60 (1992): 889-917.

Boardman, Anthony E. and Aidan R. Vining, "Ownership and performance in competitive environments: A comparison of the performance of private, mixed, and state-owned enterprises." *Journal of Law and Economics* 32 (1989): 1-33.

Bresnahan, Timothy F. and Peter C. Reiss, "Entry and Competition in Concentrated Markets." *Journal of Political Economy* 99 (1991): 977-1009.

Chaloupka, Frank J., Michael Grossman, and Henry Saffer, "The Effects of Price on Alcohol Consumption and Alcohol-Related Problems." *Alcohol Research and Health* 6(1) (2002): 22-34.

Chan, Tat Y., V. Padmanabhan, and P.B. Seetharaman, "An Econometric Model of Location and Pricing in the Gasoline Market." *Journal of Marketing Research* XLIV (2007): 622-635.

Cook, Phillip J. and Michael J. Moore, "Alcohol." *NBER Working Paper* 6905 (1999).

Cook, Phillip J., Jan Osterman, and Frank A. Sloan, "Are alcohol excise taxes good for us?: short and long-term effects on mortality rates." *NBER Working Paper* 11138 (2005).

Daskin, Mark, Network and Discrete Location: Models, Algorithms and Applications, John Wiley and Sons, Inc., New York. 1995.

Davis, Peter, "Spatial competition in retail markets: movie theaters." *RAND Journal of Economics* 37(4) (2006): 964-982.

Dixit, Avinash and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67 (1977): 297-308.

Heien, Dale M. and Greg Pompelli, "The Demand for Alcoholic Beverages: Economic and Demographic Effects," *Southern Journal of Economics* 55 (1989): 759-770.

Holmes, Thomas J., "The Diffusion of Wal-Mart and Economies of Density." *Econometrica* 79 (2011): 253-302.

Hotelling, Harold, "Stability in Competition," *The Economic Journal* 39(153) (1929): 41-57.

Houde, Jean-Francois, “Spatial differentiation and vertical mergers in retail markets for gasoline,” forthcoming, *American Economic Review*.

Jia, Panle, “What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry.” *Econometrica* 76 (2008): 1263-1316.

Mankiw, N. Gregory and Michael D. Whinston, “Free Entry and Social Inefficiency.” *RAND Journal of Economics* 17(1) (1986): 48-58.

Mazzeo, Michael, “Product Choice and Oligopoly Market Structure.” *RAND Journal of Economics* 33(2) (2002): 241-242.

McManus, Brian, “Nonlinear Pricing in an Oligopoly Market: the Case of Specialty Coffee.” *RAND Journal of Economics* 38(3) (2007): 512-532.

Miravete, Eugenio, Katja Seim, Jeff Thurk, and Joel Waldfogel, “Complexity, Efficiency, and Fairness in Multiproduct Liquor Pricing.” working paper (2012).

Perl, Jossef and Peng-Kuan Ho, “Public Facilities Location under Elastic Demand.” *Transportation Science* 24(2) (1990): 117-136.

Seim, Katja, “An Empirical Model of Firm Entry with Endogenous Product-type Choices.” *RAND Journal of Economics* 37(3) (2006): 619-640.

Small, Kenneth and Harvey Rosen, “Applied Welfare Economics with Discrete Choice Models.” *Econometrica* 49(1) (1981): 105-130.

Smith, Tom W. “Classifying Protestant Denominations.” GSS Methodological Report No. 43 October, 1986 (http://www.icpsr.umich.edu/GSS99/report/m-report/meth43.htm).

Spence, Michael, “Product Selection, Fixed Costs, and Monopolistic Competition.” *Review of Economic Studies* 43 (1976): 217-236.

Thomadsen, Raphael, “The Effect of Ownership Structure on Prices in Geographically Differentiated Industries,” *RAND Journal of Economics* 36(4) (2005): 908-929.

Toma, Eugenia, “State liquor licensing, implicit contracting, and dry/wet counties,” *Economic Inquiry* (26) (1988): 507–524.

Wang, Jirong, X. M. Gao, Eric Wailes, Gail Cramer, “U.S. Consumer Demand for Alcoholic Beverages: Cross-Section Estimation of Demographic and Economic Effects,” *Review of Agricultural Economics* 18(3) (1996): 477-489.

Young, Douglas J. and Agnieszka Bielinska-Kwapisz, “Alcohol Prices, Consumption, and Traffic Fatalities.” *Southern Economic Journal* 72(3) (2006): 690-703.

Appendix

a. Alternative Demand Models

Here, we provide additional detail and present the results of alternative demand specifications that we investigated. Table A-1 below contains estimates for seven demand specifications that employ driving distance as the measure of travel cost.

Specification (1) allows the price coefficient to vary with the log of per-capita income of the tract's black and non-black residents. Specification (2) instead lets the distance coefficient depend on the log of income.

For specification (3), we collected information from ReferenceUSA on the number of grocery stores in the tract, by downloading records for all firms listed in SIC code 541105 and reporting sales of more than $2.5 million, which effectively excludes convenience stores. We interact the number of grocery stores with distance to allow for consumer's increased willingness to travel a given distance to their liquor store if the trip allows multi-stop shopping. In unreported results, we replaced the number of grocery stores with alternative proxies for retail environment, including the tract's number of discounters (Kmart, Target, or Walmart) reported in ReferenceUSA, or the tract's density of retail stores, obtained from Spatial Insights. None of these measures significantly affect demand.

Specification (4) replaces the number of churches per capita with the more narrowly defined number of fundamentalist churches per capita. We rely on Smith (1990) who provides a classification of Protestant denominations into fundamentalist, moderate, liberal, and other. The listing of churches in Pennsylvania that we obtained from ReferenceUSA then allows us to assign each church to one of these four categories based on denomination information contained in the church name, or based on separate "franchise" information reported by ReferenceUSA. Smith (1990) describes fundamentalists as "a movement of conservative or traditionalist Protestant denominations." Our results do not provide evidence, however, that tracts with a higher presence of fundamentalist churches have statistically significantly lower alcohol consumption, even though the point estimate is negative and about twice the equivalent point estimate on church density in general.

Specifications (5) and (6) test for differences between rural and urban counties – by including the county population density – and rural and urban tracts – by including the Census Bureau's classification of tracts into urbanized and rural. We do not find statistically significant differences between rural and urban areas.

Finally, specification (7) investigates how accounting for variation in the presence of potential consumers near a store at different times of the day affects the demand results. We obtain data from the U.S. Department of Transportation Federal Highway Administration's Census Transportation Planning Products (CTPP) on the daytime population of Pennsylvania tracts. Since we focus on the population of legal drinking age, but the CTPP data does not allow classification by age, we uniformly scale down daytime population to sum to the total Pennsylvania population above the age of 21. We then estimate demand from a given tract as the weighted average of the demand of the evening population and of the daytime population, allowing for a separate demand intercept for the tract's daytime population.

Since our data do not contain sales by time of day, a challenge lies in estimating the weight to be placed on the daytime population. Rather than relying on functional form assumptions to potentially pin down the mix of daytime and evening purchases, we obtain data from the U.S. Department of Labor's American Time Use Survey (ATUS) on the share of grocery store trips that occurs during working hours. Using the 2003 through 2010 waves of the ATUS, we keep all grocery shopping activities that occur on non-holiday weekdays during the PLCB's store opening hours of 10 am to 9 pm. We then compute the share of shopping trips that occur before 5 pm and use this as the weight on the daytime population's demand. According to the ATUS data, 58% of grocery store trips occur during working hours.

As in results in Thomadsen (2005), our estimates indicate that the purchase incidence of the daytime population is lower than the purchase incidence of the evening population. The demand elasticity under this specification is −1.6, but the travel cost implied by our estimates is relatively low at 20 cents per km. We investigate in Appendix C how optimal store configurations change under this demand system.

Table A1: Alternative Demand Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
State-bundle list price	-0.4129	-0.1970	-0.2069	-0.2850	-0.1964	-0.1823	-0.2047
	(0.1708)**	(0.1116)*	(0.0618)***	(0.1261)**	(0.0699)***	(0.0406)***	(0.0851)**
State-bundle list price × ln(Income)	0.0750						
	(0.0367)**						
Driving distance	-0.0603	-0.1743	-0.0624	-0.0515	-0.0473	-0.0374	-0.0548
	(0.0183)***	(0.0991)*	(0.0246)**	(0.0182)***	(0.0174)***	(0.0208)*	(0.0111)***
(Driving distance)×(% w/o car)	-0.0032	-0.0052	-0.0059	-0.0064	-0.0065	-0.0064	
	(0.0024)	(0.0023)**	(0.0020)***	(0.0026)**	(0.0022)***	(0.0021)***	
(Driving distance)×ln(Income)		0.0400					
		(0.0315)					
(Driving dist)×I(Grocery store)			0.0160				
			(0.0153)				
Black	0.3265	0.1735	0.2178	0.2993	0.1734	0.1861	0.0183
	(0.2176)	(0.1861)	(0.2186)	(0.1998)	(0.2362)	(0.1946)	(0.0560)
Median Income	0.0084	0.0308	0.0344	0.0350	0.0336	0.0331	0.0426
	(0.0145)	(0.0057)***	(0.0054)***	(0.0039)***	(0.0052)***	(0.0044)***	(0.0051)***
Median Age	0.0015	-0.0003	-0.0009	0.0000	0.0014	-0.0001	0.0094
	(0.0012)	(0.0009)	(0.0006)	(0.0002)	(0.0002)***	(0.0001)	(0.0073)
No churches per capita	-0.0447	-0.0885	-0.1118		-0.0956	-0.0691	-0.2589
	(0.0899)	(0.1064)	(0.1003)		(0.1094)	(0.1085)	(0.0746)***
No fundamentalist churches per capita				-0.1933			
				(0.1717)			
County population density					0.0022		
					(0.0052)		
Urbanized tract						0.0016	
						(0.0012)	
Daytime							-0.3619
							(0.0062)***
Weight, daytime							0.5800
Implied elasticity of demand	-1.5121	-1.5203	-1.5969	-1.8529	-1.5160	-1.4075	-1.5866
Implied travel cost ($) per unit	0.2397	1.2025	0.6065	0.5345	0.6394	0.6260	0.1990

b. Integer Programming Techniques

One of the solution algorithms that we employ uses integer programming techniques in finding several benchmark configurations to compare to the PLCB's current store configuration. In this appendix, we provide a brief overview over the techniques used. We refer the interested reader to Land and Doig (1960) for the initial development of the branch-and-bound method to solve discrete programming problems and Winston (2003) for a more recent, detailed discussion of alternative solution methods. We begin by restating the firm's problem of choosing the optimal set of locations in which to operate stores under the assumption that the firm's objective is to maximize its profit. The benevolent planner's problem of choosing locations to maximize total surplus can be solved analogously.

Consider a market with R possible locations in which consumers reside. Each location is also available as a possible store location. The firm's problem is to decide whether to operate a store in each location s given that each store has total fixed operating costs of K and generates daily variable profit PS_{rst} from serving those consumers in locations $r = 1,\ldots, R$ who choose to frequent a store in location s at time t. We define, for $s = 1,\ldots, R$,

$$\mathrm{y}_s = \begin{cases} 1 & \text{if the firm opens a store in location s} \\ 0 & \text{otherwise} \end{cases} \tag{A.13}$$

Similarly, we define the $R \times R$ assignment matrix Y of consumer location to store matches where Y_{rs} measures the probability of a consumer in location r visiting a store in location s. Our assumption that consumers visit their closest store imply that

$$Y_{rs} = \begin{cases} 1 & \text{if store location } s \text{ is closest to consumer location } r \\ 0 & \text{otherwise} \end{cases} \tag{A.14}$$

The firm's problem consists of choosing a set of store locations y_s, as well as the associated consumer assignment matrix, to maximize total profit. Note that $\mathrm{y}_s = 1$ implies that $Y_{ss} = 1$ and that $y = \mathrm{diag}(Y)$. The assumption that consumers are assigned to their closest store with probability one transforms the store choice problem into what the Operations Research literature denotes a fixed-charge problem that can be formulated as a binary integer program (BIP). We restate the latter from Equations (9) to (12) above. The firm chooses y to maximize

$$\max_Y \Pi = \sum_{s=1}^{R} \sum_{r=1}^{R} \sum_{t=1}^{T} \frac{1}{T} PS_{rst}\, Y_{rs} - K \sum_{s=1}^{R} Y_{ss} \tag{A.15}$$

subject to

$$\sum_{s=1}^{R} Y_{rs} = 1 \quad \forall r, \tag{A.16}$$

$$Y_{ss} \geq Y_{rs} \quad \forall r, s, r \neq s, \tag{A.17}$$

$$Y_{rs} = \{0,1\} \quad \forall r, s. \tag{A.18}$$

The combinatorial optimization literature has suggested several solution approaches to binary integer programming problems. These include (1) complete enumeration; (2) implicit enumeration; (3) rounding of non-integer, linear programming (LP) solutions to the problem, which may result in a solution far from the true solution to the BIP; and (4) a branch-and-bound method combining elements of the enumeration and LP-relaxation approaches; and (4) implicit enumeration using elements of the branch-and-bound method.

Complete enumeration is impractical in our context due to the large number of possible configurations. Implicit enumeration improves upon this procedure by eliminating obviously infeasible solutions using branching diagrams similar to those used in the branch and bound method discussed below, and then evaluating only the remaining solutions to find the optimal one. The remaining configurations to evaluate may still be numerous. We instead employ the branch-and-bound method. Similar to implicit enumeration, not all, but only some - and potentially very few - of the feasible solutions are enumerated; yet, the branch-and-bound method is guaranteed to find the globally optimal solution to the BIP. It proceeds in the following steps:

1. Solve the LP program resulting from replacing the integer constraints for the solution variables in Equation (A.18) with the less stringent requirement of $Y_{rs} \in [0,1] \ \ \forall r, s$ using the simplex method. This is commonly referred to as LP-relaxation.

   Since the LP-relaxation is a less constrained version of the original BIP, the feasible solution region for the BIP is contained in the feasible solution region for the corresponding LP-relaxation. As a result, the solution to the relaxed linear programming problem provides a value for the firm's profit that is an upper bound U for the optimal solution to the original BIP. This implies that if the optimal LP answer consists of $\{0,1\}$ integers for all Y_{rs}, then it is also the optimal solution to the constrained BIP.

   Notice that the solution of the relaxation allows for fractional allocation of consumers to stores: $Y_{rs} = 1/2$ may be interpreted as allocating half of the consumers in location r to store s.

2. Starting from the solution to the initial LP-relaxation, divide the problem into sub-problems ("branching"). Choose one of the elements of the Y matrix that was assigned a fractional value in the LP solution, $Y_{r's'}$ and subdivide the feasible region of solution values into two sub-problems or nodes, adding, for the chosen Y element, the constraints of $Y_{r's'} = 0$ for sub-problem 1 or $Y_{r's'} = 1$ for sub-problem 2.

   Note also that constraints (A.16) and (A.17) imply that a large number of possible sub-problems are infeasible and can be eliminated ("pruning").

3. Fix the value of $Y_{r's'}$ to the value considered in the sub-problem and find the solution to the resulting LP under the added constraint on $Y_{r's'}$. If the resulting objective function assigned to the sub-problem is worse than an established lower bound L on profit (initially, $L = -\infty$), the entire branch – that is the current sub-problem and any descendants to the sub-problem that could be constructed by adding integer constraints on other partial-value solution elements of Y to the constraint on $Y_{r's'}$ – can be eliminated from further consideration.

4. Partition the sub-problem is again by adding a new Y element to constrain, and investigated. A solution obtained by solving a sub-problem in which all Y elements are integers is a candidate solution. If this candidate solution improves upon the current lower bound to profit, update L.

This process is repeated until no further subdivision is possible, at which point the optimal solution has been reached.

The speed with which a branch-and-bound algorithm finds the solution to a BIP problem depends greatly on finding a close approximation to the solution early, allowing pruning of many sub-problems. This requires a good heuristic for choosing the order of variables to branch on and for selecting the order of nodes to evaluate.

In choosing the sequence of sub-problems to solve, we employ depth-first search with backtracking, where we fully solve one branch of the tree before backtracking to the top of the sub-problem and finding a candidate solution for another branch of the tree. This facilitates re-optimizing the solution to each LP relaxation from the previous one. Further, the branch-and-bound approach requires specification of an order to constructing sub-problems indicating which among the variables Y_{rs} that yield fractional results in the LP-relaxation to branch on first. The Lingo software we employ to solve the store configuration problems selects the order of sub-problems arbitrarily. It further uses various preprocessing steps to detect infeasibilities and possible redundancies among the constraints to improve the lower and upper bounds to the problem.

In problems with large numbers of integer-valued variables and in cases where the LP solution is far from the optimal solution to the BIP, the number of required branching iterations of a branch-and-bound algorithm may be too large for efficient application. For such cases, we employ, as noted in the text, a variant of the "greedy" algorithms discussed in Daskin (1995).

A modeling implication of using linear-programming based techniques is that it is not possible to incorporate a more elaborate store choice model into the optimization process that would recognize the role of other store attributes beyond distance as affecting store choice, such as store size, ease of access, and other unobserved determinants of a store's popularity. LP solution techniques, such as the simplex method we employ, can easily accommodate the fractional assignment matrices for consumer-to-store locations that would result from a probabilistic model of consumer store choice. However, it is not possible in the simplex method to allow the value the solution assigns to one subset of independent variables – in our problem the assignment of consumers to stores Y – to depend on the values of sets of other independent variables to be found as part of the solution. In our problem, incorporating a probabilistic store choice model would imply that the Y matrix depends on the optimal value assigned to every element of the store opening matrix, y.

c. Optimal statewide configurations under alternative demand and fixed costs assumptions

The simulations in Section 6 rest on a number of inputs. Here, we explore the sensitivity of the results to alternative assumptions.

First, our models assume that fixed costs are the same at all current and possible alternative locations. As discussed in Section 2.c, the largest component of store operating costs is labor cost (5/7th of total). Our assumption of constant store operating cost is motivated by the fact that there is no variation across the state in PLCB pay; the PLCB uses a common, state-wide pay scale to compensate its store clerks. The remaining two components to store operating costs are distribution (1/7th of total) and rental expense (1/7th). While there may be economies of scale in distribution from serving stores that are clustered together, we have limited information on the PLCB's distribution system and are unable to examine the role of the store configuration in affecting total distribution costs.

We investigate, however, whether allowing for variation in local rental expenses significantly alters the results in Section 6.b. We assume that the rental expense contribution to store operating costs is proportional to residential median rent from the 2000 Census and predicted the rental expense at every possible location based on a factor of proportionality derived from summing scaled median rent at the actual store locations to the PLCB's total rental expense. We re-computed the optimal profit and welfare maximizing configurations under this alternative fixed cost measure. The county-by-county exact configurations are very similar in size and welfare to the constant fixed cost configurations. The magnitudes of welfare improvements over actual differ by less than 0.1 percentage points across the two fixed-cost specifications, suggesting that the role played by rental expenses is secondary given its small share in total cost.

Second, our analysis here is entirely static; we predict the optimal configuration using current demand. In practice, long-term leases and other sunk closure costs may introduce adjustments costs to the current network that are reflected in some of the apparent locational inefficiencies we detected in Section 6.b. We investigate this by testing whether the PLCB's choice of locations would look more optimal under an earlier demand distribution. We use data from the 1990 Census, together with the demand parameters from specification (1) in Table 3, to predict the optimal county-by-county configurations as of 1990. Regardless of objective, the optimal configurations in 1990 are slightly smaller than their 2000 counterparts, reflecting that real income per-capita has risen over the time period. As in the case of the 2000 configurations, however, the analysis implies significant scope for welfare gains from location adjustments: the optimal 1990 configuration of the same size as the PLCB's store network today entails welfare improvements of 7.2% of revenue relative to the actual configuration, compared to 8.5% when comparing the 2000 configuration to actual. Across Census tracts, the 1990 and 2000 configurations also imply similar deviations in the distance traveled to the closest store from that under the actual configuration. This suggests that sunk closure costs are not a primary explanation for suboptimal location choices.

Third, we investigate the sensitivity of our results to the chosen demand specification. We re-computed welfare under the optimal and actual configurations based on a demand specification that allows for systematic differences in the demand of daytime and evening population (specification (7) in Appendix Table A1) and whose estimates entail an economically low travel cost of only 20 cents per km. While this results in optimal configurations that are between 20 and 35% smaller in size than what we obtain under our main specification, the predicted magnitudes of welfare losses as a share of revenue are similar. As in our current specification, the majority of losses arise from the choice of locations, rather than the size of the network.

d. Regression models investigating political influence

Table A2 presents ordinary least squares models of the number of PLCB stores per house district on characteristics of the house representative and district characteristics. These are discussed in the paper in Section 6.d.

Table A2: OLS Models of the Number of PLCB Stores per House District

	(1)	(2)	(3)	(4)
LiquorControlCom	-0.0072		0.2049	
	(0.0569)		(0.0834)**	
LiquorControlCom × Democrat		-0.6226		0.0427
		(0.0498)**		(0.0552)
LiquorControlCom × (1−Democrat)		0.4644		0.2958
		(0.0659)*		(0.0942)**
Democrat	0.8630	0.9835	0.0878	0.0975
	(0.0496)**	(0.0492)**	(0.0489)	(0.0500)
House District Population (000)	0.0083	0.0110		
	(0.0090)	(0.0046)		
Median Family Income (000)	-0.0059	-0.0071		
	(0.0017)	(0.0014)*		
Percent Black	-2.4946	-2.4962		
	(0.0249)**	(0.0039)**		
Percent Hispanic	-2.9363	-2.8528		
	(0.0252)**	(0.1862)**		
Constant	2.6620	2.4845	1.7528	1.8402
	(0.6154)*	(0.3417)*	(0.1074)**	(0.0739)**
Observations	1,015	1,015	1,015	1,015
R-squared	0.15	0.17	0.99	0.99
District FE	No	No	Yes	Yes

Notes: Dependent variable is the number of PLCB stores in the House district. All regressions include year dummies and are clustered on the legislative session. LiquorControlCom equals one if the district's representative serves on the state's liquor control committee. Democrat equals one if the district's representative is a Democrat.

Standard errors in parentheses.* significant at 5% level; ** significant at 1% level.

Appendix References

Daskin, Mark, Network and Discrete Location: Models, Algorithms and Applications, John Wiley and Sons, Inc., New York. 1995.

Land, Ailsa H. and Alison G. Doig, "An automatic method of solving discrete programming problems." *Econometrica* 28(3) (1960): 497-520.

Smith, Tom W., "Classifying Protestant Denominations." *Review of Religious Research* 31 (1990): 225-245.

Winston, Wayne L., Operations Research: Applications and Algorithms, Duxbury Press, Belmont, California. 2003.