4-2013

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

Katja Seim
University of Pennsylvania

Joel Waldfogel
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/marketing_papers

Part of the [Behavioral Economics Commons](#), [Economic Policy Commons](#), [Marketing Commons](#), [Policy Design, Analysis, and Evaluation Commons](#), [Public Administration Commons](#), [Public Affairs Commons](#), and the [Public Policy Commons](#)

Recommended Citation

Seim, K., & Waldfogel, J. (2013). Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board's Entry Decisions. *American Economic Review, The*, 103 (2), 831-862. <http://dx.doi.org/10.1257/aer.103.2.831>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/marketing_papers/418
For more information, please contact repository@pobox.upenn.edu.

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

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.

Disciplines

Behavioral Economics | Business | Economic Policy | Marketing | Policy Design, Analysis, and Evaluation | Public Administration | Public Affairs | Public Policy

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

By KATJA SEIM AND JOEL WALDFOGEL*

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. (JEL D42, D72, L11, L12, L43, L81)

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 suboptimal locations, as in the well-known Hotelling two-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 2 allocative camps: 32 "private" or "open" states where the

*Seim: Department of Business Economics and Public Policy, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104 (e-mail: kseim@wharton.upenn.edu); Waldfogel: Carlson School of Management, University of Minnesota, 321 19th Avenue South, Minneapolis, MN 55455 (e-mail: jwaldfog@umn.edu). We gratefully acknowledge the comments and suggestions of two anonymous referees, Tim Bresnahan, and participants at various conferences and seminars and in particular we thank Thomas Krantz at the Pennsylvania Liquor Control Board and Mike Ehteshami at the National Alcohol Beverage Control State Association, Inc., for helping us to get access to the data. John Coglianese provided valuable research assistance.

[†] To view additional materials, visit the article page at <http://dx.doi.org/10.1257/aer.103.2.831>.

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 Pennsylvania 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

We explore these questions in six sections. Section I 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 II describes the data used for estimation. Section III 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 III 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 IV 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 V 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 supercompetitive 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 VI concludes with a discussion of the likely effects of the PLCB on aggregate welfare.

¹Our 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 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.

I. 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 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; such states will have fewer stand-alone 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

FIGURE 1. SCALED-UP NUMBER OF LIQUOR STORES VERSUS STATE POPULATION

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 $(0.65) \times (2,355)$, 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 percent markup and an 18 percent 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 et al. 2012), we have no conclusive evidence that retail prices vary systematically between Pennsylvania, other control,

²Calculated as the state's total apparent consumption by type of beverage divided by its population over the age of 21 (LaVallee and Yi 2011).

³The specific pricing rule is $\text{retail price} = (\text{wholesale price} (1.3) + \text{bottle fee})(1.18)$, 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 percent Pennsylvania sales tax.

⁴See American Wine Institute (2011), Distilled Spirits Council (2011), and Commonwealth of Pennsylvania (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 breakdown of sales into wines and spirits and expressing the resulting tax as a percentage of the mean marked-up Pennsylvania wholesale price.

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 III), 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 other states, we can compare the square footage of dedicated liquor stores in a random sample of zip codes in states bordering Pennsylvania (New Jersey, New York, Ohio, and West Virginia)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 percent 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, statewide 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.

⁵The data were obtained from Dun and 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.

⁶The 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/eccec.t04.htm>. We derive total labor costs inclusive of benefits by scaling wage payments by 1.68.

⁷See the PLCB Fiscal Year 09–10 Summary.

⁸Luciew reports in a 2009 article, "Big Ideas: Sell the Pennsylvania Liquor Control Board," 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.

⁹That 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.

¹⁰The 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/eccec.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 Pennsylvania 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, (i) the PLCB faces high store operation costs, and (ii) 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.

II. Data

The basic dataset 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 levels 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).

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 statewide 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 II, the PLCB uses a markup formula to calculate prices. The PLCB is further able to pass on temporary wholesale price reductions to the

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

FIGURE 2. WEEKLY AVERAGE NUMBER OF BOTTLES SOLD PER STORE

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: (i) an adjustment in wholesale prices, or (ii) 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 renegotiates 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 percent in its median week, while a store at the fifth percentile carried 72.0 percent of the products. Similarly, of the 1,000 best-selling products statewide in 2005, the median store carried 82.0 percent in its median week, while a store at the fifth percentile carried 44.2 percent 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 percent of the top 1,000 products. But most stores carry

¹²In our data, 90.26 percent 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.

¹³The 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.

most products, supporting our assumption below that differences in product availability do not drive customers' store choices to a significant 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 covariation 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, inclusion of store fixed effect, and the selection of the sample, we find that demand is relatively 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) reviews the literature on demand for alcohol, most of which use state-level time series data. According to Chaloupka, Grossman, and Safer (2002, p. 23), "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 population-weighted average great-circle distance¹⁵ 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

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

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.

TABLE 1—PRICE ELASTICITY EVIDENCE

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Weekly data (32,509 observations)</i>						
log state bundle list price	-0.963 1.230	-1.090*** 0.176	-1.318 1.356	-1.340*** 0.181		
log state bundle price					-1.893** 0.738	-1.940*** 0.105
Constant	9.888*** 3.095	10.208*** 0.444	10.690*** 3.410	10.747*** 0.454	12.227*** 1.857	12.348*** 0.265
R^2	0.06	0.06	0.07	0.07	0.06	0.07
Quadratic time trend	No	No	Yes	Yes	No	No
Store fixed effect	No	Yes	No	Yes	No	Yes
	(1)	(2)	(3) ^a	(4) ^a		
<i>Panel B. Daily data (191,921 observations unless noted)</i>						
log state bundle price	-0.733 0.387	-0.711*** 0.110	-1.487 0.947	-1.054*** 0.257		
Constant	7.759*** 0.973	6.903*** 0.276	9.768*** 2.379	7.882*** 0.645		
R^2	0.19	0.16	0.07	0.38		
Quadratic time trend	Yes	Yes	No	No		
Holidays	Yes	Yes	No	No		
Store fixed effect	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 week dummies for the weeks 1, 26, 47, 50, 51, and 52 included. We include separate time trends for the period January–October and the holiday period of November–December. Regressions using daily data include day of the week fixed effects. State-bundle prices use a constant bundle for computing the price and vary only by time and not across stores.

^a Sample consists of 23,587 days immediately prior to and following a price change.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

TABLE 2—DEMAND, POPULATION, AND DISTANCE TO THE NEAREST STORE

	Average daily sales per store (1)	Average log daily sales per store (2)	log daily sales per store (3)	log daily sales per store (4)	log daily sales per store (5)	log daily sales per store (6)
Catchment area pop./10,000	173.761*** 12.471	0.449*** 0.030	0.187*** 0.001	0.036*** 0.001	0.187*** 0.001	0.036*** 0.001
Average distance to nearest store	-23.861*** 2.853	-0.081*** 0.007	-0.077*** 4.3E-04	-0.016*** 0.001	-0.077*** 4.3E-04	-0.016*** 0.001
log state bundle list price			-0.846** 0.414	-0.657*** 0.128		
log state bundle price					-0.623* 0.352	-0.696*** 0.109
Constant	273.706*** 25.175	5.365*** 0.061	7.586*** 1.042	6.712*** 0.323	7.025*** 0.393	6.809*** 0.274
Observations	635	635	191,921	191,921	191,921	191,921
R^2	0.27	0.33	0.33	0.21	0.33	0.21
Store fixed effect	No	No	No	Yes	No	Yes

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

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: 12 new stores opened in 2005, while 6 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 percent 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 reestimate 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 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.

¹⁶Note 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. We do not, however, 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.

III. 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 (i) the sensitivity of demand to consumers' distance to stores, and (ii) 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 $C_{st} : \{d_{rst} = \min_{s'} d_{rs'} \mid \forall s' \in S_{open,t} \mid \forall r = 1, \dots, R\}$, 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

$$(1) \quad V_{ijst} = \mathbf{X}'_{jrt} \boldsymbol{\beta}_x - \beta_d d_{rs} - \beta_p p_t + \varepsilon_{ijrst}.$$

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

We aggregate consumers to demographic groups j . In equation (1), $\mathbf{X}_{jrt}$ 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 et al. 1996). Consequently, we differentiate between black (B) and other-race residents (O) and include among the $\mathbf{X}_{jrt}$ the group's per capita income and median age using data from the 2000 census.

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, π , for consumers of each demographic group j in a particular location r :

$$(2) \quad \pi_{jrst} = \frac{\exp(\mathbf{X}'_{jrt}\beta_x - \beta_d d_{rs} - \beta_p p_t)}{1 + \exp(\mathbf{X}'_{jrt}\beta_x - \beta_d d_{rs} - \beta_p p_t)}.$$

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 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:

$$(3) \quad \begin{aligned} \hat{Q}_{rst} &= \sum_{j=\{B,O\}} \pi_{jrst}(\mathbf{X}_{jrt}, d_{rs}, p_t | \beta) M_{jrt} \\ \hat{Q}_{st} &= \sum_{r \in C_{st}} \frac{\hat{Q}_{rst}}{M_{st}} M_{st} = \pi_{st} M_{st}, \end{aligned}$$

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

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

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 percent 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 statewide 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 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

¹⁸For 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.

¹⁹Due to the computational burden of computing driving distances for $3,125 \times 3,124$ 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.

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 kilometers 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 driving time in minutes, respectively. Most of the parameters are stable across specifications. The estimated price coefficients translate into an average price elasticity of -1.2 to -1.5 , 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 percent. 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 tenth and ninetieth percentiles, travel costs per kilometer range from 39 cents (when virtually all households have access to a car) to 157 cents (when 35 percent 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 kilometer 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 (2012) 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.

TABLE 3—DEMAND MODEL ESTIMATES

	(1)	(2)	(3)
State-bundle list price	−0.164*** 0.049	−0.153** 0.077	−0.124 0.098
Driving distance	−0.050*** 0.016		
(Driving distance) × (percent without car)	−0.006** 0.003		
Straight-line distance		−0.061*** 0.019	
(Straight-line distance) × (percent without car)		−0.008*** 0.003	
Driving time			−0.062*** 0.013
(Driving time) × (percent without car)			−0.001 0.001
Black	0.214 0.188	0.209 0.172	0.165 0.137
Median income	0.034*** 0.005	0.033*** 0.004	0.033*** 0.005
Median age	−2.E-04 2.E-04	1.E-04 2.E-04	0.0128 0.0127
Number of churches per capita	−0.105 0.089	−0.154** 0.060	−0.226 0.428
Monday	0.501*** 0.057	0.517*** 0.058	0.593*** 0.067
Tuesday	0.546*** 0.061	0.563*** 0.061	0.627*** 0.073
Wednesday	0.663*** 0.059	0.681*** 0.060	0.748*** 0.073
Thursday	0.819*** 0.058	0.834*** 0.060	0.919*** 0.069
Friday	1.368*** 0.060	1.384*** 0.058	1.455*** 0.067
Saturday	1.393*** 0.056	1.414*** 0.056	1.487*** 0.065
Implied elasticity of demand	−1.267	−1.184	−1.478
Implied travel cost (\$) per unit	0.606	0.841	0.498

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

***Significant at the 1 percent level.

**Significant at the 5 percent level.

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 more elastic demand elasticity of -1.6 .

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 kilometer 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 kilometer 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 kilometer of driving distance. Further, we 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. Online 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

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

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 :

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

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

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

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 $\mathbf{Y}$ 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 $\mathbf{Y}$ matrix also indicates S , the total number of stores operating, as $\text{trace}(\mathbf{Y}) = \sum_{s=1}^R Y_{ss}$. 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

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

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; online Appendix B provides an overview of the branch-and-bound algorithm we employ in finding optimal store configurations.

²⁰We employ LINGO 13.0 to solve these problems.

²¹Our problem is closely related to the facilities location problem analyzed in Perl and Ho (1990). Chan, Padmanabhan, and Seetharaman (2007) employs 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.

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

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

subject to

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

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

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

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 stand-alone 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 inframarginal 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, and each existing store is profitable.²⁴ This deviates from Nash equilibrium because

²²This 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).

²³We 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.

²⁴Our 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

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).

IV. Effect of State Control on Liquor Retailing

We have already seen, in Section II, 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 III, 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 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 II 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

of \$618 eventually cycles among 8 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.

²⁵These 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).

TABLE 4—STATEWIDE COMPARISONS: ACTUAL AND FREE ENTRY CONFIGURATIONS

Configuration	Number of stores	Bottles sold	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
Private Pennsylvania (liquor tax = average tax of private states)					
Free entry (FC = \$618)	1,527	303,192	10,797	699	11,495
Private Pennsylvania (no liquor tax)					
Free entry (FC = \$549)	2,230	322,197	10,921	409	11,330
Free entry targeting size of welfare-maximizing configuration					
Free entry	1,130	289,533	10,706	848	11,554
Welfare-maximizing planner configurations of the size of privatized configurations					
NFE under FC = \$1,110	527	274,507	10,615	1,102	11,717
FE under FC = \$713	906	303,454	10,799	1,041	11,840
NFE under FC = \$549	1,290	322,287	10,920	926	11,846

Notes: Consumer surplus, total producer rents, and total surplus reported in thousands. Producer rents calculated as the sum of variable profit under the given tax structure, labor surplus $((FC - \$549) \times \text{number of stores})$, 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.

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 one-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 percent 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 one-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 percent of expenditure under the PLCB to its value under the two forms of free entry considered here.

V. 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. As discussed in Section III, however, 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 redo this exercise at the state level, before turning to alternative planner objectives and assessing the PLCB relative to these objectives.

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

²⁶Performing 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 I is computationally taxing. We therefore rely on the free entry configuration resulting from the threshold of 145 bottles as an approximation.

²⁷Under 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).

TABLE 5—PERFORMANCE OF MYOPIC ALGORITHM: COMPARISON OF WELFARE-MAXIMIZING CONFIGURATION FOR FIVE COUNTIES

County	Number of stores	Net welfare	Bottles sold	Profit	Net welfare deviation
<i>Panel A. Exact equilibrium configuration</i>					
Berks	22	341,055	8,772	20,055	
Blair	5	112,859	2,167	5,435	
Lancaster	20	406,433	9,130	24,090	
Lycoming	6	102,885	2,072	3,843	
Schuylkill	10	142,324	2,932	3,766	
<i>Panel B. Sequential myopic entry configuration</i>					
Berks	21	340,977	8,668	20,637	−0.02%
Blair	5	112,791	2,161	5,405	−0.06%
Lancaster	21	405,455	9,141	23,033	−0.24%
Lycoming	6	102,529	2,040	3,682	−0.35%
Schuylkill	10	142,324	2,932	3,766	0%

Note: Net 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.

sequential myopic entry algorithm for each county. The results are similar: the maximum welfare under the myopic algorithm is within 0.5 percent 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 welfare- and profit-maximizing configurations from the two approaches differ in size by 0.4 and 4 percent, respectively, 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).

²⁸Our trade-off between consumer surplus and profit is reminiscent of the framework employed in Armstrong, Cowan, and Vickers (1994).

TABLE 6—COMPARISON OF EXACT COUNTY AND SME ALGORITHMS: STATEWIDE ESTIMATES

Configuration	Solution algorithm	Number of stores	Bottles sold	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 thousands. 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.

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, CS rises and PS 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 trade-off between the two.

Calculating the frontier requires an assumption about the store operation cost facing the planner. One interpretation of the PLCB's current supercompetitive 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 $PS + 0.18 \times CS$. 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

FIGURE 3. CONSUMER SURPLUS/PRODUCER SURPLUS FRONTIER, ALTERNATIVE CONFIGURATION SIZES

TABLE 7—STATEWIDE COMPARISONS: ACTUAL AND EFFICIENT CONFIGURATIONS

Configuration	Number of stores	Bottles sold	Consumer surplus	Total producer rents	Total surplus
Actual	621	256,502	10,498	960	11,458
<i>Efficient configurations</i>					
N = locations with at least one 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 thousands. Producer rents equal the sum of variable profit, taxes, and labor surplus $((FC - \$549) \times N)$, less fixed costs. All efficient configurations calculated using exact county algorithm.

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 $PS + 1.17 \times CS$. 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, forgoing 8.7 percent and 7.1 percent 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 percent of expenditure, in daily CS and \$0.13 in daily profit. Consumption declines by 10 percent, 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

²⁹The amounts of CS and profit forgone 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.

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

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 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 percent smaller in size than the ones in Table 7, the majority of welfare losses continue to stem from locational inefficiencies. Online 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; 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 suboptimality 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 (Standard Error = 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 that 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 (Standard Error = 0.20), while the 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. The coefficient is less than a quarter as large, however, 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.

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

TABLE 8—DISTANCE TRAVELED: ACTUAL AND FREE ENTRY VERSUS OPTIMAL CONFIGURATIONS

	$d_t^{PLCB} - d_t^*$ (1)	$d_t^{FE} - d_t^*$ (2)
Rural Share	1.873*** 0.196	8.122*** 0.320
Median income (000)	0.012*** 0.004	-0.045*** 0.006
Constant	-0.630*** 0.177	2.225*** 0.289
Observations	3,123	3,123
R^2	0.03	0.18

Note: Dependent variables defined in text.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

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 Control committee member-represented districts had a store, although this difference is not statistically significant (p -value = 0.37).

We explore this more systematically by regressing the number of liquor stores in a House district on population, median income, 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

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 12 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.

Liquor Control committee member do not have 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 committee 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 party (Republican) impact remains significant (0.3 additional stores). Online Appendix Table A2 contains detailed results. Overall, there is only modest evidence of a political impact on store location decisions.³⁵

VI. 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 percent 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 percent.

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 supercompetitive 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 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

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

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.

REFERENCES

- American Wine Institute.** 2011. "Wine Consumption in the US." <http://www.wineinstitute.org/resources/statistics/article86> (accessed January 13, 2011).
- Armstrong, Mark, Simon Cowan, and John Vickers.** 1994. *Regulatory Reform: Economic Analysis and British Experience*. Cambridge, MA: MIT Press.
- Berry, Steven T.** 1992. "Estimation of a Model of Entry in the Airline Industry." *Econometrica* 60 (4): 889–917.
- Boardman, Anthony E., and Aidan R. Vining.** 1989. "Ownership and Performance in Competitive Environments: A Comparison of the Performance of Private, Mixed, and State-Owned Enterprises." *Journal of Law and Economics* 32 (1): 1–33.
- Bresnahan, Timothy F., and Peter C. Reiss.** 1991. "Entry and Competition in Concentrated Markets." *Journal of Political Economy* 99 (5): 977–1009.
- Chaloupka, Frank J., Michael Grossman, and Henry Saffer.** 2002. "The Effects of Price on Alcohol Consumption and Alcohol-Related Problems." *Alcohol Research and Health* 26 (1): 22–34.
- Chan, Tat Y., V. Padmanabhan, and P. B. Seetharaman.** 2007. "An Econometric Model of Location and Pricing in the Gasoline Market." *Journal of Marketing Research* 44 (4): 622–35.
- Commonwealth of Pennsylvania.** 2011. "Liquor Privatization Analysis: Final Report." http://www.portal.state.pa.us/portal/server.pt/community/liquor_privatization_analysis__final_report/4575 (accessed January 13, 2011).
- Cook, Phillip J., and Michael J. Moore.** 1999. "Alcohol." National Bureau of Economic Research Working Paper 6905.
- Cook, Phillip J., Jan Osterman, and Frank A. Sloan.** 2005. "Are Alcohol Excise Taxes Good for Us?: Short and Long-Term Effects on Mortality Rates." National Bureau of Economic Research Working Paper 11138.
- Daskin, Mark.** 1995. *Network and Discrete Location: Models, Algorithms and Applications*. New York: Wiley-Interscience.
- Davis, Peter.** 2006. "Spatial Competition in Retail Markets: Movie Theaters." *RAND Journal of Economics* 37 (4): 964–82.
- Distilled Spirits Council of America.** 2011. "Distilled Spirits Industry Review Support Tables—2010." http://www.discus.org/assets/1/7/Spirits_Category_Tables_2010.pdf (accessed February 2, 2013).

- Dixit, Avinash, and Joseph E. Stiglitz.** 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67 (3): 297–308.
- Heien, Dale M., and Greg Pompelli.** 1989. "The Demand for Alcoholic Beverages: Economic and Demographic Effects." *Southern Journal of Economics* 55 (3): 759–70.
- Holmes, Thomas J.** 2011. "The Diffusion of Wal-Mart and Economies of Density." *Econometrica* 79 (1): 253–302.
- Hotelling, Harold.** 1929. "Stability in Competition." *The Economic Journal* 39 (153): 41–57.
- Houde, Jean-Francois.** 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review* 102 (5): 2147–82.
- Jia, Panle.** 2008. "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry." *Econometrica* 76 (6): 1263–1316.
- LaVallee, Robin A., and Hsiao-ye Yi.** 2011. *Surveillance Report #92: Apparent Per Capita Alcohol Consumption: National, State, and Regional Trends, 1977–2009*. Bethesda, MD: NIAAA.
- Leung, Siu Fai, and Charles E. Phelps.** 1993. "My Kingdom for a Drink...? A Review of Estimates of the Price Sensitivity of Demand for Alcoholic Beverages." In *Economics and the Prevention of Alcohol-Related Problems: Proceedings of a Workshop on Economic and Socioeconomic Issues in the Prevention of Alcohol-Related Problems*, October 10–11, 1991, Bethesda, MD, edited by Michael E. Hilton and Gregory Bloss, 1–31. Rockville, MD: National Institutes of Health.
- Mankiw, N. Gregory, and Michael D. Whinston.** 1986. "Free Entry and Social Inefficiency." *RAND Journal of Economics* 17 (1): 48–58.
- Mazzeo, Michael.** 2002. "Product Choice and Oligopoly Market Structure." *RAND Journal of Economics* 33 (2): 221–42.
- McManus, Brian.** 2007. "Nonlinear Pricing in an Oligopoly Market: the Case of Specialty Coffee." *RAND Journal of Economics* 38 (2): 512–32.
- Miravete, Eugenio, Katja Seim, Jeff Thurk, and Joel Waldfogel.** 2012. "Complexity, Efficiency, and Fairness in Multiproduct Liquor Pricing." Unpublished.
- Pennsylvania Liquor Control Board.** 2005a. "Daily Units Sold and Revenue by UPC and Store." Obtained from the Liquor Control Board under a Right-to-Know Law request.
- Pennsylvania Liquor Control Board.** 2005b. "Daily Wholesale Price by UPC." Obtained from the Liquor Control Board under a Right-to-Know Law request.
- Pennsylvania Liquor Control Board.** 2009–2010. Summary of Financial Statements FY 2009–10, <http://www.lcb.state.pa.us/PLCB/About/FactsAndFigures/FinancialReports/index.htm> (accessed January 31, 2013).
- Perl, Jossef, and Peng-Kuan Ho.** 1990. "Public Facilities Location under Elastic Demand." *Transportation Science* 24 (2): 117–36.
- Reference USA.** 2009. "Directory of Pennsylvania Religious Organizations (NAICS code 813110)." <http://www.referenceusa.com/> (accessed December 4, 2009).
- Seim, Katja.** 2006. "An Empirical Model of Firm Entry with Endogenous Product-Type Choices." *RAND Journal of Economics* 37 (3): 619–40.
- Seim, Katja, and Joel Waldfogel.** 2013. "Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board's Entry Decisions: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.103.2.831>.
- Small, Kenneth, and Harvey Rosen.** 1981. "Applied Welfare Economics with Discrete Choice Models." *Econometrica* 49 (1): 105–30.
- Smith, Tom W.** 1986. "Classifying Protestant Denominations." General Social Survey Methodological Report 43.
- Spence, Michael.** 1976. "Product Selection, Fixed Costs, and Monopolistic Competition." *Review of Economic Studies* 43 (2): 217–35.
- Thomadsen, Raphael.** 2005. "The Effect of Ownership Structure on Prices in Geographically Differentiated Industries." *RAND Journal of Economics* 36 (4): 908–29.
- Toma, Eugenia.** 1988. "State Liquor Licensing, Implicit Contracting, and Dry/Wet Counties." *Economic Inquiry* 26 (3): 507–24.
- US Census Bureau.** 2001. "United States Census 2000." United States Department of Commerce. <http://www.census.gov/main/www/cen2000.html> (accessed February 1, 2013).
- Wang, Jirong, X. M. Gao, Eric Wailes, Gail Cramer.** 1996. "US Consumer Demand for Alcoholic Beverages: Cross-Section Estimation of Demographic and Economic Effects." *Review of Agricultural Economics* 18 (3): 477–89.
- Young, Douglas J., and Agnieszka Bielinska-Kwapisz.** 2006. "Alcohol Prices, Consumption, and Traffic Fatalities." *Southern Economic Journal* 72 (3): 690–703.