

NBER WORKING PAPER SERIES

ONE MARKUP TO RULE THEM ALL:
TAXATION BY LIQUOR PRICING REGULATION

Eugenio J. Miravete
Katja Seim
Jeff Thurk

Working Paper 24124
<http://www.nber.org/papers/w24124>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2017

We thank Thomas Krantz at the Pennsylvania Liquor Control Board as well as Jerome Janicki and Mike Ehtesham at the National Alcohol Beverage Control Association for providing us with access to the data. We are grateful for many comments and suggestions, and in particular to Jeff Campbell, Kenneth Hendricks, Louis Kaplow, Julie Mortimer, Aviv Nevo, Joel Waldfogel, and participants at the 2017 Cowles Conference in Structural Microeconometrics and the 2017 NBER Summer Institute Industrial Organization workshop. This research was supported in part by the Notre Dame Center for Research Computing. We are solely responsible for any remaining errors. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Eugenio J. Miravete, Katja Seim, and Jeff Thurk. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

One Markup to Rule Them All: Taxation by Liquor Pricing Regulation
Eugenio J. Miravete, Katja Seim, and Jeff Thurk
NBER Working Paper No. 24124
December 2017
JEL No. D42,D63,H23,L43,L66

ABSTRACT

Government often chooses simple rules to regulate industry even when firms and consumers are heterogeneous. We evaluate the implications of this practice in the context of alcohol pricing where the regulator uses a single markup rule that does not vary across products. We estimate an equilibrium model of wholesale pricing and retail demand for horizontally differentiated spirits that allows for heterogeneity in consumer preferences based on observable demographics. We show that the single markup increases market power among upstream firms, particularly small firms whose portfolios are better positioned to take advantage of the policy. For consumers, the single markup acts as a progressive tax by overpricing products favored by the rich. It also decreases aggregate consumer welfare though 16.7% of consumers are better off under the policy. These consumers tend to be older, less wealthy or educated, and minorities. Simple policies therefore generate significant cross-subsidies and may be an effective tool for government to garner favor of key constituencies.

Eugenio J. Miravete
The University of Texas at Austin
Department of Economics
BRB 1.116
1 University Station C3100
Austin, TX 78712-0301
eugenio@eugeniomiravete.com

Jeff Thurk
University of Notre Dame
Department of Economics
Notre Dame, IN 46556.
jthurk@nd.edu

Katja Seim
The Wharton School
University of Pennsylvania
1458 Steinberg-Dietrich Hall
3620 Locust Walk
Philadelphia, PA 19104
and NBER
kseim@wharton.upenn.edu

“The paramount role traditionally assigned by economists to government regulation was to correct the failures of the private market (...), but in fact the premier role of modern regulation is to redistribute income.”

George J. Stigler: Preface to *Chicago Studies in Political Economy*

1 Introduction

Policymaking when agents have heterogeneous preferences over outcomes has been widely studied in the context of theoretical agency models applied to regulatory design (Laffont and Tirole, 1993). Often however, practical policy interventions fail to follow the intricate nonlinear optimal contract design needed to account for moral hazard and adverse selection and, instead, simple policy rules prevail. While technological and administrative limitations may impede the implementation of efficient public policy in practice, Kahn (1970, §7) points out that political objectives or distributional goals also affect policy outcomes.

We study the use of a simple policy rule in the pricing of alcohol in Pennsylvania where the Pennsylvania Liquor Control Board (*PLCB*) operates all retail stores and, by law, applies the same markup formula to all products regardless of differences in consumer demand or upstream market power. Our objective is to identify the winners and losers of the single markup among both upstream distillers and downstream consumers. As the *PLCB*'s single markup policy amounts to a consumption tax, our analysis seeks to broadly address common traits among winners and losers of uniform taxation.¹

Our analysis combines sales data of 312 spirit products sold in 454 local markets across Pennsylvania, covering all retail stores operated by the *PLCB* in 2003 and 2004. We merge our detailed sales data with local market demographics to estimate a random coefficient discrete choice model of demand for horizontally differentiated spirits based on Berry, Levinsohn and Pakes (1995), *BLP* hereafter, where we allow for, but do not impose, market power among upstream firms.² We document the existence of important preference heterogeneity among consumers of different demographic traits as well as large differences in product portfolios operated by the upstream firms. Given our estimates of consumer demand, we estimate the marginal cost of each product assuming a non-cooperative Bertrand-Nash equilibrium among upstream distillers.

¹ Dasgupta and Stiglitz (1971, 4.3) detail the conditions on preferences, demand elasticities, and labor supply that ensure that the optimal tax scheme consists of a single consumption tax or markup rate. The theoretical literature has also established that carefully designed consumption taxes can, similar to income taxation, achieve a given degree of income redistribution (Atkinson and Stiglitz, 1972). Our paper can thus also be viewed as a first empirical analysis of the redistributive effects of uniform consumption taxation.

² The *PLCB* sells both wine and spirits but we focus on the latter as spirits represent the most significant category both in terms of tax revenue and ethanol content. U.S. sales of spirits have also been growing fast at 13.3% between 2007 and 2012, compared to 9% for wine and -4.3% for beer, e.g., see *2012 Beverage Information Group Handbook Advance* available at <http://beverage-handbook-store.myshopify.com/>.

Many features of the regulation of alcohol sales in Pennsylvania help set this study apart. First, the pricing rule is fixed and we observe both wholesale and retail price for each product sold. We thus do not need to assume an objective for the *PLCB* when estimating demand in order to recover the retail marginal costs, as wholesale prices are known to us. Instead, we evaluate the ability of the current policy to attain the state’s dual goals in selling alcoholic beverages – generating tax revenue and managing alcohol consumption to internalize the associated external costs. Second, a product’s price is fixed across the state at a given point in time; upstream firms sell products to the state, which in turn sells them in its state-run stores. This facilitates identification of preferences across markets with widely varying demographic profiles. Third, rules regarding the frequency and advance notice of sales, together with the inability of upstream firms to respond to spatial variation in demand, alleviate common concerns over price endogeneity. Finally, our data include the entire spirit retail industry in Pennsylvania. This provides an opportunity to explore not only how simple regulation impacts welfare among heterogenous consumers but also how it impacts sophisticated multi-product upstream firms offering diverse product portfolios.

We measure the implications of the single pricing rule by comparing the observed equilibrium to alternative Stackelberg equilibria in which the *PLCB* chooses its pricing policy anticipating the profit-maximizing behavior of upstream firms and utility-maximizing responses of downstream consumers. This reflects that imperfect competition in the upstream market enables distillers to impact the equilibrium outcome in the event of a policy change; we allow firms to re-optimize their wholesale pricing decisions following any change to the *PLCB*’s markup policy. In other words, we account for both the *mechanical effect* and *behavioral response* to a change in policy (Saez, 2001) by not only consumers, but also firms. In contrast, the majority of the optimal taxation literature assumes perfect competition among firms.³

Our demand estimates indicate significant preference heterogeneity across products and consumer groups while upstream firms enjoy considerable market power, earning 36 cents in profit for every dollar in revenue. Taken together, these results indicate that a single markup policy that ignores these differences among products and the strategic response of firms could generate large cross-subsidies and have significant consequences for consumers, firms, as well as the state.

We measure these cross subsidies, as well as identify winners and losers of the single markup policy, by comparing the current regulation to equilibrium outcomes generated by alternative policies in which the regulator internalizes demand heterogeneity and upstream market power through product-specific markups. As our demand estimation does not require us to take a stand on the *PLCB*’s objective, we first test whether the observed 30% markup is consistent with maximizing tax revenue for the state’s general fund finding the single markup policy that maximizes tax revenue. We find that the *PLCB* can maximize tax revenue by reducing its uniform tax rate from 30% to 23.7%. Such a reduction generates increases in industry profits – the sum of distiller profit and

³ The computational re-optimization of distillers’ wholesale pricing following a change in markup policy amounts to evaluating the pass-through effect of a tax in a non-competitive model of horizontally differentiated products, an empirical extension of the work of Weyl and Fabinger (2013).

PLCB profit / tax revenue – of \$16.15 million, but 94.4% of this increase is captured by upstream firms, mirroring findings in Miravete, Seim and Thurk (2017). The current 30% markup rule thus captures 99.6% of potential tax revenues attainable with a single markup policy. This result is remarkable because it provides evidence that, conditional on employing a single markup, the current policy leaves little money on the table.

If instead the *PLCB* chose the vector of 312 product-specific markups that maximize tax revenue, aggregate tax revenue would increase 9.8% (\$24.70 million) as the state captures nearly all of the incremental industry profits. Thus, product-level markups provide flexibility for the regulator to use differences in product elasticities to increase competition and lower margins in the upstream market and extract a greater share of industry profits. Interestingly, consumers also benefit from moving to this policy as aggregate consumer welfare improves \$9.47 million, or 1.5% of the average drinking-age resident’s liquor expenditure. While we choose the product-level markups to weakly increase tax revenue ex ante, the welfare effect to consumers is theoretically ambiguous. One contribution of the paper is therefore to showcase an empirical example where a uniform consumption tax not only leaves money on the table, but also decreases consumer welfare.

The single markup also redistributes rents within the upstream firm and downstream consumer segments. For firms, the uniform tax transfers profits from large to small firms and a significant beneficiary of the current policy is a local firm based in Philadelphia, Jacquin. The firm plays a minor role in markets outside of Pennsylvania and we argue that Jacquin’s success under the single markup stems from its ability to lever local market knowledge to sell niche products that are under-priced under the single markup policy. In contrast, large firms such as Diageo and Bacardi operate broad product portfolios. For consumers, the single markup policy amounts to a progressive tax that overprices the spirits favored by higher income consumers who are also less sensitive to changes in price. While most Pennsylvania residents are better off under the equilibrium generated by product-level markups, 16% of residents prefer the single markup policy. These consumers tend to be older, less wealthy, less educated, and minorities. We show that these results are robust to a less sophisticated pricing policy in which the *PLCB* sets markups by spirit type (6 types) and bottle size (3 sizes).

We evaluate the optimality of the single markup by comparing the observed equilibrium to equilibria generated by product-level markups that maximize consumer surplus but are tax revenue neutral as in Ramsey (1927). Absent any externalities, consumers are far better off under the Ramsey solution as aggregate consumer welfare increases \$95.08 million (15.4% of liquor expenditure) due to lower equilibrium retail prices on average and an increase in quantity sold.

Our initial focus on tax revenue generation is motivated by both the significant tax revenue the *PLCB* generates for the state’s general fund in practice and recent state legislation designed increase the *PLCB*’s profitability. The original goal for the *PLCB* at its creation in 1933 was, however, to “discourage the purchase of alcoholic beverages by making it as inconvenient and expensive as possible”, as stated by then Governor Gifford Pinchot. If we assume, therefore, that the *PLCB* uses the tax revenue generated by the 30% single markup exclusively to offset the external

costs of alcohol consumption, we are able to obtain an estimate of the implicit average external cost per unit of ethanol consumed. Given this cost estimate, we show that it is possible to design product-level markups that increase consumption and consumer welfare, but also generate sufficient additional tax revenue to cover the associated larger external costs. Relative to the single markup, this alternative policy would leave older, less wealthy, less educated, and minority consumers, and smaller firms worse off.

In a famous paper, Posner (1971) referred to the use of *uniform rates* as an example of the so-called *taxation-by-regulation* effect. In our case, the taxation-by-regulation effect of a single markup rule is its induced income redistribution through cross-subsidization of spirit prices. Posner’s taxation-by-regulation has long been acknowledged in the theoretical literature, though there exist few empirical studies. To the best of our knowledge, ours is the first to measure taxation-by-regulation using an equilibrium model that not only addresses the best responses of consumers to changes in government policy but also incorporates optimal responses of suppliers.⁴ Moreover, we identify common traits among winners and losers of this regulation and find that current markup regulation generates significant distortions that favor particular constituencies: a small local distiller and well-defined consumer groups.⁵

Whether or not the state legislature explicitly targeted these groups in designing the single markup is difficult to say. Our point is that regardless of the state government’s intent, these groups benefited from the retail price distortions generated by the single markup policy. Indeed, our results do not require us to take a stand regarding the objective of the state legislature in crafting the single markup policy. Instead, we show that by restricting the *PLCB* to a 30% single markup policy, the state legislature leaves money on the table, decreases aggregate welfare, and generates significant redistribution across consumers and firms. We thus identify the implicit effects without addressing whether they are the intended outcome of the regulator or the unintended consequence of a policy that ignored the economic incentives underlying Posner’s taxation-by-regulation argument.

If redistribution was the regulator’s motivation, our results hint that simple policy can be used to garner favor among key constituencies: older, poorer, and less educated consumers in our case. If redistribution was only the unintended consequence of regulation, our results amount to a measurement of the cross-subsidies induced by such policy. In either case our results demonstrate that a uniform tax can generate significant redistribution among consumers and firms, while proponents of uniform taxes often argue that these taxes are inherently fair as they treat everyone equally. Our results indicate just the opposite.

⁴ For instance, Linneman (1980, §5) provides back-of-the-envelope estimates of profit redistribution (in favor of large manufacturers) and uneven price increases (hurting low income families) after the 1973 Mattress Flammability Standard. Finkelstein, Poterba and Rothschild (2009) make use of numerical simulations to evaluate the allocation efficiencies and redistribution effects of not allowing for gender-based pricing (different mortality risk) among participants of the U.K. annuity market. In another related paper, Hausman (1998) does not address redistribution effects and focuses only on the allocation inefficiency of taxing interstate telephone service to subsidize internet access for schools and libraries.

⁵ Seim and Waldfogel (2013, §1.C) also show that the unionized store clerks who operate the state-run network of liquor stores are overpaid (both in wages and benefits) relative to private retailing.

Our work complements other recent papers dealing with the retail regulation of alcoholic beverages. For instance, Seim and Waldfogel (2013), Aguirregabiria, Ershov and Suzuki (2016), and Illanes and Moshary (2015) address the effects of entry restrictions, e.g., no entry in Pennsylvania, partial entry in the wine segment in Ontario, and privatization of the alcohol distribution system in Washington State, respectively. Others focus on the possibility that regulation facilitates collusion among wholesalers: Conlon and Rao (2015) analyzes “post and hold” regulation and Miller and Weinberg (2015) use mergers in the brewing industry to identify collusive behavior. In our analysis we evaluate product-specific markups driven by the elasticity of demand of each product rather than by its alcohol content. Griffith, O’Connell and Smith (2017) do the opposite and introduce differentiated product markups based on corrective taxes that minimize health externalities due to ethanol consumption. Finally, in a related paper (Miravete et al., 2017), we show how the strategic pricing response of firms interacts with the regulator’s ability to use taxation to generate tax revenue, finding that upstream market power flattens the Laffer Curve.

Our paper is also related to classic questions around the value foregone through uniform pricing over sophisticated pricing and price discrimination, though our case focuses on uniform markups across products, rather than uniformly priced, heterogeneous products or uniform prices across locations for a given product. Cho and Rust (2010), Shiller and Waldfogel (2011), and Orbach and Einav (2007) provide evidence on uniform pricing of products of different attributes – car rentals, song downloads, and movies, respectively. Empirical evidence on uniform pricing across heterogeneous geographic markets includes Adams and Williams (2017) for a single product in home improvement supply stores, Chintagunta, Dubé and Singh (2003) for a small set of products in supermarkets in Chicago, and Della Vigna and Gentzkow (2017) for one million products across 35,000 stores in the U.S. The latter work provides convincing evidence that uniform pricing is likely due to the existence of important managerial decision-making costs (Bloom and Reenen, 2007), rather than due to a lack of sufficient heterogeneity in demand across local markets. Our work complements these papers by considering cross-subsidies among heterogeneous consumers and firms due to legislatively mandated uniform markups, an important feature of consumption taxes generally.

The paper is organized as follows. Section 2 presents our data, documents heterogeneous consumption patterns across demographic groups, and discusses the rules currently governing *PLCB*’s pricing and its interaction with the upstream market firms, the distillers. Section 3 presents an equilibrium discrete choice model of demand of horizontally differentiated spirits that incorporates the features of the current pricing regulations while accounting for oligopoly competition in the upstream distillery market. Section 4 reports our estimates and documents significant heterogeneity of preferences for spirits as well as significant market power in the upstream distillery market. In Section 5 we use the estimated model to measure the effects of the single markup on firms, consumers, and state tax revenues. We conclude in Section 6. Additional information on data construction, detailed descriptive statistics, estimation algorithm, robustness, and detailed results are included in the Appendices.

2 The Pennsylvania Market for Spirits

In this section, we briefly summarize the regulation of alcoholic beverages in Pennsylvania that informs our theoretical modeling and econometric identification. We then describe the data we obtained from the *PLCB* and other sources on sales, prices, characteristics of products sold by the *PLCB* and the upstream distillery market. Finally, we document the heterogeneity of consumer preferences for different types of spirits and quantity of alcohol consumed that is behind the differentiated incidence of the current 30% taxation rule.⁶

2.1 The Mechanics of the Pricing Regulation

Pennsylvania adheres to the common three-tier alcoholic beverage distribution system adopted by most states after the repeal of Prohibition: distillers sell their products to wholesale distributors who then sell to retailers, and only retailers may sell to consumers. Unlike most states, the *PLCB* vertically integrates and monopolizes wholesale and retail distribution of wine and spirits.⁷

The Pennsylvania State Legislature provides guidance on the prices the *PLCB* charges both across products and across stores, e.g., Pennsylvania Liquor Code (47 P.S. §1-101 *et seq.*) and the Pennsylvania Code Title 40. Since its inception in 1933, the *PLCB* has regulated alcohol, particularly spirits and wines, using a simple pricing rule that transforms upstream prices into retail prices. This rule has been modified only a limited number of times. From 1937 until 1980 the *PLCB* implemented a uniform markup rule that amounted to a markup (an ad valorem tax) over wholesale cost of 55% for all gins and whiskeys and 60% for other products. In 1980, the Pennsylvania legislature introduced a per-unit handling fee, the *Logistics, Transportation, and Merchandise Factor (LTMF)* and reduced the markup to 25% for all products. In 1993, the markup increased to 30% and the handling fee began varying by bottle size, resulting in a per-unit charge of \$1.05, \$1.20, and \$1.55 for 375 ml, 750 ml, and 1.75 L bottles, respectively.⁸ Consumers also pay an 18% ad valorem tax on all liquor purchases – the so-called “Johnstown Flood Tax” that was instituted as a temporary emergency relief measure in 1936 but has never been repealed.⁹ Summarizing these elements, the uniform pricing rule employed by the *PLCB* during our 2003-2004 sample period is

⁶ Miravete et al. (2017) provide additional detail on the market environment, in particular pertaining to pricing interaction along the supply chain.

⁷ This was the case over the period from 1933 to 2016, when Pennsylvania’s Act 39 allowed for the controlled entry of private retailers into wine sales, although the *PLCB* continues to serve as their supplier. Pennsylvania also has a private system for the sale of beer, allowing the controlled entry of private retailers.

⁸ In 2016 the Pennsylvania state legislature relaxed the requirement that the 30% tax be applied categorically to all products by allowing the *PLCB* to depart from uniform pricing on the top 150 wine and top 150 spirits products. At the time of this writing, the agency has chosen to exercise this option for only 17 of the 300 eligible products, focusing on wines and spirits outside the top 100 and lowering the price below the one prescribed by the uniform pricing rule.

⁹ The *PLCB* collects an additional 6% Pennsylvania sales tax on the posted price to generate the final price paid by the consumer.

$$p^r = [p^w \times 1.30 + LTMF] \times 1.18, \quad (1)$$

where p^r is the retail price of a given product with wholesale price, p^w . Given the simple structure of the pricing regulation and the vertical integration of wholesale and retail segments in Pennsylvania, the pricing rule is simply a combination of ad valorem and unit taxes. Our focus in this paper is the 30% uniform markup; we take at face value that the fixed fees simply represent the storage and transportation costs of bottles of different sizes.

Due to the legislated pricing formula, the wholesale pricing decisions of the *PLCB*'s suppliers, e.g., distillers, are largely responsible for inducing retail price changes; any change in their prices is immediately passed on to retail prices. For the mature products we consider, the *PLCB* places certain restrictions on the number and frequency of wholesale price changes over the course of the year. Temporary wholesale price changes, typically price reductions or sales, amount to 89.7% of price changes in the sample and last for four or five weeks from the last Monday of each month. A product can be put on sale up to four times a year, or once per quarter. A product can thus go on sale for one month, but not two months in a row. Permanent price changes take effect at the beginning of one of the *PLCB*'s four-week long reporting periods for accounting purposes, a slightly different periodicity from the sale periods. See Appendix A for further details. Finally, distillers need to inform the *PLCB* of any proposed temporary price changes at least five months ahead of the desired sale period. Because they need to decide far ahead of time when to run temporary sales, the distillers' ability to respond to unexpected demand swings is limited, facilitating a cleaner identification of demand responses to price changes.

2.2 Data: Quantities Sold, Prices, and Characteristics of Spirits

Our data, obtained under the Pennsylvania Right-to-Know Law, contain daily information on quantities sold and gross receipts at the UPC and store level for all spirits carried by the *PLCB* during 2003 and 2004. We aggregate the daily data to the monthly sales periods, resulting in 22 periods for our sample.¹⁰ The *PLCB* also provided the wholesale cost of each product, which is constant across stores, but varies over time reflecting the temporary or permanent price changes discussed above.

Each store carries a vast variety of products among which we focus on popular 375 ml, 750 ml, and 1.75 L spirits products, representing 64.1% of total spirit sales measured in bottles and 70.1% of total spirit sales by revenue.¹¹ The resulting sample contains 312 products across the two-year sample. We see little variation in the product set across stores indicating that consumers have equal access to the products we consider in the sample.

¹⁰We drop periods that overlap with January due to a significant reduction in temporary price changes.

¹¹Many products are available to consumers but are seldom purchased. The 375 ml, 750 ml, and 1.75 L bottle sizes account for 80.9% of total spirit sales by volume and 91.6% of total spirit sales by revenue. Within these bottle sizes we further focus on popular products which we define as those products that account for 80% of bottle sales in each spirit type-bottle size pair. We also dropped tequilas as there were few products. In total, these restrictions allow us to drop a total of 1,313 products from our sample.

We classify products into six categories: brandy, cordials, gin, rum, vodka, and whiskey. For each product, the *PLCB* provided its alcohol content measured as proof (100 proof corresponds to 50% alcohol content by volume) and whether or not the product is imported or contains flavorings. Table 1 summarizes these product characteristics. Vodkas and whiskeys have significantly larger market shares (32.1% and 24%, respectively), than rum (16.3%), cordials (13.6%), or brandy (7.3%), even though cordials in particular is one of the top categories in terms of number of products. Flavored spirits, which represent 16.3% of products, are primarily cordials and brandies and, to a lesser extent, rums and vodkas. Most whiskeys and cordials are imported while other spirits are predominantly domestic. There is significant variation in proof across product categories; while the average across all products is 75.33, it ranges from 55.82 for cordials to 83.42 for gins.¹²

Table 1: Product Characteristics by Spirit Type

	Products	Price	Share	% Flavored	% Imported	Proof
By Spirit Type:						
BRANDY	26	13.90	7.26	30.77	26.92	76.15
CORDIALS	62	15.10	13.59	32.26	51.61	55.82
GIN	28	15.59	6.72	3.57	28.57	83.42
RUM	40	14.32	16.31	10.00	17.50	74.03
VODKA	66	13.76	32.10	21.21	40.91	81.60
WHISKEY	90	16.74	24.03	0.00	58.89	80.98
By Price:						
CHEAP	162	10.50	47.00	17.90	22.84	72.46
EXPENSIVE	150	19.91	53.00	12.00	64.67	77.82
By Bottle Size:						
375 ml	48	7.15	15.21	8.33	47.92	75.10
750 ml	170	14.49	50.29	21.76	44.71	72.95
1.75 L	94	18.83	34.50	6.38	37.23	78.77
ALL PRODUCTS	312	14.87	100.00	16.30	37.40	75.33

Notes: “Price” is the average price weighted by quantity sold (bottles). “Share” is based on number of bottles sold. “Cheap” (“Expensive”) products are those products whose mean price is below (above) the mean price of all spirits of the same spirit type and bottle size.

We characterize spirits as expensive when their simple averaged price exceeds the mean price of all spirits of the same type and bottle size. Table 1 shows that expensive spirits are purchased nearly as often as cheaper ones, but are less likely to be flavored or domestically produced and have higher proof. The 750 ml bottle is the most popular size in terms of unit sales and product variety, accounting for 50.3% of unit sales and 54.5% of available spirits products, followed by the 1.75 L bottle with a share of 34.5% of unit sales and 30.1% of products. The smallest bottles we consider, those in the 375 ml format, account for the remaining 15.2% of units sold and 15.4% of products.¹³

¹²In addition to these characteristics, the empirical analysis also includes a measure of product quality obtained from Proof66.com, a spirits ratings aggregator. We do not include it in Table 1 because the product score is primarily informative within, but not across, spirit categories.

¹³Table 1 and Table B.1 in Appendix B show that average bottle prices increase with bottle size, both in total and within spirit category although it decreases by per ounce. This explains why the revenue share of 1.75 L bottles, 43.7%, exceeds the share of bottles sold, 34.5%.

Some of this variation is driven by the product set as upstream firms manufacture brands in particular bottle sizes. For instance, our final sample is composed of 198 brands (e.g., Captain Morgan) but 88 of these brands were available only in the 750 ml bottle size while 1 and 31 brands were offered only in the 375 ml and 1.75 L size, respectively. The remaining 78 brands were offered in several bottle sizes (e.g., Diageo sold Captain Morgan in 375 ml, 750 ml, and 1.75 L sizes).¹⁴

Table 2: Frequency of Sales

	Holiday	Year	Times
By Spirit Type:			
BRANDY	34.62	61.54	2.47
CORDIALS	41.94	64.52	2.70
GIN	39.29	71.43	2.62
RUM	40.00	72.50	2.46
VODKA	48.48	78.79	2.44
WHISKEY	47.78	77.78	2.74
By Price:			
CHEAP	35.19	62.96	1.36
EXPENSIVE	53.33	83.33	2.47
By Bottle Size:			
375 ml	4.17	50.00	3.08
750 ml	49.41	78.24	2.53
1.75 L	54.26	74.47	2.68
ALL PRODUCTS	43.91	72.76	2.60

Notes: The “Holiday” season is defined as the two pricing periods that encompass Thanksgiving through the end of the year. Statistics reflect the percent of products with a temporary price reduction during the corresponding period except for “Times,” which denotes the average number of times that spirits in each category are on sale during a year.

While Table 1 summarized mean price differences between products of different spirit types, Table 2 summarizes price changes over time. As discussed above, these largely originate from distillers temporarily adjusting their products’ wholesale price resulting in monthly retail sales. Distillers temporarily change a product’s price 2.6 times per year on average, and most products, 73%, go on sale at least once a year. Distillers place vodkas, whiskeys, cheaper products, and products in larger bottle sizes (750ml or 1.75L) on sale more frequently than average. A significant share of all spirits products – 44% – go on sale at some point during the holidays, which we define as pricing periods that overlap with Thanksgiving through the end of the year. While over 63% of cheap products go on sale at some point in the year, the average product is placed on sale only 1.36 times, far less than other product categories. In contrast, 375 ml products are also less frequently subject to any temporary price reductions, on average, but those 375 ml products that do go on sale at least once are placed on sale three times per year on average. The monthly sales activity, together with variation in the amount of price reduction, are our primary source of price variation to identify consumers’ price responsiveness.

¹⁴This pattern is reflected in the raw data as well where 958 of the potential 1,192 brands were offered only in one bottle size, usually the 750 ml format.

2.3 The Upstream Distillers

Our objective in this paper is to assess the differential effect of the uniform markup policy on equilibrium behavior of different groups of agents to uncover Posner’s taxation-by-regulation. As the pricing formula in equation (1) highlights, how any alternative markup policy affects retail prices and consumer behavior depends on how wholesale prices respond to the markup policy. Weyl and Fabinger (2013, §4) show that the pass-through of a tax depends on the market conduct and responses of the taxed upstream firms. In our setting, the upstream firms’ markups depend on the characteristics and positioning of their product portfolios. This section provides a brief summary of the features of the upstream distillery market.

During our sample period, the upstream distillery market was composed of 34 firms with broad arrays of product portfolios summarized in Tables 3 and 4. The market leader, Diageo, accounted for 22% of total unit sales and 25% of revenue, while a large set of small fringe producers (29) accounted for 42% and 46% of total quantity sold and revenue, respectively. Nineteen of these firms operated product portfolios of less than five products and seven were single product firms.

Table 3: The Upstream Market

Firm	Products	Share of Spirit Market	
		By Quantity	By Revenue
Diageo	63	22.10	25.02
Bacardi	22	8.89	9.72
Beam	32	9.69	8.84
Jacquin	27	9.52	5.97
Sazerac	18	7.95	4.69
Other Firms (29)	150	41.85	45.76

Notes: Sample of distilleries with national brands. Firms ordered according to market share based on revenue.

In Table 4 we show that large firms such as Diageo and Bacardi operated extensive product portfolios although there was also substantial heterogeneity across firms in their product offerings. For example, Diageo had a relatively balanced portfolio where rums, vodkas, and whiskeys generated approximately 21%, 31%, and 25% of its total revenue, respectively. In contrast, Bacardi operated a much more concentrated portfolio as 71% of its revenue came from its rum products compared to 19% from whiskey products. Among the larger competitors with unit shares of the Pennsylvania market of at least 5%, only the Pennsylvania-based firm Jacquin sells brandies, where it faces only seven smaller competitors and generates 22% of its revenue. While the company offers products in all bottle sizes and spirit types, its portfolio focuses exclusively on cheap products.

Table 4 documents that a significant number of competitors are present in all product categories. Nevertheless, the upstream distiller market is best characterized as oligopolistic, with Herfindahl-Hirshman indices ranging from 1,023 for cordials to 3,087 for rums (951 for spirits in total). We therefore assume that distillers set wholesale prices non-cooperatively as in a Bertrand-

Table 4: Upstream Product Portfolios

	DIAGEO	BACARDI	BEAM	JACQUIN	SAZERAC	FIRMS
By Spirit Type:						
BRANDY	0.00	0.00	0.00	21.57	0.00	8
CORDIALS	11.22	1.98	16.22	9.77	13.29	18
GIN	11.67	8.26	5.14	1.77	3.62	10
RUM	21.27	70.89	4.06	22.12	0.00	10
VODKA	30.99	0.00	10.4	43.26	63.9	14
WHISKEY	24.85	18.88	64.18	1.52	19.18	20
By Price:						
CHEAP	19.60	34.20	52.61	100.00	85.30	25
EXPENSIVE	80.40	65.80	47.39	0.00	14.70	25
By Bottle Size:						
375 ml	5.82	5.90	2.93	10.24	15.07	18
750 ml	54.67	50.59	38.13	27.05	20.53	31
1.75 L	39.51	43.51	58.95	62.71	64.59	25
ALL PRODUCTS	100.00	100.00	100.00	100.00	100.00	34

Notes: Table presents share of total revenues for each firm by spirit type, price, and bottle size. “Firms” is the total number of firms with at least one product in the given category.

Nash oligopoly with horizontally differentiated products, which allows us to recover the distillers’ marginal costs and predict their optimal behavior under alternative *PLCB* pricing strategies.

2.4 Heterogeneity of Preferences for Quantity and Varieties of Spirits

Pennsylvania is a demographically diverse state, which allows us to trace consumer preferences across a wide range of demographic profiles. We geocode the 624 stores’ street addresses to link their geographic location to data on population and demographic characteristics for nearby consumers based on the 2000 Census. We combine Pennsylvania’s Census block groups into markets, assigning each to the operating store that is closest to them in any time period. We further consolidate stores in the same ZIP code resulting in 454 total local markets.¹⁵ Table 5 summarizes these demographic characteristics. About 39% of households in Pennsylvania make more than \$50,000 a year, but the income distribution differs significantly across markets, with rich households comprising anywhere between 10% to 76% of the population across markets. Similarly, the share of minority households in a market ranges from virtually zero to 99%, with minorities comprising 13% of citizens in the average market. We see similar diversity in educational attainment with 44% of residents in the average market reporting at least some college education but this varies from 13% to 87% across the state. Finally, there is less disparity in age where the average age in the average market is 40 while the youngest market has an average age of 31 and the oldest market has an age of 43.

¹⁵ Because stores open and close during the year, the characteristics of stores’ ambient consumers also change over time, which helps identify the effect of demographic interactions. See Appendix A for further detail on store openings and closings, ambient demographics, and widespread availability of most spirits in every store.


Table 5: Demographic Attributes of Pennsylvania Markets

	Percentage of Population				
	POPULATION	AGE	EDUCATION	INCOME	MINORITY
Mean	26,241	40.07	44.06	38.81	13.07
Std. Dev.	16,386	1.45	14.41	14.33	18.83
Min	1,469	31.34	12.56	9.90	0.44
Max	112,065	43.06	87.01	75.92	98.95

Source: 2000 Census of Population. AGE reports the average age in the market. EDUCATION reports the share of population with some college education. MINORITY defined as share of non-white population. INCOME is the share of households with income greater than \$50,000. Figure B.1 in Appendix B we present the spatial distribution of demographics.

Do these wide differences in the demographics of local markets translate into heterogeneity in terms of ethanol consumed and tax revenue generated, both measures linked to the *PLCB*'s objectives? We answer this question by dividing the store markets into quintiles based on each of the four demographic traits – the share of households with incomes above \$50,000, the share of non-white or minority households, the share of residents with some college education, and the average age. As markets differ in their consumption of ethanol, we also split the markets according to per capita ethanol consumption where a market in the top (bottom) quintile is designated as a high (low) “Risk” market.

Figure 1: Demographics and the PLCB’s Objectives



Notes: Graphs compare shares of total ethanol consumption and tax revenue garnered by the bottom and top 20% of markets according to the demographic trait considered. We denote the difference between the retail price and the wholesale cost as the single amount of tax revenue per bottle, not differentiating between the ad valorem tax (markup), the handling fee, and the liquor tax. Total tax revenue is then tax revenue per bottle times number of bottles purchased. Since tax revenue per bottle is near proportional to the final retail price of equation (1), variation in total tax revenue across markets also maps into differences in consumer spending.

In Figure 1 we present the shares of total tax revenue and total ethanol consumption for markets in the top and bottom quintiles for each trait, relative to a 20% benchmark (the dashed line) corresponding to the case where market attributes do not correlate with consumption and expenditures. We observe high ethanol consumption and spending in markets with high

concentrations of wealthy, well-educated, young, and non-white consumers.¹⁶ Interestingly, the difference in the consumption shares of the top and bottom markets in the distribution is always greater than the difference in the share of taxes paid. High risk markets (i.e., the top 20% of markets by per capita consumption) consume 33% of the total ethanol in the state and generate 29% of total tax revenue. In contrast, ethanol consumption in low risk markets accounts for only 5% of total ethanol sales and generates 11% of total tax revenue.

Table 6: Connecting Consumer Preferences and Demographics

	RISK		AGE		MINORITY		EDUCATION		INCOME	
	Low	High	Low	High	Low	High	Low	High	Low	High
By Spirit Type:										
BRANDY	6.6	7.3	9.8	4.6	5.2	12.3	11.4	5.4	11.9	4.2
CORDIALS	14.7	12.9	13.2	13.4	15.3	12.4	14.7	12.2	13.0	12.9
GIN	5.3	8.0	8.1	6.4	4.7	9.1	7.8	7.6	8.4	6.9
RUM	18.4	15.4	17.5	14.2	16.6	16.8	18.2	14.2	17.5	13.7
VODKA	27.7	34.5	32.5	34.1	27.4	31.8	26.9	36.9	29.6	37.3
WHISKEY	27.3	21.9	18.8	27.4	30.9	17.6	21.0	23.6	19.7	25.0
By Price:										
CHEAP	56.1	50.7	55.0	49.1	54.8	56.6	57.5	47.7	59.0	46.6
EXPENSIVE	43.9	49.3	45.0	50.9	45.2	43.4	42.5	52.3	41.0	53.4
By Bottle Size:										
375 ml	15.1	15.7	20.1	10.6	10.5	22.8	19.3	13.6	22.5	11.6
750 ml	49.9	50.6	51.1	49.1	49.3	50.8	51.0	50.6	49.9	49.7
1.75 L	35.0	33.7	28.8	40.3	40.2	26.4	29.6	35.7	27.5	38.6

Notes: Statistics represent market shares (based on bottles sold) by product characteristic. Demographic categories are defined in Table 5 or, in the case of “Risk,” in the text.

In Table 6 we complete our analysis by documenting the heterogeneity of preferences for different types of spirits across consumers with different demographic characteristics and alcohol consumption habits. We do so by calculating the market shares (by unit sales) in each demographic cohort. For example, from Table 6 we see that 38.6% of bottles purchased in the wealthiest 20% of markets were 1.75 L bottles. We can then infer differences in preferences by reading across columns (i.e., demographics) for a given product characteristic (e.g., 1.75 L bottles). Recall that retail prices at a point in time are identical across the state and stores carry similar product sets so differences in purchase patterns across demographics reveals differences in preferences (see Appendix A).

The data indicate that minorities strongly favor brandy, gin, and 375 ml products, but not whiskey or 1.75 L products. In markets with high income and a highly educated population, vodka is far more popular than rum and brandy while consumers also tend to buy more expensive spirits. Markets dominated by young, less educated, and lower income populations show a clear preference for cheap products. The popular 750 ml bottle amounts to almost exactly half of all bottle sales across demographic traits, but between the 375 ml and 1.75 L sizes, higher-income markets clearly

¹⁶The finding that per capita ethanol consumption and income are positively correlated has been demonstrated in earlier work though the exact mechanism is unclear, e.g., see Cerda, Johnson-Lawrence and Galea, 2011. Researchers often cite social norms or even social networking as causes.

favor 1.75 L products. Finally, heavy drinkers, i.e., consumers in high “risk” markets, prefer expensive and vodka products but are unlikely to purchase 375 ml bottles; preferences that reflect the positive correlation between per capita ethanol consumption and income.

3 Model

In this section we describe a static model of oligopoly price competition with differentiated goods. We envision a two-stage Stackelberg game where the regulator first commits to a markup (or tax) policy. Upstream distillers observe the policy and then simultaneously choose wholesale prices p^w to maximize profits each period. The chosen wholesale prices translate into specific retail prices based on the regulator’s chosen markup policy. Finally, consumers in each market choose the product that maximizes their utility given the retail prices and characteristics of all products.

The purpose of the model is twofold. First, the model lays out how consumers respond to the chosen retail prices. We rely on a flexible utility specification to estimate consumer preferences from observed consumer choices over time and across markets. The estimated preferences facilitate predicting consumer responses to counterfactual pricing policies. Second, the model highlights that the retail prices the consumer is exposed to are driven to a significant degree by the strategic wholesale price choices in the upstream market. If the *PLCB* were to alter its pricing rule, upstream distillers would choose different wholesale prices leading to different retail prices and consumer purchase decisions. We thus need to incorporate upstream behavior in the model in order to appropriately conduct our counterfactual analysis. We begin with a description of the consumer’s choice problem and then continue with the distillers’ pricing problem.

3.1 A Discrete Choice Model of Demand for Spirits

We follow the large literature on discrete-choice demand system estimation using aggregate market share data (Berry, 1994; Berry et al., 1995; Nevo, 2001) to model demand for spirits as a function of product characteristics and prices. By mapping the distribution of consumer demographics into preferences, the model enables us to estimate realistic substitution patterns between products while accounting for the heterogeneity in preferences exhibited in Table 6. We assume that consumer i in market l in period t obtains the following indirect utility from consuming a bottle of spirit $j \in J_{lt}$ ¹⁷

$$u_{ijlt} = x_j \beta_i^* + \alpha_i^* p_{jt}^r + H_t \gamma + \xi_{jlt} + \epsilon_{ijlt}, \tag{2}$$

where $i = 1, \dots, M_{lt}; \quad j = 1, \dots, J_{lt}; \quad l = 1, \dots, L; \quad t = 1, \dots, T.$

The $n \times 1$ vector of observed product characteristics x_j is identical in all markets l where product j is available and fixed over time, though the availability of different spirits changes over

¹⁷In the absence of individual purchase information we opt to treat bottles of different sizes of the same spirit as different products with identical observable characteristics other than size and focus on horizontal differentiation between products.

time due to product introductions or removals. The $T \times 2$ matrix $H_t = [q3_t \ m12_t]$ includes a summer dummy for periods in July, August, and September and a holiday dummy for periods t coinciding with the holiday season from Thanksgiving to the end of the year. The price of product j at time t is denoted p_{jt}^r and is the same across all markets l . We further allow utility to vary across products, markets, and time via the time and location-specific product valuations ξ_{jlt} , which are common knowledge to consumers, upstream firms, and the *PLCB* but unobserved by the econometrician. Lastly, ϵ_{ijlt} denotes the idiosyncratic unobserved preferences by consumer i for product j in market l and period t , which we assume to be distributed Type-I extreme value across all available products J_{lt} .

We characterize consumer i in market l by a d -vector of observed demographic attributes, D_{il} including age, education, income, and race. To allow for individual heterogeneity in purchase behavior, we model the distribution of consumer preferences over characteristics and prices as multivariate normal with a mean ν_{il} that shifts with these consumer attributes

$$\begin{pmatrix} \alpha_i^* \\ \beta_i^* \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_{il} + \Sigma \nu_{il}, \quad \nu_{il} \sim N(0, I_{n+1}), \quad (3)$$

where Π is a $(n + 1) \times d$ matrix of coefficients that measures the effect of observable individual attributes on the consumer valuation for spirit characteristics including price, while Σ measures the covariance in unobserved preferences across characteristics. We restrict $\Sigma_{jk} = 0 \ \forall k \neq j$, and estimate only the variance in unobserved preferences for characteristics.¹⁸ Similarly, the contribution of demographic and product characteristic interactions (Π) allows cross-price elasticities to vary differentially in markets with observed differences in demographics. Exogenous store openings/closings alter the demographic composition of demand for surrounding stores and help us identify these demographic interactions. We make the common assumption that during a particular time period, each consumer selects either one bottle of the J_{lt} spirits available in her market, or opts to purchase beer or wine for off-premise consumption, i.e., not consumed in a restaurant or bar.¹⁹ We denote this outside option, denominated in 750 ml bottle-equivalents, by $j = 0$ with zero mean utility.

The annual potential market for location l is the number of drinking-age residents scaled by per-capita off-premise consumption. Hence, a consumer opting for the outside option is choosing to

¹⁸Introducing unobserved preferences for a given characteristic j allows cross-price elasticities among products with similar characteristics (e.g., proof) to be higher than in a simple Logit model, thereby relaxing the restrictive substitution patterns generated by the Independence of Irrelevant Alternatives (IIA) property of the multinomial logit model.

¹⁹Nevo (2000, p.401) discusses limitations of the present discrete choice approach when instead, individuals purchase several products or multiple bottles of the same product at the same time. If such consumer behavior were important, Hendel (1999) and Hendel and Nevo (2006) show that assuming single-unit purchases could understate price elasticities in the case of assortment decisions, but overstate own-price elasticities in the case of stockpiling. In Miravete et al. (2017) we test for stockpiling using a similar dataset and find no evidence. Similarly, Seim and Waldfogel (2013) present suggestive evidence that the *PLCB*'s demand does not respond to price declines more strongly than average in areas with higher travel costs to the store where customers have a higher incentive to buy larger quantities or assortments.

consume a 750 ml “bottle” of beer (i.e., roughly two bottles of beer) or a 750 ml bottle of wine.²⁰ The average drinking-age Pennsylvanian consumed 96.2 liters of alcoholic beverages through off-premise channels in 2003, or 128.2 750 ml bottle equivalents according to Haughwout, Lavalley and Castle (2015). The 2003 potential market for location l is then the number of drinking-age residents scaled by 128.2, which we allocate evenly across pricing periods according to the periods’ lengths to give us the potential market per period, M_{lt} . To put this figure in perspective, beer accounts for approximately 90% of total consumption by volume so the average drinking-age Pennsylvanian consumed slightly less than five 375 ml bottles of beer per week, but only approximately thirteen 750 ml bottles of both wine and spirits annually. We follow a similar approach in constructing the potential market for 2004.

Consumer utility-maximization connects the set of individual-specific attributes and the set of product characteristics as follows

$$A_{jt}(x., p_{.t}^r, \xi_t; \theta) = \{(D_{il}, \nu_{il}, \epsilon_{i.lt}) | u_{ijlt} \geq u_{iklt} \quad \forall k = 0, 1, \dots, J_{lt}\}, \quad (4)$$

where we summarize all model parameters by θ . We follow the literature in decomposing the deterministic portion of the consumer’s indirect utility into a common part shared across consumers, δ_{jlt} , and an idiosyncratic component, μ_{ijlt} , given by

$$\delta_{jlt} = x_j \beta + \alpha p_{jt}^r + H_t \gamma + \xi_{jlt}, \quad (5a)$$

$$\mu_{ijlt} = \begin{pmatrix} x_j & p_{jt}^r \end{pmatrix} (\Pi D_{il} + \Sigma \nu_{il}). \quad (5b)$$

In estimating the model, we integrate over the distribution of $\epsilon_{i.lt}$ analytically. The probability that consumer i purchases product j in market l in period t is then

$$s_{ijlt} = \frac{\exp(\delta_{jlt} + \mu_{ijlt})}{1 + \sum_{k \in J_{lt}} \exp(\delta_{klt} + \mu_{iklt})}. \quad (6)$$

Deriving product j ’s aggregate market share in each location requires integrating over the distributions of observable and unobservable consumer attributes D_{il} and ν_{il} , denoted by $P_D(D_i)$ and $P_\nu(\nu_i)$, respectively. The market share for product j in market l at time t is:

$$s_{jlt} = \int_{\nu_l} \int_{D_l} s_{ijlt} dP_D(D_i) dP_\nu(\nu_i), \quad (7)$$

which we evaluate using simulating techniques. See Appendix C for detail.

²⁰Note that this definition accounts for the total volume of alcoholic beverages but not for the different ethanol contents of beer (4.5% on average), wine (12.9% on average), and spirits (37.7% on average in our sample).

3.1.1 Consumer Welfare

An advantage of a structural model is that it enables the researcher to assess equilibrium changes in welfare. We identify beneficiaries of the single markup policy by evaluating changes in consumer welfare via compensating variation, i.e., the amount of income necessary to keep individuals in a given market indifferent between the counterfactual set of prices p' and the current ones p under the single markup rule,

$$CV_{ilt}(p, p') = \frac{1}{\alpha_i^*} \ln \left[\frac{\sum_{j \in J_{lt}} \exp[V_{ijlt}(p')]}{\sum_{j \in J_{lt}} \exp[V_{ijlt}(p)]} \right], \quad (8)$$

where $V_{ijlt}(\cdot)$ is given by (2). The mean compensating variation for agents living in location l is

$$CV_l(p, p') = \sum_t M_{lt} \int_{\nu_l} \int_{D_l} CV_{ilt}(p, p') dP_D(D_i) dP_\nu(\nu_i). \quad (9)$$

As defined, residents in location l are on average better off under the current policy when $CV_l(p) > 0$, i.e., when the mean compensating variation is greater than zero, consumers at this location would be willing to pay the government to not adopt the policy that generates retail prices p' .²¹

3.2 An Oligopoly Model for Distillers

Given optimal consumer choices, we now consider competition between upstream distillers. A total of F firms compete in the upstream market where each firm $f \in F$ produces a subset J_t^f of the $j = 1, \dots, J_t$ products. In each period t the upstream firms simultaneously choose the vector of wholesale prices $\{p_{jt}^w\}_{j \in J_t^f}$ to maximize period t profit

$$\max_{p_{jt}^w} \sum_{j \in J^f} \left[(p_{jt}^w - c_{jt}) \times \sum_{l=1}^L M_{lt} s_{jlt}(p^r(p^w), x, \xi; \theta) \right], \quad (10)$$

where c_{jt} denotes the marginal cost of product j in period t . To simplify the notation of this static problem, we omit the period t subscripts going forward.²² Define as $s_j(p^r, x, \xi; \theta)$ the state-wide demand for product j , $\sum_{l=1}^L M_{lt} s_{jlt}(p^r, x, \xi; \theta)$. Profit maximization in the upstream market implies that distiller f chooses prices $p_j^w \forall j \in J^f$ to solve the following set of first-order conditions,

$$s_j(p^r(p^w), x, \xi; \theta) + \sum_{m \in J^f} (p_m^w - c_m) \times s_m(p^r(p^w), x, \xi; \theta) \times \frac{\partial s_m}{\partial p_j^w} = 0. \quad (11)$$

²¹ Implicitly we are assuming that a change in *PLCB* spirit pricing policy has no effect on the utility of consumption for beer and wine (the outside good).

²² We ignore any dynamic considerations to distillers' pricing decisions based on regulatory restrictions on their pricing. While the *PLCB* limits the number of times distillers can temporarily change a product's price to four, this regulation does not constrain upstream pricing for the majority of products. In the data, 73.5% of products go on sale three times or less.

The final term $\frac{\partial s_m}{\partial p_j^w}$ is the response in quantity sold for product m to a change in the wholesale price and, through the pricing rule, the retail price of product j . Assuming a pure-strategy Bertrand–Nash equilibrium in wholesale prices, the vector of profit-maximizing wholesale prices solves

$$p^w = c + \underbrace{[O^w * \Delta^w]^{-1} \times s(p^r(p^w), x, \xi; \theta)}_{\text{vector of \$ markups}}, \quad (12)$$

where O_t^w denotes the ownership matrix for the upstream firms with element (j, m) equal to one if goods j and m are in J^f and upstream firm f chooses these prices jointly. We define Δ^w as a matrix that captures changes in demand due to changes in wholesale price,

$$\Delta^w = \begin{bmatrix} \frac{\partial s_1}{\partial p_1^r} & \cdots & \frac{\partial s_1}{\partial p_j^r} \\ \vdots & \ddots & \vdots \\ \frac{\partial s_j}{\partial p_1^r} & \cdots & \frac{\partial s_j}{\partial p_j^r} \end{bmatrix} \times \begin{bmatrix} \frac{dp_1^r}{dp_1^w} & \cdots & \frac{dp_1^r}{dp_j^w} \\ \vdots & \ddots & \vdots \\ \frac{dp_j^r}{dp_1^w} & \cdots & \frac{dp_j^r}{dp_j^w} \end{bmatrix}' = \Delta^d \Delta^p, \quad (13)$$

where Δ^d is the matrix of changes in quantity sold in period t due to changes in retail price with element (k, m) equal to $\frac{\partial s_k}{\partial p_m^r}$ and Δ^p is the matrix of changes in retail price due to changes in wholesale price with element (m, j) equal to $\frac{dp_m^r}{dp_j^w}$.

In vertical retail pricing markets, Δ^p can be a complicated object (see Villas-Boas, 2007). In our context, however, the state’s regulation of alcohol sales simplifies this matrix significantly by committing downstream stores to a uniform pricing policy and by eliminating off-diagonal terms. For example, under the current pricing rule, $\frac{dp_j^r}{dp_j^w}$ is simply $1.30 \times 1.18 = 1.534$, reflecting the 30% uniform tax and the 18% liquor tax that translate a change in the wholesale price for product j to a change in the product’s retail price. Further, the retail price for product m does not respond to a change in the wholesale price for product j ; the markup rule implies that $\frac{dp_m^r}{dp_j^w} = 0 \forall m \neq j$. The limited ability of the *PLCB* to respond to changes in wholesale prices chosen by the upstream firms results in a diagonal price response matrix Δ^p . The full demand response matrix Δ^w of equation (13) thus becomes

$$\Delta^w = \begin{bmatrix} \frac{\partial s_1}{\partial p_1^r} & \cdots & \frac{\partial s_1}{\partial p_j^r} \\ \vdots & \ddots & \vdots \\ \frac{\partial s_j}{\partial p_1^r} & \cdots & \frac{\partial s_j}{\partial p_j^r} \end{bmatrix} \times \begin{bmatrix} 1.534 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1.534 \end{bmatrix}. \quad (14)$$

Given estimates of consumer demand, data on retail prices, and the observed *PLCB* policy. We use this model of upstream behavior to recover product-level marginal costs via (12). Marginal cost estimates facilitate the evaluation of alternative pricing regulations to analyze the aggregate and distributional consequences of the simple uniform markup policy. Modifications to the pricing rule also require adjustment of Δ^p , as more flexible, product-based policies enable the *PLCB* to set markups that vary by product.

4 Estimation Results

We adapt the estimation approach of Nevo (2000) to the institutional features surrounding the price regulation of spirits in Pennsylvania. As in Miravete et al. (2017), we follow a three-step estimation procedure that takes advantage of the fact that the *PLCB* charges the same retail price for a given product in all local markets. This allows us to separately identify the contribution to demand of demographic taste heterogeneity across the state at a point in time and the contribution of time varying shifters of demand that are common across demographic groups, including price. In the first-stage, we use generalized method of moments techniques to estimate the determinants of deviations from a product’s mean utility μ_{ijlt} , controlling for market and product by pricing period fixed effects that absorb the mean effects of price, product characteristics, and seasonality. In the second stage, we employ instrumental variables techniques to project the estimated product by pricing period fixed effects onto price, seasonality, and product fixed effects, using contemporaneous prices in distant control states and input prices as instruments. Last, we project the estimated product fixed effects from the second stage onto time-invariant product characteristics.

Identification of random coefficients Σ relies on correlation between a product’s market share and its characteristics relative to other more or less similar products; see Berry and Haile (2014). We construct two instruments similar to those used in Bresnahan, Stern and Trajtenberg (1997). First, we employ the number of products in the market that share product j ’s characteristic. For example, to identify taste variation for brandies, we use the total number of competing brandies of the same bottle size in location l in period t as the instrument for a given brandy. Second, we use spirit product scores from [Proof66.com](#) as a measure of product quality and compute the average distance, measured in squared deviations, of product j to other products that share its characteristic. Thus, for the above brandy, this would be the average distance in product score space from other brandies available in market l at time t . This instrument provides additional identifying power since it captures differential effects on market shares of offering a high-quality brandy in a market with other high quality brandies relative to an alternative market largely populated by lower quality brandies.

Similarly, identification of demographic interactions Π is based on correlation between the market shares of products with particular characteristics in a given store market and the demographics of the population served by each store. We interact the above two instruments with the prevalence of a given demographic attribute in each market. For example, we would identify the differential taste of young households for the above brandy by interacting the number of brandies and the brandies’ quality level with the share of young consumers in each market. To identify the effect of the income and price interaction, we construct similar instruments based on the number of products sharing a given product’s price category (cheap vs. expensive) interacted with the share of households in the market with income above \$50,000. Berry and Haile (2010) point out that these “Waldfoegel” instruments (Waldfoegel, 2003) are valid provided there exist no demand spillovers from consumers in other similar markets.

Variation in prices over time identifies the price coefficient α , exploiting the fact that distillers do not change the wholesale prices p^w for all products at the same time. We identify seasonality and mean preferences β for time-invariant product characteristics such as proof and spirit type from systematic variation in market shares of spirits by time period or characteristic. We present the estimation algorithm in detail in Appendix C where we also discuss potential sources of price endogeneity stemming from responses by distillers to unobserved demand shocks (via ξ), and the use of input costs and retail prices in other control states as instruments to address such possible confounding.

4.1 Parameter Estimates

Table 7 presents the demand estimates of our preferred specification of the mixed-logit model.²³ Parameters are very precisely estimated and the estimated demand specification captures the documented patterns of spirit consumption across demographic groups (Table 6).

We allow for rich variation across demographics by interacting consumer age and indicators for minority and high educational attainment with proof and indicators for spirit type, bottle size, and import status. The estimates of Π reveal significant differences in tastes for spirits across demographic groups. While minority consumers favor brandy, cordials and rum over gin, the reference category, older and college-educated consumers prefer gin to cordials and rum, all else equal. We also find that older consumers and to a lesser extent college educated consumers, are more likely to purchase 1.75 L than 750 ml bottles, our reference category, while minority households are more likely to purchase 375 ml bottles. Estimated demand for wealthier consumers is steeper, which is consistent with the increased consumption of expensive spirits by “high income” consumers reported in Table 6. Older and particularly minority consumers also favor spirits with higher proof.

We allow for unobserved variation in preferences for a number of the product characteristics, including proof and certain bottle sizes, product categories, and import status. The estimated random coefficients are large, in particular for brandies and for the 375 ml size, indicating that even after controlling for the significant degree of observed differences in tastes on average and by demographic groups, there still exist further similarities between products in these categories that influence their substitution patterns. Lastly, we find that demand increases during the summer and the holiday season and that, on average, consumers prefer products of higher quality and lower proof and favor cordials, rums, and vodkas over gins and brandy.

4.2 Elasticities

Next, we translate the parameter estimates into own-price elasticities that we summarize in Table 8. At the product level, we estimate an average own-price elasticity of -3.75 , which is within the range of median elasticities of Conlon and Rao (2015, Table 8) for spirits (from -2.61 for whiskeys

²³In Appendix D we show that our estimation results are robust to arbitrage on the state border (i.e., the “border bleed”), other samples, and alternative instrumentation approaches.

Table 7: Mixed-Logit Demand

	Mean Utility	Random Coeff.	Demographic Interactions (II)			
	(β)	(Σ)	AGE	EDUCATION	INCOME	MINORITY
PRICE	-0.2763 (0.0046)				0.0787 (0.0026)	
CONSTANT	-34.8299 (0.8218)	0.1759 (0.3653)	6.2002 (0.5176)	5.7245 (0.3197)		-7.3124 (0.6198)
375 ml	4.8700 (0.2451)	2.1181 (0.5896)	0.3947 (0.1487)	-0.7853 (0.1320)		0.8109 (0.1283)
1.75 L	9.0752 (0.2330)	0.0204 (1.1874)	3.2208 (0.7540)	0.9151 (0.1621)		-0.9581 (0.0530)
BRANDY	-60.4569 (0.3636)	10.0606 (0.5963)	13.6819 (1.1016)	-2.5638 (0.1221)		3.4660 (0.1696)
CORDIALS	20.7050 (0.3343)	0.7215 (0.3827)	-6.6553 (0.5039)	-3.5539 (0.1362)		4.4148 (0.3320)
RUM	24.8060 (0.3506)		-7.3399 (0.4846)	-2.4288 (0.0972)		1.9946 (0.1501)
VODKA	24.5847 (0.2760)	0.0819 (0.4193)	-7.2702 (0.4868)	0.6748 (0.1683)		-0.2070 (0.0481)
WHISKEY	-0.9444 (0.3187)	0.3425 (0.3474)	0.8279 (0.1885)	-1.2156 (0.0770)		-1.2939 (0.0435)
FLAVORED	-0.6278 (0.2130)		-0.1211 (0.0482)	-0.1971 (0.0709)		0.6804 (0.0745)
IMPORTED	-0.6931 (0.1960)	0.4807 (0.3356)	0.1772 (0.0810)	1.6193 (0.1162)		0.1681 (0.0417)
PROOF	-14.6819 (0.7327)	0.2244 (0.4845)	1.3367 (0.1999)	-3.9805 (0.3061)		15.9646 (0.7387)
QUALITY	4.0281 (1.2690)					
HOLIDAY	0.4483 (0.0075)					
SUMMER	0.0820 (0.0065)					

Notes: Robust standard errors are reported in parentheses. Estimates for random coefficients (Σ) and demographic interactions (II) based on *GMM* estimation using 2,237,937 observations in 8,470 store-periods and 1,000 simulated agents in each market. AGE is $\log(\text{age} - 20)$, EDUCATION is an indicator variable equal to one if the agent has some level of college education, INCOME is $\log(\text{income})$, and MINORITY is an indicator equal to one if the agent is non-white. Appendix A contains further details. Mean utility contributions of price, holiday, and summer are based on a product fixed effect regression of the product-period fixed effects from the *GMM* estimation after controlling for price endogeneity. Estimates of the contribution of time-invariant product characteristics to mean utility result from a the projection of the estimated product fixed effects onto these observable characteristics. We document robustness of our results in Appendix D.

to -3.80 for vodkas), and somewhat more elastic than the -2.41 own-price elasticity of wine estimate of Aguirregabiria et al. (2016). Our model estimates imply an estimated price elasticity of off-premise spirit demand of -2.48 overall, more elastic than the -1.5 estimate of Leung and Phelps (1993) in their review of the literature on demand estimation for alcoholic beverages. We attribute this difference to our exclusion (due to lack of data) of the on-premise consumption in bars and restaurants, which is likely less price sensitive than off-premise consumption.²⁴

²⁴Perhaps more important, most earlier studies use aggregate consumption data at the state-level whereas we have detailed local data on consumption choices. To corroborate this hypothesis, Appendix D shows that aggregation

Table 8: Own Price Elasticities by Spirit Type, Price, and Size

	Price	Price Elasticity	
		Average	SD
By Spirit Type:			
BRANDY	14.41	-3.64	1.80
CORDIALS	14.08	-3.46	1.35
GIN	15.15	-3.90	1.82
RUM	13.72	-3.38	1.15
VODKA	16.82	-3.95	1.60
WHISKEY	16.77	-3.98	1.63
By Price:			
EXPENSIVE	20.43	-4.74	1.54
CHEAP	10.96	-2.81	0.84
By Bottle Size:			
375 ml	8.94	-2.36	0.89
750 ml	14.53	-3.58	1.32
1.75 L	20.65	-4.74	1.61
ALL PRODUCTS	15.16	-3.75	1.57

“Price” denotes the average price in the relevant category measured in dollars.

Most relevant to our research question is that the own price elasticity estimates exhibit substantial heterogeneity across spirit types and bottle sizes, thus questioning the optimality of a single markup across products. The empirical distribution of estimated price elasticities has the largest spread for brandy and gin and the lowest for rums. A uniform increase in the *PLCB*’s markup will thus trigger a non-uniform response of demand across spirit varieties. The bottom of Table 8 shows that demand responses to a uniform increase of retail prices vary more widely across bottle sizes and price segments than across spirit types. The demand for 375 ml bottles is less elastic than for 1.75 L bottles, with the 750 ml bottles in-between. Similarly, we also find that demand for expensive products is more elastic than the nearly half as expensive cheap spirits.

In Table 9 we summarize demand heterogeneity by comparing sales-weighted average elasticities across markets, thereby providing evidence supporting Posner’s claim that a one-size-fits-all policy leads to cross-subsidies in the presence of a heterogenous population. In markets with a predominance of high-income, college-educated, and, to a lesser extent, high-risk consumers, demand is significantly less price responsive than in markets with concentrations of lower income, less educated, and moderate drinkers. Across products, the demand responsiveness is similar in markets with low and high concentrations of minority consumers, but this aggregate pattern masks significant heterogeneity by spirit types: demands for cordials, gins, and rums are slightly less elastic, and more elastic than average for the remaining spirit types in markets with large shares of minorities. Overall, the large degree of variation in estimated elasticities across spirit types and demographics corroborates the existence of substantial preference heterogeneity.

in our data set drives the price coefficient (and consequently the estimated elasticity for spirits) toward zero. Appendix D also evaluates the sensitivity of our estimates to alternative samples and instrumentation strategies.

Table 9: Estimated Product Elasticities Across Demographics

	RISK		AGE		EDUC.		INCOME		MINORITY	
	Low	High	Low	High	Low	High	Low	High	Low	High
By Spirit Type:										
BRANDY	-3.65	-3.57	-3.42	-3.87	-3.76	-3.44	-3.93	-3.33	-3.52	-3.91
CORDIALS	-3.61	-3.34	-3.27	-3.55	-3.74	-3.15	-3.77	-3.08	-3.67	-3.56
GIN	-4.08	-3.78	-3.75	-3.97	-4.15	-3.62	-4.17	-3.57	-4.09	-3.89
RUM	-3.52	-3.27	-3.20	-3.47	-3.65	-3.08	-3.68	-3.02	-3.56	-3.47
VODKA	-4.02	-3.87	-3.75	-4.08	-4.13	-3.66	-4.21	-3.60	-3.99	-4.09
WHISKEY	-4.06	-3.93	-3.78	-4.17	-4.20	-3.73	-4.28	-3.66	-4.01	-4.22
By Price:										
EXPENSIVE	-4.93	-4.63	-4.50	-4.89	-5.05	-4.37	-5.12	-4.30	-4.89	-4.92
CHEAP	-2.89	-2.74	-2.65	-2.92	-3.01	-2.57	-3.06	-2.51	-2.90	-2.92
By Bottle Size:										
375 ml	-2.39	-2.32	-2.45	-2.27	-2.51	-2.19	-2.55	-2.13	-2.42	-2.45
750 ml	-3.71	-3.49	-3.72	-3.43	-3.82	-3.30	-3.89	-3.22	-3.68	-3.74
1.75 L	-4.87	-4.67	-4.96	-4.52	-5.00	-4.42	-5.11	-4.33	-4.75	-5.03
ALL PRODUCTS	-3.85	-3.67	-3.57	-3.89	-3.98	-3.48	-4.04	-3.42	-3.85	-3.90

Notes: Statistics are average product elasticity in the relevant category-demographic pair. “High” refers to markets in the top 20% while “Low” refers to markets in the bottom 20% for the corresponding demographic trait. Demographic categories defined in Section 2.4.

Random coefficients (Σ) and demographic interactions (Π) allows for flexible substitution patterns across observable characteristics, particularly bottle size and select spirit types. Our estimates indicate that the cross-price elasticity for two products in the same bottle size category is, on average, 1.79 times greater than for product pairs of different bottle sizes (see Table E.1). We also estimate cross-price elasticities for two products of the same spirit type to be 8.9 times greater than for product pairs across spirit types. Brandies appear to be a particularly well defined segment as the large estimated random coefficient in this category implies cross-price elasticities between two brandies that are 39.21 times greater, on average, than those between a brandy and a product of a different spirit type. Cordials, on the other hand, are heterogenous as reflected by an average cross-price elasticity within spirit type of only 1.3 times that across spirit types.

4.3 Implied Upstream Marginal Cost

The theoretical model from Section 3 laid out how consumers and upstream distillers respond to any chosen retail pricing policy by the *PLCB*. To consider the response in upstream behavior to retail pricing policy alternatives, we require an estimate of the upstream firms’ marginal costs. Beginning with our demand estimates, we assume Bertrand–Nash pricing behavior by upstream firms to recover the marginal cost that renders the observed wholesale prices optimal under the current pricing policy and observed ownership structure according to the first-order conditions in equation (12). We rely on these marginal cost estimates in conducting our counterfactual analysis.

We find that the marginal costs of expensive products are on average 2.7 times higher than of inexpensive products and that brandies and whiskey are the least and most costly products,

Table 10: Estimates of Upstream Market Power (Select Firms)

	ALL	DIAGEO	BACARDI	BEAM	JACQUIN	SAZERAC
By Spirit Type:						
BRANDY	45.25	-	-	-	50.73	-
CORDIALS	35.79	30.87	17.19	47.63	57.55	38.69
GIN	37.07	32.52	20.37	39.87	37.80	54.03
RUM	37.36	32.25	38.06	40.81	48.30	-
VODKA	35.82	38.58	-	37.33	38.81	43.01
WHISKEY	34.30	32.01	18.86	36.55	36.17	39.47
By Price:						
CHEAP	45.70	46.45	44.46	45.53	46.48	44.26
EXPENSIVE	27.12	28.92	24.48	27.57	-	26.52
By Bottle Size:						
375 ml	59.18	59.57	65.96	72.07	91.76	47.71
750 ml	36.99	34.16	33.78	44.45	56.80	54.25
1.75 L	29.16	23.19	21.43	29.07	35.21	37.12
ALL PRODUCTS	36.49	34.25	34.99	39.43	46.48	42.33

Notes: Numbers reflect average Lerner indices weighted by quantity sold. Lerner index defined as $100 \times \frac{p^w - \hat{c}}{p^w}$.

respectively, to manufacture on average (see Table E.2 in Appendix E). If brandies and whiskeys are aged more than four years, their marginal costs are approximately 1.5 times higher than of non-aged products. Imported products have 1.8 times the marginal cost of non-imported products on average, reflecting increased transportation costs and import tariffs that the upstream firms pay.

Our cost estimates further indicate that upstream firms have significant market power (Table 10) as they, on average, enjoy a Lerner index of 36.5% – that is for every dollar in revenue the average product yields 36 cents in profit. Products manufactured by larger firms like Diageo and Bacardi tend to have lower Lerner indices (roughly 34%) while smaller manufacturers such as Jacquin and Sazerac operate more niche product portfolios (Table 4) and earn 46.5% and 42.3% of each dollar in revenue, respectively. Across products, CHEAP and 375 ml products tend to be more profitable for all manufacturers while EXPENSIVE and 1.75 L products tend to have lower Lerner indices.

The presence of upstream market power illustrates two points. First, current policy may implicitly redistribute rents between firms or even enable these firms to extract additional rents from consumers and/or the *PLCB*. Second, upstream firms possess the ability to respond to changes in policy – a factor that we must account for in solving for the equilibrium of any alternative policy.

5 What Does the Single Markup Rule Achieve?

In this section we use the estimates of consumer demand and distillers’ marginal costs to identify the winners and losers of the single markup policy among upstream distillers and downstream consumers. As the *PLCB*’s single markup policy amounts to a uniform consumption tax, our analysis seeks to broadly identify common traits of winners and losers of uniform taxation.

We identify winners and losers by comparing the observed equilibrium to a series of alternative policies. In all of the policies we restrict attention to variation in the ad valorem markup and hold unit fees fixed at current levels as they are purportedly linked to transportation costs only. Focusing on the ad valorem markup also allows for generalization of our results beyond Pennsylvania since nearly all states in the US (liquor control or otherwise) employ simple ad valorem excise taxes on alcohol. We also assume the value of the outside option is invariant to the policies we consider, which implicitly fixes retail beer and wine prices. This is a reasonable simplification in our context since the retail beer distribution is competitive and retail wine prices are controlled by the *PLCB*.²⁵

We restrict attention to *subsidy-free* Stackelberg equilibria in which our alternative policies generate at least as much tax revenue as generated by the current policy. The implicit assumption therefore is that state government must meet its budget constraint. We model each alternative policy as a two-stage Stackelberg game in which the regulator chooses its policy anticipating the distillers’ optimal price responses. We allow upstream distillers to re-optimize their wholesale pricing decisions to maximize their profit in response to policy changes and solve for the markup policy that maximizes the government’s objective. We thus account not only for the direct effect of changing the pricing rule, the *mechanical effect*, but also for the adjustment of firms and consumers or the *behavioral response* to policy (Saez, 2001). Accounting for the strategic response of firms stands in contrast to the optimal taxation literature where assuming perfect competition among firms is common. For all policies we assume upstream firms maximize profits by choosing the set of wholesale prices for the products in their portfolio (Equation 12).

5.1 Does the Single Markup Policy Maximize Tax Revenue?

The *PLCB* was originally founded to discourage alcohol consumption to limit the external costs of alcohol. Today, the *PLCB* generates significant revenue for the state’s general fund. Recent legislation furthermore supports the hypothesis that the state government uses it solely to generate tax revenue.²⁶ As our demand estimation allowed us to be agnostic about the state’s objective, we begin by evaluating whether the single markup is effective at generating tax revenue.

We do so by comparing the current equilibrium to three alternative *PLCB* policies and present the results in Table 11. We initially compare the current 30% single markup with the

²⁵ We do not capture effects of alternative spirit pricing policies on tax revenue from wine sales or beer distributor and restaurant license fees. Relative to spirits, wine purchases account for only 40% of the tax revenue generated by state liquor stores. Substitution to off-premise beer or consumption of alcohol in restaurants is likely limited, due to increasing product dissimilarity and different points of sale. Meng, Brennan, Purshouse, Hill-McManus, Angus, Holmes and Meier (2014) find very small elasticities of substitution between the on- and off-premise markets (suggesting that optimal license fees remain largely unaffected by changes in spirit pricing). They also find small cross-price elasticities between beer and spirits (0.113) and wine and spirits (0.163).

²⁶ In June 2016 Pennsylvania governor Tom Wolf signed House Bill 1690 into law as Act 39 of 2016. Act 39, which took effect in August 2016, granted the *PLCB* “common retail marketing abilities including pricing flexibility” in allowing it to price its best-selling items, defined as the top-selling 150 SKUs, “in a manner that maximizes the return on the sale of those items.” Full text available at <http://www.legis.state.pa.us/cfdocs/legis/li/uconsCheck.cfm?yr=2016&sessInd=0&act=39>. At the time of this writing, the *PLCB* has taken advantage of this pricing flexibility only a limited number of times.

Table 11: Aggregate Effects of the single markup

	UNIT	CURRENT	SINGLE	TYPE	PRODUCT
PLCB Markup Rule:	Percent	30.00	23.73	38.82	55.51
Prices:					
Wholesale Price (\$)	Dollars	8.69	8.82	8.62	8.40
Retail Price (\$)	Dollars	14.87	14.41	14.69	15.01
Alcohol Consumption:					
Bottles	Bottles, millions	41.34	44.69	43.97	42.23
Ethanol	Liters, millions	16.54	18.14	19.23	18.06
Profits:					
Tax Revenue	Dollars, millions	255.70	256.60	264.80	280.38
Distillers	Dollars, millions	113.13	128.38	121.55	112.11
Industry	Dollars, millions	368.83	384.98	386.35	392.49
PLCB Share	Percent	69.33	66.65	68.54	71.44

Notes: “Current ” denotes the current scenario with the 30% markup rule. “Single” is the scenario where the *PLCB* maximizes tax revenue with just one markup. “Type” adds flexibility by using 9 common markups defined by spirit type and bottle size while “Product” employs 312 product-specific markups to maximize tax revenue. All counterfactuals Stackelberg equilibria where the regulator accounts for distillers’ re-optimization of wholesale pricing for each alternative markup policy. The average markup we display for “PLCB Markup Rule” does not include the 18% Johnstown Flood Tax. “Industry” denotes the sum of *PLCB* and distiller profits, in millions of dollars. “Tax Revenue” is total *PLCB* tax revenue (profits), in millions of dollars.

single markup that maximizes tax revenue – a policy we call “Single.” We find that the current markup policy achieves 99.6% of the potential tax revenue afforded by a single markup rule. The small increase in tax revenue requires a modest decrease in the single markup from 30% to 23.7% as the regulator anticipates distillers’ response: an increase in average wholesale price from \$8.69 to \$8.82 per bottle. The upstream distiller response partially undoes the reduction in markup as average retail price declines from \$14.87 to \$14.41 per bottle, or 3.1%. This decrease in retail price leads to an 8.1% increase in the number of bottles sold. Interestingly, the distiller response enables these firms to increase their profits by 13.5% – a substantial improvement. Miravete et al. (2017) find similar results under broader assumptions of upstream conduct. Hence, reducing the *PLCB*’s uniform tax rate to 23.7% increases industry profits by \$16.15 million, but 94.4% is captured by upstream firms. We conclude that the *PLCB*’s 30% single markup appears broadly to maximize tax revenue within the confines of uniform markup policies but that small deviations are largely captured by upstream firms.

The use of a uniform rate is common practice for the taxation of goods and services although it is unlikely that such a policy would maximize tax revenue when demand elasticity estimates vary across products, consumers, and firms. This suggests the *PLCB* is leaving money on the table by not leveraging these differences – a point reinforced by the “Single” Stackelberg equilibrium where distillers leveraged their product-level pricing power to extract the majority of any incremental profits generated by deviations in the single markup.

In the “Product” column of Table 11, we present the Stackelberg equilibrium when the *PLCB* leverages the differences across products by choosing the vector of 312 product-specific (e.g., Bacardi Dry, 750 ml) state-wide markups that maximize aggregate tax revenue. Adding flexibility to the pricing rule leads to an increase in the average markup from 30% to 55.5%. This average change masks significant variation among the heterogeneous products. For instance, much of the increase stems from products with relatively inelastic demand: under the new policy the *PLCB* applies markups of 62.1% and 53.4% on average for 375 ml bottle and brandy products, respectively. Products with estimated demands that are more elastic, however, are marked up less: expensive products and 1.75 L bottle products are taxed 37.3% and 39.3% on average, respectively. Distillers respond by decreasing average wholesale price by 3.3% from \$8.69 to \$8.40; the large wholesale price reductions occur in the 375 ml bottle category (−23.9%), among cheap products (−9.9%), and brandies (−9.9%).²⁷ These price reductions ultimately limit the average retail price increase to consumers to only 14 cents from \$14.87 to \$15.01. The aggregate number of bottles sold increases by 2.2%.

The increased flexibility in the markup policy enables the *PLCB* to increase tax revenues \$24.68 million or 9.7% of current revenues – a significant improvement. Integrated industry profits rise by 6.4% while total upstream profits decrease about 1%. Hence, the *PLCB* is able to capture the bulk of the incremental industry profits generated by the shift in its policy, which stands in stark contrast to the case when the regulator is constrained to a single markup. By allowing markups to vary by product and leveraging differences in demand across products, the *PLCB* extracts market power from upstream firms.

One may argue that the product-specific markups we consider in “Product” are unnecessarily complex and thus not feasible (despite being common practice in the private sector). In column “Type” in Table 11 we report the results of a policy where rather than charging a markup per product we restrict the *PLCB* to only 9 markups based on bottle size (3) and spirit type (6) where the markup m_j for a product j of bottle size b and spirit type c is $m_j = m_b + m_c$, where $b = 1, \dots, 3$, and $c = 1, \dots, 6$. This is a similar policy to the one employed in Vermont and enables the *PLCB* to target the heterogeneous preferences of consumers, albeit in a coarser way.

This less flexible pricing policy falls between the “Single” and “Product” policies: average retail prices fall by less than under the optimal uniform markup policy and distillers reduce their wholesale prices by a lesser magnitude. Using only 9 markups to address consumer heterogeneity enables the *PLCB* to increase tax revenue \$8.20 million, or 34.5% of the potential tax revenue increase from moving to fully flexible, product-specific markups. Distillers are much better off under this policy than under product-level markups as they retain their ability to price products individually, increasing profits by 7.4%, or \$8.42 million, relative to the “Product” scenario.

In summary, does the *PLCB* maximize tax revenue? Our analysis offers a two-pronged answer: Within the set of single markup policies, the *PLCB*’s 30% markup rule generates almost the

²⁷See Table E.3 for a decomposition of wholesale and retail prices across policies and product segments.

maximum tax revenue possible. However, a more flexible product-specific markup policy increases tax revenue by 9.7% even after accounting for the pricing response of upstream distillers. We show that product-specific markups enable the *PLCB* to increase tax revenue by leveraging differences in demand elasticities that not only leads to an increase in quantity sold, but also enables the *PLCB* to capture a greater percentage of industry profits. A simpler policy based on markups by spirit type and bottle size captures about a third of these potential incremental profits.

5.2 Measuring Redistribution from Single Markup Pricing

In this section we use the estimated structural model to identify winners and losers of the single markup policy among the regulated upstream firms and consumers, i.e., we measure Posner’s taxation-by-regulation. We do so by comparing the observed equilibrium under the uniform single markup to the *subsidy-free* Stackelberg equilibrium where the *PLCB* chooses product-level markups (i.e., “Product”) and internalizes the diverse set of own and cross-price demand elasticities of the regulated products. This comparison therefore illuminates the implicit cross-subsidies envisioned by Posner.

The idea that regulation can benefit a select few is of course not new. Stigler (1971) first articulated the hypothesis that regulation mostly served the interests of the firms in the regulated industry.²⁸ However, at least some consumers might also benefit from regulation. Laffont and Tirole (1993, §3.9) note that the cross-subsidization of certain population groups by others, for instance via a single markup rule, can be optimal if the regulator intentionally distorts prices to favor a targeted class of consumers.

5.2.1 Taxation by Regulation Among Upstream Distillers

Our analysis begins in the upstream distillery market where we evaluate the induced redistribution of profits among upstream distillers. We hypothesize that *PLCB* policy impacts each of these firms differently – a point we confirm in Table 12 – since distillers operate very different product portfolios (Table 4) and face consumers with different tastes (Table 6).

We find the *PLCB*’s single markup policy increases market power in the upstream market and transfers profits from large firms like Diageo and Bacardi to small firms like Jacquin. As discussed above, the implementation of product-specific markups over the current policy leads to an average reduction in wholesale prices of 3.3%. Much of this decline is driven by smaller firms like Jacquin and Sazerac who lower their prices by 13.2% and 14%, respectively. In contrast, the reduction in wholesale price for large multi-product distillers like Diageo, Beam, and Bacardi is much smaller, 1.77%, 4.1%, and 1.3% on average, respectively. This reflects the ability of these

²⁸Peltzman (1976) built on this view to consider consumers and other interest groups that may influence the design of regulatory rules and eventually the redistribution of rents among constituencies through the political process. These theories of regulation build upon the influential work by Olson (1965) on collective action and politics. Noll (1989) summarizes the political economy aspects of regulatory capture and Laffont and Tirole (1993, §11) elegantly formalize it within a principal-agent model of regulation.

Table 12: Upstream Performance under Alternative Policies

FIRM	Wholesale Price (\$)		Bottles Sold (M)		Margins (%)		Profits (\$M)	
	CURRENT	PRODUCT	CURRENT	PRODUCT	CURRENT	PRODUCT	CURRENT	PRODUCT
Diageo	9.95	9.77	9.14	9.49	34.33	32.15	26.59	28.31
Beam	7.89	7.56	4.01	3.56	39.37	35.81	10.87	9.02
Jacquin	5.21	4.53	3.94	3.44	46.28	40.75	10.11	6.24
Bacardi	9.57	9.45	3.68	4.19	35.11	33.06	9.98	11.97
Sazerac	4.86	4.18	3.29	3.21	42.30	38.02	8.29	5.61
All Firms	8.69	8.40	41.34	42.23	36.54	33.59	113.13	112.11

Notes: “Current” corresponds to the current equilibrium. “Product” corresponds to the Stackelberg equilibrium with 312 product-specific markups. Upstream firms ordered according to total profits observed in the data. “Wholesale Price” is the average wholesale price (p^w) weighted by current bottles sold. “Margin” is defined as the average Lerner index (i.e., $100 \times \frac{p-c}{p}$) for each firm, weighted by current bottles sold.

firms to modify the wholesale prices of their large product portfolios to offset reduced profits from a subset of products.

In terms of profitability, product-level markups entail declines in price cost margins for all firms, though not all firms are worse off as increases in quantity sold offset lower margins for some: e.g., profits for Diageo and Bacardi increase under product-level markups by 6.4% (\$1.71 million) and 20% (\$1.99 million), respectively. This growth comes at the expense of smaller, more specialized firms like Jacquin, whose profits decrease significantly under this alternative policy; Jacquin is a clear beneficiary of the single markup.


While Jacquin’s success under the single markup fits the narrative of Jordan’s *Producer Protection* argument for regulation (Jordan, 1972), we hypothesize that this result amounts to a “home bias” effect where the firm’s success stems from its focus on selling CHEAP products (Table 4) that are relatively inelastic (Table 8). In Table E.3 we show that the *PLCB*’s single markup policy under-prices these products, leading to greater upstream market power in this segment (Table 10). Thus, it could be that the company’s management chose a business strategy that exploits a regulatory policy in its home market. In contrast, large firms like Diageo and Bacardi sell products throughout the world and are subject to a diverse set of regulations, so they gain less from developing a product set to exploit any one particular state’s regulatory policy. The fact that Jacquin is at best the 18th largest distiller based on unit sales in any of the other liquor control states supports this hypothesis.

5.2.2 Taxation by Regulation Among Consumers

In this section we continue our analysis of taxation-by-regulation and ask: Does the single markup benefit certain consumer groups? As above, we use product-level markups as the alternative policy to identify the implicit redistribution generated by the single markup and note that our results are robust to a large set of alternative policies (See Table E.4).

We begin with a discussion of changes in retail price (Figure 2, panel a) where we compare equilibrium prices under the single markup to the equilibrium prices under product-level markups.

Figure 2: Taxation by single markup Regulation Among Consumers



Notes: Statistics represent percent change from the Stackelberg equilibria under the alternative markup policy to the one observed under the current 30% markup rule. Percent change in retail price is weighted by market shares in the data to capture differences in demand across demographic groups. Average of markets with top and bottom 20% of AGE, MINORITY, EDUCATION, and INCOME as defined in Table 5. “Risk” is defined in Section 2.4. In panel (b) we present the distribution of compensating variation $\{CV_i\}_{i=1}^{454}$, denominated in thousands of dollars.

To reflect differences in preferences across consumer types, we first calculate the average retail price of each market under each policy using quantity sold under the current single markup rule as weights. We then compute the percent change of the average retail price in each market under the 30% single markup relative to the alternative policy. Thus, a positive change indicates that consumers pay a higher price under the current 30% markup for the products they favor than under product-level markups.

Panel (a) indicates that by ignoring the heterogeneous preferences documented in Tables 6 and 8, the current 30% single markup rule induces substantial price distortions that vary systematically with demographic traits. For instance, the current regulation results in underpricing of spirits by 2.4% in the least affluent markets while overpricing them by 0.7% in wealthy neighborhoods – a 3.1 percentage point difference across consumers of top and bottom income quintiles. A similar result arises when we classify markets by the share of population with some college education (and to a lesser extent with population age), which is highly correlated with income.

The intuition behind these results is as follows. Recall that lower income consumers prefer 375 ml and CHEAP products (Section 2.4) and that the estimated demand for these products are relatively inelastic (Section 4.2). Upstream firms set wholesale prices for products in their product set taking into account differences in demand across products and the *PLCB*’s markup policy (Equation 12), but they are limited in the degree to which they can extract consumer surplus because they compete with other firms. The *PLCB* on the other hand faces no downstream competition and its policy affects all upstream firms and products. When it enacts the more flexible “Product” markup policy, it is therefore able to more fully internalize these differences in demand. The net result is that the regulator chooses markups that increase the average equilibrium retail prices of 375 ml and CHEAP products (Table E.3), which adversely affects consumers who prefer these products.

Table 13: Beneficiaries of the Single Markup

	$\mathbf{I}(\mathbf{CV}_l > 0)$	\mathbf{CV}_l
AGE	0.9726 (0.4931)	-3.7956 (1.7853)
EDUCATION	-0.4083 (0.1886)	-7.5975 (0.6468)
INCOME	-0.8294 (0.1951)	-3.7976 (0.6915)
MINORITY	0.4048 (0.1000)	1.7433 (0.3382)
N	454	454
R^2	0.2901	0.6616

Notes: In the first column we report the marginal effects of a probit model relating observable demographics to market mean compensating variations of the product-specific relative to the current market outcome. Dependent variable is equal to one when a market has mean compensating variation less than zero. Robust standard errors reported in between parentheses. The R^2 statistic is McFadden’s pseudo- R^2 . In the rightmost column we regress compensating variation divided by total liquor expenditure in market l on market demographic characteristics. All regressions include a constant (not reported). AGE is $\log(\text{average age})$ in the location. EDUCATION, INCOME, and MINORITY are defined in Table 5.

These changes in retail price translate to changes in consumer welfare. We define welfare as compensating variation; a negative mean compensating variation in location l indicates the average amount a consumer in the market would be willing to pay the state to switch to the product-level markup policy. See Section 3.1.1. In Figure 2, panel (b) we see that some (16.7%) consumers benefit from the single markup though most (83.3%) consumers are worse off under the current policy.²⁹

In Table 13 we project the distribution of mean compensating variation onto consumer demographics. In the left column we look for systematic differences between areas where mean compensating variation is positive: i.e., we look to identify common characteristics among the 16.7% of residents who benefit from the single markup. We do so via a simple probit model where the dependent variable is equal to one when a market has mean compensating variation greater than zero. Again, we find that less educated, lower income, and minority households tend to benefit from the current 30% single markup policy.

In the rightmost column we measure how mean compensating variation in each market varies with demographics where we normalize market l compensating variation (\mathbf{CV}_l) by liquor expenditure in market l control for differences in market size. We find that the effect of education is nearly double that of income: on average, increasing the share of households with some college education by one percentage point is associated with a reduction in consumer surplus of twice the

²⁹To make these numbers more concrete, we find that 18% of our store locations have mean compensating variation greater than zero. After weighting each location by the local drinking-age population, we find that 16.7% consumers benefit from the single markup.

amount associated with a one percentage point increase in the share of high-income households. The effect of income, in turn, is approximately twice the size of the effect associated with the share of minority households; an increasing prevalence of minority households is, however, associated with increases in consumer surplus. In summary, our results indicate that less educated, lower income, and minority households tend to benefit from the single markup as the policy generates lower equilibrium retail prices for the products they prefer. To our knowledge, the redistribution due to alcohol taxation has not been documented before.

5.3 Optimal Taxation

In this section we evaluate the aggregate welfare consequences of the single markup. Of course, a product-specific markup policy that maximizes tax revenue is not likely to also maximize consumer surplus. Ramsey (1927) showed that to minimize efficiency costs associated with raising a given amount of tax revenue, the regulator should tax each product as a function of its demand elasticity – a similar mechanism to the product-level markups we consider above. In Table 14, we therefore evaluate the single markup’s optimality by comparing it to alternative policies that set tax rates by product. In “Ramsey” we solve for the Pareto improving policy that maximizes consumer surplus by optimally setting a vector of 312 product markups subject to the budget constraint of generating at least the same revenues as current policy.

In both the earlier “Product” and the current “Ramsey” policies, we compare the current equilibrium to one in which the *PLCB* maximizes an objective – either tax revenue or consumer surplus given a tax revenue target – that does not account for any negative externalities from alcohol consumption. Implicitly, we have thus assumed that the external costs of alcohol consumption are negligible and largely internalized by consumers, or that, for whatever reason, the state chooses not to manage the externality.

Here, we consider the opposite alternative, where we assume that the primary objective of the *PLCB*’s pricing is to manage alcohol consumption. This is motivated by the fact that identifying and measuring the external costs of ethanol consumption is difficult and an open topic of research in the public health literature (see Greenfield et al., 2009; Bouchery et al., 2011; Sacks et al., 2013). We therefore incorporate the externality into the government’s problem in a simple way: we assume that the government uses the full amount of *PLCB* tax revenue to offset the external costs of ethanol consumption. This allows us to estimate the *average* equilibrium shadow externality as the *PLCB*’s observed tax revenue per liter of ethanol contained in its aggregate sales. We denote this average externality cost by λ ; it amounts to the *PLCB*’s tax revenue of \$255.70 million divided by the 16.54 million liters of ethanol sold, or $\lambda = \$15.46$.

Is this a reasonable estimate of the externality? Bouchery et al. (2011) and Sacks et al. (2013) provide estimates of the external costs of alcohol consumption borne by the Pennsylvania state government during this period. Taken together, their results indicate that the state government paid \$0.43 per U.S. standard drink (i.e., 17.7 ml of ethanol) in alcohol-related costs, or

equivalently \$24.29 per liter of ethanol.³⁰ If we further exclude estimated costs attributable to “lost productivity” due to premature mortality and absenteeism (45.4% of total expense), the estimated external cost falls to \$13.36 per liter of ethanol.³¹ Thus, our simple estimate of \$15.46 appears to be a reasonable approximation.

We use this estimated average external cost to formulate an alternative Ramsey pricing problem that considers whether the government can improve consumer welfare by moving to product-specific pricing while at least covering the associated external costs with tax revenue. Sandmo (1976) and, more recently, Kopczuk (2003) show that the optimal Ramsey pricing problem with externalities can be separated into two independent problems jointly defining the markup charged for each product: a Pigouvian tax targeting the externality and deviations from the resulting prices based on demand elasticities to meet revenue targets.³² We find optimal prices that allow the government to maximize aggregate consumer surplus, subject to an *externality-adjusted* budget constraint:

$$T(p^r) - \lambda \times E(p^r) \geq 0$$

where $T(p^r)$ is aggregate tax revenue, $E(p^r)$ is aggregate ethanol consumption (in liters), p^r is the vector of equilibrium retail prices, and $\lambda = \$15.46$ is the external cost of ethanol consumption. In a slight abuse of nomenclature, we call this policy “Pigou” and note that rather than assigning markups / taxes according to the marginal external cost of ethanol consumption, we rely on the average. We compare the aggregate effects of this alternative policy to the base Ramsey policy in Table 14.

We find that aggregate consumer welfare increases under all of the more flexible taxation policies we consider, including under product-level markups designed to maximize tax revenue with no emphasis on consumer welfare (“Product”). This is an interesting result since the aggregate impact to consumers of more flexible pricing is theoretically ambiguous. One contribution of the paper therefore is to provide an empirical example where a uniform consumption tax not only left money on the table, but also decreased consumer welfare.

Moreover, our results indicate the 30% single markup is far from optimal. Absent any externalities, consumers are far better off under the “Ramsey” solution as aggregate consumer welfare increases \$95.08 million, or 15.4% of total observed liquor expenditure, due to lower equilibrium retail prices on average and an increase in quantity sold. Accounting for the external costs associated with alcohol consumption leads the *PLCB* to choose markups closer to the earlier product-level markups that maximize tax revenue (i.e., “Product”) as the \$18.32 million increase in

³⁰Sacks et al. (2013) presents total alcohol-related costs by state. In Bouchery et al. (2011) they decompose the total expense attributable to alcohol consumption by payer, including state and federal government.

³¹The remaining costs are due to health care expenses from treatment of alcohol-related diseases (12.7% of total) plus “other costs” ranging from property damage to incarceration for alcohol-related crimes (41.7% of total). See Bouchery et al. (2011, Table 2).

³²Conlon and Rao (2015, Appendix 2) present the formal solution of the Pigouvian Tax and Ramsey Pricing solutions in the context of alcohol taxation in Connecticut while Griffith et al. (2017) study alcohol taxation using consumer survey data and an assumed functional form for the externality function.

Table 14: Pricing Rules and Welfare

	Unit	CURRENT	PRODUCT	RAMSEY	PIGOU
Avg. Markup Rule	Percent	30.00	55.51	30.66	52.92
Avg. Wholesale Price	Dollars	8.69	8.40	8.85	8.45
Avg. Retail Price	Dollars	14.87	15.01	13.38	14.79
Tax Revenue	Dollars, millions	255.70	280.38	255.70	274.02
Ethanol	Liters, millions	16.54	18.06	24.91	17.72
Distiller Profits	Dollars, millions	113.13	112.11	179.31	122.61
Compensating Variation	Dollars, millions	-	-9.47	-95.08	-56.68
Consumers Better Off	Percent	-	16.68	0.00	0.00

Notes: “Ethanol” is total ethanol content consumed during the sample. “Compensating Variation” is aggregate mean compensating variation $CV = \sum_i CV_i$ where CV_i is defined in Section 3.1.1. “Consumers Better Off” is the percent of Pennsylvania drinking-age residents who benefit from the current 30% markup policy, i.e., have compensating variation greater than zero.

tax revenue compensates for the 7.1% increase in ethanol consumption. In aggregate consumers are better off by \$56.68 million, or 9.2% of their liquor expenditure under “Pigou” taxes. Interestingly, all consumers prefer the “Ramsey” and “Pigou” taxes whereas 16.7% of consumers preferred the 30% single markup to “Product” taxes.

5.4 What is the *PLCB*’s Objective?

Thus far we have not taken a stand regarding the objective of the state legislature in crafting the single markup policy. Instead, we have shown that by restricting the *PLCB* to a 30% single markup policy, the Pennsylvania state legislature leaves money on the table, decreases aggregate welfare, and generates significant redistribution across consumers and firms. We have therefore identified the implicit effects without addressing whether these results are the explicit choice of the regulator or merely the unintended consequence of a policy adopted by a regulator. In this section we interpret our results under different government objectives.

First, it could be that the state relies on the *PLCB* solely to generate tax revenue for the state’s general fund. If true, our results indicate that allowing more flexible pricing (e.g., “Product” and “Ramsey” policies) would increase tax revenue and consumer welfare. The state legislature’s 2016 passage of Act 39 granted the *PLCB* “common retail marketing abilities including pricing flexibility [to set prices on popular products] in a manner that maximizes the return on the sale of those items,” which provides supportive evidence of this hypothesis. If true, our results highlight that the simple uniform markup both fails to meet this objective and generates redistribution among heterogenous consumers and firms.

A second hypothesis is that the regulator indeed cares about the external costs of alcohol consumption. In Section 5.1 we established that the 30% single-markup forgoes 24.68 million in tax revenue, but it also reduces alcohol consumption by 2.1% in terms of bottles and 8.5% in terms of liters of ethanol. Our demand estimates imply that the 30% markup employed by the *PLCB* is statistically significantly different from the tax revenue maximizing single markup

of 23.7%. We therefore cannot reject the possibility that the state overprices spirits in order to reduce consumption. If we define the average shadow value of consumption as foregone tax revenue divided by foregone consumption, the 30% single-markup implies a shadow value of \$16.24 per liter of ethanol. Interestingly, this figure is remarkably close to our estimated average external cost under the assumption that *PLCB* tax revenue is used to fully offset government expenditure related to alcohol consumption. Our results in the “Pigou” alternative policy indicate that more flexible pricing would simultaneously increase aggregate consumer welfare while generating sufficient tax revenue to cover the costs of alcohol consumption. Thus, the single markup is again sub-optimal.

In both of the above two hypotheses our results amount to measuring the cross-subsidies as well as identifying winners and losers of a simple policy that was not seeking redistribution explicitly. Interestingly, proponents of simple policies often argue that such policies are inherently fair as they treat all citizens equally. Our results indicate just the opposite. If the managerial decision-making costs or political feasibility of implementing and maintaining product-level taxes, or something close to it, were sufficiently high, one could also rationalize the use of the single markup. The recent change in Pennsylvania’s pricing regulation stands in contrast to this argument.

An alternative hypothesis is that redistribution might actually be the objective of a regulator (or state legislature) that understands the cross-subsidies simple policy generate. If so, our results provide evidence that governments can use the cross-subsidies generated by simple policy to favor key constituent groups. We find that, compared to all of the alternative policies we consider, the single markup favors less educated, lower income, and minority households by generating lower equilibrium retail prices for the products they prefer, albeit at the expense of significant tax revenue. It is therefore possible that the *PLCB*’s markup rule is a purposeful effort to curry favor with pivotal constituent groups. Testing this hypothesis however requires formally rationalizing the 30% single markup as the outcome of a political equilibrium where state representatives weigh foregone tax revenue against voter approval – a difficult task that we leave as an interesting and important avenue for future research.

6 Concluding Remarks

In this paper we studied the redistributive effects of simple policy when agents have heterogeneous preferences over outcomes. We focus on the regulation of spirits where we observe the pricing decisions of all upstream firms as well as the final retail prices faced by heterogeneous consumers. In our setting, the regulator is legally-mandated to charge a uniform 30% markup on all horizontally-differentiated products it sells in state-run stores. We find significant heterogeneity in the consumption patterns of different consumer types as well as heterogeneity in the product portfolios offered by the upstream suppliers; hence, a uniform policy likely generates winners and losers among both firms and consumers. Restricting policy to a single markup or tax has significant effects on firms, consumers, as well as the state. This paper therefore not only contributes to our

understanding of optimal taxation, it also amounts, to the best of our knowledge, as the first empirical evidence of Posner’s taxation-by-regulation argument.

Our counterfactual analysis allowed product-level markups to account for differences in demand across products but we maintained the *PLCB*’s policy of charging the same price for each product across all local stores at a given point in time. Common pricing across nearby areas, but not across products, is also standard in private retailing industries as it avoids arbitrage concerns across stores and, as shown by Della Vigna and Gentzkow (2017), reduced managerial decision-making costs. Allowing for such zone pricing is an interesting area for future research as coarse spatial price discrimination might provide regulators with sufficient policy instruments to increase tax revenues without necessarily increasing alcohol consumption.

References

- Adams, B. and K. R. Williams (2017) “Zone Pricing in Regional Oligopoly,” Discussion Paper 2079, Cowles Foundation.
- Aguirregabiria, V., D. Ershov, and J. Suzuki (2016) “Estimating the Effect of Deregulation in the Ontario Wine Retail Market,” Mimeo, University of Toronto.
- Atkinson, A. B. and J. S. Stiglitz (1972) “The Structure of Indirect Taxation and Economic Efficiency,” *Journal of Public Economics*, Vol. 1, pp. 97–119.
- Berry, S. (1994) “Estimating Discrete-Choice Models of Product Differentiation,” *RAND Journal of Economics*, Vol. 25, pp. 242–262.
- Berry, S., J. Levinsohn, and A. Pakes (1995) “Automobile Prices in Market Equilibrium,” *Econometrica*, Vol. 63, pp. 841–890.
- Berry, S. T. and P. A. Haile (2010) “Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers,” Discussion Paper 1718, Cowles Foundation. Yale University.
- (2014) “Identification in Differentiated Products Markets Using Market Level Data,” *Econometrica*, Vol. 82, pp. 1749–1797.
- Bloom, N. and J. V. Reenen (2007) “Measuring and Explaining Management Practices Across Firms and Countries,” *Quarterly Journal of Economics*, Vol. 122, pp. 1351–1408.
- Bouchery, E. E., H. J. Harwood, J. J. Sacks, C. J. Simon, and R. D. Brewer (2011) “Economic Costs of Excessive Alcohol Consumption in the U.S., 2006,” *American Journal of Preventive Medicine*, Vol. 41, pp. 516–524.
- Bresnahan, T. F., S. Stern, and M. Trajtenberg (1997) “Market segmentation and the sources of rents from innovation: Personal Computers in the Late 1980s,” *RAND*, Vol. 28, pp. S17–S44.

- Cerda, M., V. Johnson-Lawrence, and S. Galea (2011) “Lifetime income patterns and alcohol consumption: Investigating the association between long and short-term income trajectories and drinking,” Technical report, National Institute of Health.
- Chintagunta, P., J.-P. Dubé, and V. Singh (2003) “Balancing Profitability and Customer Welfare in a Supermarket Chain,” *Quantitative Marketing & Economics*, Vol. 1, pp. 111–147.
- Cho, S. and J. Rust (2010) “The Flat Rental Puzzle,” *Review of Economic Studies*, Vol. 77, pp. 560–594.
- Conlon, C. T. and N. Rao (2015) “The Price of Liquor is Too Damn High: The Effects of Post and Hold Pricing,” Mimeo, Columbia University.
- Dasgupta, P. and J. E. Stiglitz (1971) “Differential Taxation, Public Goods, and Economic Efficiency,” *Review of Economic Studies*, Vol. 38, pp. 151–174.
- Della Vigna, S. and M. Gentzkow (2017) “Uniform Pricing in U.S. Retail Chains,” Working Paper 23996, NBER.
- Dubé, J.-P., J. T. Fox, and C.-L. Su (2012) “Improving the Numerical Performance of Static and dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation,” *Econometrica*, Vol. 80, pp. 2231–2267.
- Finkelstein, A., J. Poterba, and C. Rothschild (2009) “Redistribution by Insurance Market Regulation: Analyzing a Ban on Gender-Based Retirement Annuities,” *Journal of Financial Economics*, Vol. 91, pp. 38–58.
- Greenfield, T. K., Y. Yu, W. Kerr, J. Bond, J. Rehm, and N. Giesbrecht (2009) “Externalities from Alcohol Consumption in the 2005 US National Alcohol Survey: Implications for Policy,” *International Journal of Environmental Research and Public Health*, Vol. 6.
- Griffith, R., M. O’Connell, and K. Smith (2017) “Design of Optimal Corrective Taxes in the Alcohol Market,” Mimeo, Institute for Fiscal Studies.
- Haughwout, S. P., R. A. Lavalley, and I.-J. P. Castle (2015) “Apparent Per Capita Alcohol Consumption: National, State, And Regional Trends, 1977-2013,” Technical report, National Institute on Alcohol Abuse and Alcoholism.
- Hausman, J. (1998) “Taxation by Telecommunications Regulation,” *Tax Policy and the Economy*, Vol. 12, pp. 29–48.
- Hendel, I. (1999) “Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns,” *Review of Economic Studies*, Vol. 66, pp. 423–446.
- Hendel, I. and A. Nevo (2006) “Measuring the Implications of Sales and consumer Inventory Behavior,” *Econometrica*, Vol. 74, pp. 1637–1673.
- Illanes, G. and S. Moshary (2015) “Estimating the Effect of Potential Entry on Market Outcomes Using a Licensure Threshold,” Mimeo, Northwestern University.
- Jordan, W. A. (1972) “Producer Protection, Prior Market Structure and the Effects of Government Regulation,” *The Journal of Law and Economics*, Vol. 15, pp. 151–176.

- Kahn, A. E. (1970) *The Economics of Regulation. Vol. I: Economic Principles*, New York, NY: John Wiley & Sons.
- Kopczuk, W. (2003) “A Note on Optimal Taxation in the Presence of Externalities,” *Economics Letters*, Vol. 80, pp. 81–86.
- Laffont, J.-J. and J. Tirole (1993) *A Theory of Incentives in Procurement and Regulation*, Cambridge, MA: The MIT Press.
- Leung, S. F. and C. E. Phelps (1993) “My Kingdom for a Drink?: A Review of the Price Sensitivity of Demand for Alcoholic Beverages,” in M. Hilton and G. Bloss eds. *Economics and the Prevention of Alcohol-Related Problems*, Washington, DC: National Institute on Alcohol Abuse and Alcoholism, NIH Publication No. 93-3513.
- Linneman, P. (1980) “The Effects of Consumer Safety Standards: The 1973 Mattress Flammability Standard,” *The Journal of Law and Economics*, Vol. 23, pp. 461–479.
- McDonald, J. B. (1984) “Some Generalized functions for the Size Distribution of Income,” *Econometrica*, Vol. 52, pp. 647–663.
- Meng, Y., A. Brennan, R. Purshouse, D. Hill-McManus, C. Angus, J. Holmes, and P. S. Meier (2014) “Estimation of own and cross price elasticities of alcohol demand in the UKA pseudo-panel approach using the Living Costs and Food Survey 20012009,” *Journal of Health Economics*, Vol. 34, pp. 96–103.
- Miller, N. H. and M. Weinberg (2015) “Mergers Facilitate Tacit Collusion: Empirical Evidence from the U.S. Brewing Industry,” Mimeo, Georgetown University.
- Miravete, E. J., K. Seim, and J. Thurk (2017) “Market Power and the Laffer Curve,” Mimeo, University of Texas at Austin.
- Nevo, A. (2000) “Mergers with differentiated Products: The Case of the Ready-to-Eat Cereal Industry,” *RAND Journal of Economics*, Vol. 31, pp. 395–421.
- (2001) “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, Vol. 69, pp. 307–342.
- Noll, R. G. (1989) “Economic Perspectives on the Politics of Regulation,” in R. Schmalensee and R. Willig eds. *Handbook of Industrial Organization*, Vol. 2, New York, NY: North-Holland.
- Olson, M. (1965) *The Logic of Collective Action*, Cambridge, MA: Harvard University Press.
- Orbach, B. Y. and L. Einav (2007) “Uniform Prices for Differentiated Goods: The Case of the Movie-Theater Industry,” *International Review of Law and Economics*, Vol. 27, pp. 129–153.
- Peltzman, S. (1976) “Toward a More General Theory of Regulation,” *Journal of Law and Economics*, Vol. 19, pp. 211–240.
- Posner, R. A. (1971) “Taxation by Regulation,” *BELL Journal of Economics and Management Science*, Vol. 2, pp. 22–50.
- Ramsey, F. P. (1927) “A Contribution to the Theory of Taxation,” *Economic Journal*, Vol. 37, pp. 47–61.

- Sacks, J. J., J. Roeber, E. E. Bouchery, K. R. Gonzales, F. J. Chaloupka, and R. D. Brewer (2013) “State costs of excessive alcohol consumption, 2006,” *American Journal of Preventive Medicine*, Vol. 45.
- Saez, E. (2001) “Using Elasticities to Derive Optimal Income Tax Rates,” *Review of Economic Studies*, Vol. 68, pp. 205–229.
- Sandmo, A. (1976) “Optimal Taxation in the Presence of Externalities,” *Scandinavian Journal of Economics*, Vol. 77, pp. 86–98.
- Seim, K. and J. Waldfogel (2013) “Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Boards Entry Decisions,” *American Economic Review*, Vol. 103, pp. 831–862.
- Shiller, B. and J. Waldfogel (2011) “Music for a Song: An Empirical Look at Uniform Pricing and its Alternatives,” *Journal of Industrial Economics*, Vol. 59, pp. 630–660.
- Somaini, P. and F. Wolak (2015) “An Algorithm to Estimate the Two-Way Fixed Effect Models,” *Journal of Econometric Methods*, Vol. 4.
- Stigler, G. J. (1971) “The Theory of Economic Regulation,” *Bell Journal of Economics and Management Science*, Vol. 2, pp. 3–21.
- Villas-Boas, S. B. (2007) “Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data,” *Review of Economic Studies*, Vol. 74, pp. 625–652.
- Waldfogel, J. (2003) “Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets,” *RAND Journal of Economics*, Vol. 34, pp. 557–68.
- Weyl, E. G. and M. Fabinger (2013) “Pass-through as an economic tool: Principles of incidence under imperfect competition,” *Journal of Political Economy*, Vol. 121, pp. 528–583.

Appendix

A Data


In this section we discuss the data in more detail. We begin with a discussion of how we aggregate the initial daily, store-level *PLCB* data and how we define market areas served by each store. We also address the possibility of stock-outs and how we link the available demographic information to our geographic market definition.

To reduce the size of the estimation sample, we consider the periodicity with which we observe price changes in the data. *PLCB* regulation in place during our sample period allows price to change only for two reasons: permanent and temporary wholesale price changes. Both follow set timing requirements. Permanent price changes can take effect on the first day of one of the *PLCB*'s four-week long accounting reporting periods. Temporary sales, on the other hand, begin on the last Monday of each month and last for either four or five weeks until the day before the last Monday of the following month. Reporting periods and temporary sales periods thus align largely, but not perfectly. To recognize that temporary price reductions are more prevalent than permanent ones (89.7% of price changes in the sample are temporary in nature) and avoid having multiple very short periods, we use sales periods as our time interval. In case of permanent price changes that take effect at the beginning of a reporting period that bisects two sales periods, we assume that the price change takes effect in the sales period that most overlaps with the given reporting period. This results in 22 “pricing periods” during which prices remain constant. In aggregating our daily sales data to the level of sales during a pricing period, we treat a product as being available in a store if we observe a sale at least once during a given pricing period. The length of the pricing period alleviates concern about distinguishing product availability from lack of sales in the period.

Stores exhibit significant variation in the product composition of purchases. These differences reflect heterogeneity in consumer preferences more than differences in the availability of products across stores: Of the 100 best selling products statewide in 2003, the median store carried 98.0%, while a store at the fifth percentile carried 72.0% of these products. Similarly, of the 1000 best selling products statewide in 2003, the median store carried 82.0%, while a store at the fifth percentile carried 44.2% of the products. The product availability at designated “premium” stores is somewhat better than the average, with the median premium store carrying all of the top 100 products and 95.1% of the top 1000 products. In addition, a customer can request to have any regular product in the *PLCB*'s product catalog shipped to his local store for free, should that store not carry the product. In Figure A.1 we demonstrate the product set available to consumers in wealthier markets is greater for 1.75 L and EXPENSIVE products though the difference is small and consumers in poor neighborhoods clearly have access to a large set of these products.

The fact that most stores carry most popular products and can provide access to all products in the catalog easily, together with the absence of price differences across stores, supports an important assumption underlying our demand model: Differences in product availability do not drive customers' store choices to a significant degree and as a result, consumers visit the store

Figure A.1: Product Availability and Income



Notes: A product is considered in the product set of a geographic market if it is ever sold during 2003-2004. Numbers reflect the average number of products in each category (e.g., 1.75 L products) carried by stores in the relevant income group.

closest to them. In making this assumption, which allows us to focus on the consumer’s choice between different liquor products available at the chosen store, we follow previous studies using scanner data such as Chintagunta et al. (2003).³³

To define the population served by each store, we calculate the straight-line distance to each store from each of Pennsylvania’s 10,351 regular block groups and assign consumers to the closest open store for each pricing period. In instances where the *PLCB* operates more than one store within a ZIP code, we aggregate sales across stores to the ZIP code level; there are 114 such ZIP codes out of a total of 1,775. Note that these instances include both store relocations, where a store moved from one location in a ZIP code to another during 2003, but the data contain separate records for the store in the two locations, and instances where the *PLCB* operates two stores simultaneously within a ZIP code.³⁴ We consider the resulting block group zones as separate markets.

We derive consumer demographics for the stores’ market areas by calculating the total population and population-weighted average demographics. We obtained detailed information on each block group’s discrete income distribution by racial identity of the head of household, with household income divided into one of 16 categories. We aggregate across racial groups and across block groups in a store’s market area to derive the discrete income distribution separately for white and non-white households. We construct two income measures. First, we calculate the share of high-income households by minority status, defined as households with incomes above \$50,000. Second, we fit continuous market-specific distributions to the discrete distributions of income conditional on minority status. We employ generalized beta distributions of the second kind

³³Near the state’s borders, the *PLCB* runs seven outlet stores that sell products, such as multi-packs, not available in regular stores to reduce the so-called ‘border bleed’ of consumers’ shopping in lower-priced neighboring states. The addition of these stores to the sample has little qualitative or quantitative effects on the results. See Appendix D

³⁴We drop wholesale stores, administrative locations, and stores without valid address information, for a total of 13 stores.

to fit the empirical income distributions. McDonald (1984) highlights that the beta distribution provides a good fit to empirical income data relative to other parametric distributions. Similarly, used a generalized beta distribution to estimate the continuous market-specific age distribution though due to data census limitations could not condition this on race or income.

We also obtained information on educational attainment by minority status and aggregated across several categories of educational attainment to derive the share of the population above the age of 25 with at least some college education, by minority status and market.

Our price instruments come from two sources. First, the data on retail prices in other liquor control states consists of monthly product-level shelf prices by liquor control state. We assign a month to our Pennsylvania pricing periods to facilitate a match between the two data sets. Second, we attained historical commodity prices for corn and sugar from Quandl, a data aggregator. The prices are the monthly price of a “continuous contract” for each commodity where a “continuous contract” is defined as a hypothetical chained composite of a variety of futures contracts and is intended to represent a the spot market price of the given commodity. We also attained prices for rice, sorghum, wheat, barley, oats, and glass (as a cost input for bottle size) but found these input costs provided little additional explanatory power.

B Additional Descriptive Statistics


Table B.1 presents the distribution of bottle prices contained in our sample of 312 products. Average price is increasing across bottle sizes both within a category and for the whole sample. Vodkas tend to be the most expensive products, while rums are least expensive.

Table B.1: Average Price and Market Shares by Type and Size

	Products	Avg. Price	Share of Market	
			By Quantity	By Revenue
BRANDY	26	14.41	7.26	6.75
375 ml	7	8.54	1.75	1.09
750 ml	13	15.56	4.28	4.13
1.75 L	6	18.76	1.22	1.52
CORDIALS	62	14.08	13.59	13.71
375 ml	13	10.76	2.11	1.49
750 ml	46	14.16	10.80	11.05
1.75 L	3	27.34	0.67	1.17
GIN	28	15.15	6.72	7.04
375 ml	4	7.80	0.62	0.33
750 ml	46	12.40	3.19	2.92
1.75 L	3	21.06	2.91	3.79
RUM	40	13.72	16.31	15.70
375 ml	5	6.59	1.65	0.73
750 ml	23	12.66	9.56	8.11
1.75 L	12	18.71	5.11	6.86
VODKA	66	16.82	32.10	29.80
375 ml	8	8.14	6.76	2.34
750 ml	33	15.54	10.85	11.08
1.75 L	25	21.29	14.50	16.37
WHISKEY	90	16.77	24.03	27.01
375 ml	11	9.12	2.33	1.37
750 ml	42	15.50	11.61	11.70
1.75 L	37	20.49	10.10	13.94
ALL PRODUCTS	312	16.35	100.00	100.00

Notes: “Quantity” market share is based on bottles while “Revenue” is based on dollar values.

Figure B.1: Spatial Distribution of Consumer Characteristics



Notes: Maps correspond to the spatial distribution of characteristics in Pennsylvania during the sample. Outlined polygons correspond to geographic markets (i.e., “stores” in the text). Dark shaded regions correspond to markets in the top 20% for the corresponding demographic trait (“High” in the text). Lightly shaded regions correspond to markets in the bottom 20% for the corresponding demographic trait (“Other” in the text). Remaining markets (“Other” in the figures) are not shaded.

C Estimation Procedure

In this Appendix, we lay out the three-stage estimation procedure we adopt to estimate contributions to the consumer’s mean utility from a given product, δ_{jlt} , and individual-specific contributions to utility, μ_{ijlt} . We discuss each stage in turn, highlighting the variation in the data that allows us to identify the relevant parameters in each stage.

C.1 Stage 1: Random Coefficients and Demographic Interactions

In the first of the three estimation stages, we estimate the contributions of unobserved (Σ) and observed (Π) demographic interactions to deviations from mean utility, μ_{ijlt} , controlling for location and product by time fixed effects. We decompose the unobserved product valuations, ξ_{jlt} , as follows

$$\xi_{jlt} = \zeta_l^1 + \xi_{jt} + \Delta\xi_{jlt}. \quad (\text{C.1})$$

In equation (C.1), ζ_l^1 is a market fixed effect that captures systematic variation across locations in the preference for spirits consumption, relative to beer and wine.³⁵ We control for systematic variation in preferences for a given product over time via ξ_{jt} , to reflect the fact that across the state, a product’s mean demand varies over the course of the year. The remaining structural error $\Delta\xi_{jlt}$ represents deviations in unobserved product valuations within a store from these mean product-time valuations, controlling for the average taste for spirits in market l .

This decomposition of ξ_{jlt} simplifies the mean utility of product j , δ_{jlt} in equation (5a), to

$$\delta_{jlt} = \zeta_l^1 + \zeta_{jt}^2 + \Delta\xi_{jlt}, \quad (\text{C.2})$$

where the product and time specific fixed effect ζ_{jt}^2 comprises the effect of product characteristics ($x_j\beta$), seasonal buying ($H_t\gamma$), price (αp_{jt}^r), and ξ_{jt} on a product’s mean utility.

Equation (C.2) highlights an advantage to our setting: since price does not vary across locations l , we are able to control for its mean contribution to utility via product by time fixed effects, which we then use in a second stage estimation to isolate α .

Given a guess at $\theta_A = \{\Sigma, \Pi\}$, we solve for the structural error $\Delta\xi_{jlt}(\theta_A)$ using the following algorithm. We first find the mean-utility levels $\delta_{jlt}(S_{jlt}; \theta_A)$ that set the predicted market share of each product, s_{jlt} in equation (7), equal to the market share observed in the data, S_{jlt} .³⁶ To evaluate the integral in equation (7) we simulate for each market l the purchase probabilities of 1,000 randomly drawn heterogeneous consumers who vary in their demographics and income. We construct the sample of simulated consumers for each market by relying on the empirical

³⁵This accounts for the fact that the potential market is defined based on the average Pennsylvanian’s consumption as disaggregated per-capita consumption of alcoholic beverages is not available.

³⁶We make use of the contraction mapping procedure outlined in Appendix I of *BLP*, imposing a tolerance level for the contraction mapping of 1e-14 as advised by Dubé, Fox and Su (2012, §4.2) to ensure convergence to consistent stable estimates.

distributions of the demographic attributes considered above – whether a consumer is young, non-white, college educated, and their income level – incorporating correlations between demographic attributes where possible. Conditional on a realization of a consumer’s minority status, we take random draws from the corresponding income and educational attainment distributions and assign the consumer to an age bin based on the unconditional distribution of age above 21 years in the relevant location. See Appendix A for further details. Since the ambient population of stores changes with store openings and closings over the course of the year, the simulated set of agents changes in each pricing period. Lastly, we account for the unobserved preferences via Halton draws.

Given mean utility levels that equate predicted and actual market shares, we then follow Somaini and Wolak (2015) and use a within transformation of δ to remove the store and product-period fixed effects ζ_t^1 and ζ_{jt}^2 , leaving only $\Delta\xi_{jlt}$. We follow the earlier literature in using a generalized method of moments (*GMM*) estimator that interacts $\Delta\xi$ with within-transformations of suitable instruments Z . Define Z^+ as the within transformation of the instruments matrix; e.g., for instrument k , $Z_{jlt}^{+,k} = Z_{jlt}^k - \overline{Z_{jlt}^k} - \overline{Z_t^k}$. The *GMM* estimator exploits the fact that at the true value of parameters $\theta^* = (\Sigma^*, \Pi^*)$, the instruments Z^+ are orthogonal to the structural errors $\Delta\xi(\theta^*)$, i.e., $E[Z^{+'} \Delta\xi(\theta^*)] = 0$, so that the *GMM* estimates solve

$$\hat{\theta}_A = \underset{\theta_A}{\operatorname{argmin}} \left\{ \Delta\xi(\theta_A)' Z^+ W^+ Z^{+'} \Delta\xi(\theta_A) \right\}, \quad (\text{C.3})$$

where W^+ is the weighting matrix, representing a consistent estimate of $E[Z^{+'} \Delta\xi \Delta\xi' Z^+]$.³⁷ To increase the likelihood of achieving a global minimum, we employed the Knitro Interior/ Direct algorithm suggested by Dubé et al. (2012) starting from several different initial conditions.

C.2 Stage 2: Mean Utility – Price and Seasonality Coefficients

In the second of the three stages of the estimation procedure, we decompose the mean utility implied by the estimated first-stage coefficients $\hat{\theta}_A$, $\delta_{jlt}(\hat{\theta}_A)$, into the associated location and product by type fixed effects, $\zeta_t^1(\hat{\theta}_A)$ and $\zeta_{jt}^2(\hat{\theta}_A)$. We then project ζ_{jt}^2 onto price and the seasonal indicators, controlling for product fixed effects ζ_j ,

$$\zeta_{jt}^2 = H_t \gamma + \alpha p_{jt} + \zeta_j + \xi_{jt}. \quad (\text{C.4})$$

Equation (C.4) highlights the potential for price endogeneity, to the extent that price responds to time varying preference variation for a given product that is common across locations, in the form of, for example, category-specific seasonal variation in consumption. The *PLCB* pricing cannot respond to unobserved demand shocks. However, the predictable link between wholesale and retail prices opens the possibility to spirit prices being endogenous because of the pricing behavior

³⁷In constructing our optimal weighting matrix, we first assume homoscedastic errors and use $W^+ = [Z^{+'} Z^+]^{-1}$ to derive initial parameter estimates. Given these estimates, we solve for the structural error $\Delta\xi$ and construct $E[Z^{+'} \Delta\xi \Delta\xi' Z^+]^{-1}$ as a consistent estimate for W^+ .

of upstream distillers whose wholesale prices reflect, through their products' market shares, the unobserved common tastes for product characteristics of spirits, ξ_{jt} . Recall the pricing optimality conditions in equation (12).

In principle, such endogeneity concerns are mitigated by the fact that distillers need to request both temporary and permanent changes to their wholesale price a number of months before the new price takes effect. Prices thus only respond to predictable variation in a product's demand over time. At the same time, none of the available product characteristics vary across time, limiting our ability to flexibly represent such time varying preference heterogeneity at the level of the product. We therefore use instrumental variables techniques to estimate the parameters in equation (C.4) using the contemporaneous average price of a given product from liquor control states outside of the Northeast and Mid-Atlantic regions (Alabama, Iowa, Idaho, Michigan, Mississippi, North Carolina, Oregon, Utah, and Wyoming) as an instrument for price denoted as Z_B . Our identifying assumption is that cost shocks are national (since products are often produced in a single facility) but demand shocks are at most regional, perhaps due to differences in demographics or climate.³⁸ We add to this instrument changes in input prices, sugar and corn, interacted with spirit-type dummies to account for exogenous cost shifts across spirit types. For instance, a major input for rums is sugar while corn is an input to gins, vodkas, and whiskeys. We found that contemporaneous futures prices worked best while including price-type interactions for barley, glass, oats, rice, rye, sorghum, and wheat did not improve our estimates. Collapsing the second stage parameters into vector θ_B , this implies the following parameter estimates

$$\hat{\theta}_B = (\hat{X}'_B \hat{X}_B)^{-1} \hat{X}'_B \zeta^2, \quad (\text{C.5})$$

where $\hat{X}_B = Z_B(Z'_B Z_B)^{-1} Z'_B X_B$, with $X_B = [H_t \ p_{jt} \ \zeta_j]$. The price coefficient is identified by variation in prices over time, benefiting from the fact that distillers do not change the wholesale prices p^w for all products simultaneously.

C.3 Stage 3: Mean Utility – Product Characteristics Coefficients

In the third and final estimation stage, we recover product fixed effects ζ_j from equation (C.5) and project them onto observable product characteristics x_j , resulting in

$$\hat{\theta}_C = (x'x)^{-1} x' \zeta. \quad (\text{C.6})$$

where mean preferences for these product characteristics are identified by variation in market shares of spirits of differing characteristics, e.g., proof or spirit type.

³⁸For example, whiskey consumption, more so than the consumption of other spirits, peaks during the colder fall and winter months. Whiskey consumption also varies significantly across demographic groups; for example, African American households consume larger amounts of whiskey than other racial groups relative to their baseline levels of spirit consumption.

D Robustness

In this Appendix, we present the results of several alternative demand specifications.

In Table D.1 we demonstrate the robustness of our demand results to alternative samples using a simple OLS multinomial logit demand system. For each model, we regress the logged ratio of product to outside share on product-period and store fixed effects, including interactions between mean demographics and product characteristics (e.g., % minority-X-rum dummy). In Column (i) we presents results using the sample in the main text. This model generates product elasticities that are similar to our preferred mixed-logit model while the elasticity for spirits as a category is more elastic reflecting the IIA problem of logit demand systems (see *BLP*). In Columns (ii)-(iv) we vary the number of markets to show that including markets with premium (i.e., large stores) and border stores (i.e., stores located within five miles of the PA border) as well as the holiday period has little effect on our estimated price coefficient and elasticities. This indicates that restricting the sample has little effect on our results.

**Table D.1: OLS Demand Estimates Based on Different Samples
(Multinomial Logit Demand)**

	(i)	(ii)	(iii)	(iv)
PRICE	-0.2396 (0.0032)	-0.2469 (0.0033)	-0.2238 (0.0032)	-0.2341 (0.0028)
Product FEs	Y	Y	Y	Y
Premium Stores	Y	N	Y	Y
Border Stores	Y	Y	N	Y
Holiday Period	Y	Y	Y	N
Statistics:				
R^2	0.9584	0.9589	0.9564	0.9736
N	6,852	6,852	6,852	5,606
Elasticities:				
Average	-3.7454	-3.8610	-3.4916	-3.6618
% Inelastic	0.7353	0.3626	0.7477	0.7389
Spirits	-3.3936	-3.5374	-3.1225	-3.3134

Notes: The dependent variable for all models is the estimated product-period fixed effect from a first-stage regression of $\log(S_{jmt}) - \log(S_{0mt})$ onto product-period fixed effects and demographic-product interactions. Robust standard errors in parentheses. “% Inelastic” is the percentage of products with inelastic demand. “Spirits” is the price elasticity of total *PLCB* off-premise (i.e., sold in a state-run store) spirit sales. “Premium Stores” are a *PLCB* designation. These stores typically carry greater number of products. “Border Stores” are stores located within five miles of the Pennsylvania border.

In Table D.2, we show that our estimation approach based on disaggregated data provides superior identification. In Model (i) we deviate from our multi-step approach and estimate the model in a single step, regressing the logged ratio of product share to outside share on price, brand fixed effects, bottle size fixed effects, pricing period fixed effects, market fixed effects, and mean demographic interactions, where brand refers to all bottle sizes of a particular “brand name”, e.g., “Absolut Vodka”. Demand becomes steeper relative to the Model (i) in Table D.1 when following this alternative approach leading to less elastic demand. We see even steeper effects when aggregating product demand across the state (Models iii and iv).

Interestingly, we see that not conducting the estimation via the steps outlined in the text leads to price elasticity estimates found by Leung and Phelps (1993) as well as other studies. Less elastic product demands increase estimated dollar markups for upstream firms, ultimately driving down estimated upstream marginal costs. Miravete et al., 2017 show using similar data that spirit category elasticities presented in the health literature (e.g., Leung and Phelps, 1993) imply negative marginal costs for these upstream firms. Table D.2 therefore suggests that such studies may suffer from an aggregation bias that leads to less elastic estimated demand.

**Table D.2: OLS Demand Estimates Using Different Approaches
(Multinomial Logit Demand)**

	(i)	(ii)	(iii)	(iv)
PRICE	-0.1224 (0.0004)	-0.0513 (0.0003)	-0.0822 (0.0022)	-0.0103 (0.0016)
Brand FEs	Y	N	Y	N
Statistics:				
R^2	0.5129	0.2420	0.8218	0.1441
N	2,237,937	2,237,937	6,852	6,852
Elasticities:				
Average	-1.9133	-0.8028	-1.2853	-0.1610
% Inelastic	12.9738	77.7657	39.1113	100.0000
Spirits	-1.7512	-0.7393	-1.1805	-0.1488

Notes: The dependent variable for models (i)-(ii) is $\log(S_{jmt}) - \log(S_{0mt})$ while it is $\log(S_{jt}) - \log(S_{0t})$ for models (iii)-(iv). Robust standard errors in parentheses. “% Inelastic” is the percentage of products with inelastic demand. “Spirits” is the price elasticity of total *PLCB* off-premise (i.e., sold in a state-run store) spirit sales.

In Model (ii) we replace the product fixed effects with observable characteristics (e.g., dummies for spirit type, imported). Demand becomes even steeper and demand becomes much more inelastic due to the coarseness of our observable characteristics. For example, two brands of imported rum could have different unobservable quality to consumers thereby leading to different product shares and firms choosing to charge different prices but in this specification, the estimation wrongly correlates differences in price with the differences in shares (quantity sold). In Models (iii)-(iv) we aggregate consumption to the state-level requiring us to drop the demographic interactions but otherwise using the same controls as Models (ii)-(iii). Again, we see the inclusion of brand fixed effects is important to absorbing differences in unobservable (to the econometrician) characteristics across brands. We further see that aggregation drives the elasticity of off-premise spirits to become more inelastic, well within the set of estimates included in Leung and Phelps (1993).

As discussed in Section C.2, we use the contemporaneous average price in distant control states as an instrument for price in the second step. In Table D.3, we consider the sensitivity of our results to the particular instrumentation strategy. We compare the estimated price coefficient from alternative two-stage least squares regression models of the estimated first stage product-period fixed effects underlying the estimates in Table 7 projected onto price, seasonal dummies, and product fixed effects.

Relative to *IV1*, our preferred specification, the estimated price coefficients are stable across alternative instruments, and, as expected, entail larger price responses than an uninstrumented *OLS* specification. Each estimated price coefficient is significant at the 95% level and the sets of *IVs* generate significant F-statistics for all specifications. This suggests that, as expected, our time-invariant observable characteristics are unable to fully capture the variation in a product's preferences that distillers base their pricing decisions on.

Table D.3: Price Endogeneity

	<i>OLS</i>	<i>IV1</i>	<i>IV2</i>	<i>IV3</i>
PRICE	-0.2412 (0.0038)	-0.2763 (0.0046)	-0.2781 (0.0046)	-0.2775 (0.0046)
Instruments:				
Input Prices		Y	Y	Y
Alabama		Y		Y
Iowa		Y	Y	
Idaho		Y	Y	Y
Michigan		Y		
Mississippi		Y	Y	
North Carolina		Y	Y	
Oregon		Y	Y	Y
Utah		Y	Y	
Wyoming		Y	Y	Y
F-Statistic		1,280.2	1,235.1	1,235.8
N	6,852	6,852	6,852	6,852

Notes: Specifications include the same covariates as in Table 7.

E Additional Results and Figures

Table E.1: Estimated Cross-Price Elasticities

	Ratio
By Type:	
BRANDY	39.21
CORDIALS	1.30
GIN	2.09
RUM	1.97
VODKA	2.07
WHISKEY	1.88
By Size:	
375 ml	1.88
750 ml	0.89
1.75 L	2.61
By Characteristic:	
FLAVORED	0.83
NOT FLAVORED	2.43
IMPORTED	1.05
DOMESTIC	1.28

Notes: For each product we calculate the average cross-price elasticity for (1) products within the corresponding group and (2) products outside the group. “Ratio” is the average ratio of (1) to (2).

Table E.2: Estimated Marginal Costs (Select Firms)

	ALL	DIAGEO	BACARDI	BEAM	JACQUIN	SAZERAC
By Spirit Type:						
BRANDY	5.34	-	-	-	3.66	-
CORDIALS	6.16	7.18	15.00	3.21	1.98	5.08
GIN	6.43	7.72	12.51	4.90	4.29	2.67
RUM	5.66	7.05	5.48	4.47	3.45	-
VODKA	6.37	6.07	-	4.64	4.24	3.83
WHISKEY	7.11	8.17	14.82	6.05	4.67	4.88
By Price:						
CHEAP	3.67	3.66	3.67	3.62	3.59	3.70
EXPENSIVE	9.04	8.53	10.46	8.08	-	7.78
By Bottle Size:						
375 ml	2.39	2.13	1.45	1.02	0.23	2.89
750 ml	5.81	6.13	5.83	3.51	2.02	2.97
1.75 L	8.24	11.28	11.84	7.54	5.00	4.68
ALL PRODUCTS	6.33	6.33	6.89	5.14	3.59	4.14

Notes: Estimated upstream marginal costs weighted by sales.

Figure E.1: Distribution of Demand Elasticities





Table E.3: Retail and Wholesale Prices by Product Category


	ELAST.	WHOLESALE PRICE (p^w)			RETAIL PRICE (p^r)		
		CURRENT	SINGLE	PRODUCT	CURRENT	SINGLE	PRODUCT
By Spirit Type:							
BRANDY	-3.64	8.12	8.25	7.31	13.90	13.48	15.68
CORDIALS	-3.46	8.93	9.05	8.78	15.10	14.62	15.03
GIN	-3.90	9.13	9.26	8.82	15.59	15.09	16.10
RUM	-3.38	8.34	8.46	8.14	14.32	13.88	14.46
VODKA	-3.95	7.95	8.07	7.58	13.76	13.35	13.73
WHISKEY	-3.98	9.89	10.01	9.70	16.74	16.19	16.62
By Price:							
CHEAP	-2.81	5.85	5.97	5.27	10.50	10.24	11.41
EXPENSIVE	-4.74	11.97	12.10	12.00	19.91	19.21	19.15
By Bottle Size:							
375 ml	-2.36	3.85	3.97	2.93	7.15	7.04	7.72
750 ml	-3.58	8.52	8.64	8.23	14.49	14.04	15.01
1.75 L	-4.74	11.08	11.21	11.06	18.83	18.19	18.23
ALL PRODUCTS	-3.75	8.69	8.82	8.40	14.87	14.41	15.01

Notes: “Elast.” corresponds to the average estimated demand elasticities from Table 8. Other reported statistics are average wholesale and retail price (\$). “Cheap” (“Expensive”) products are those products whose mean price is below (above) the mean price of other spirits in the same spirit type and bottle size. “Single” indicates the counterfactual where the *PLCB* chooses the revenue-maximizing markup. In “Product” the *PLCB* employs 312 product-specific markups to maximize tax revenue.


Figure E.2: Compensating Variation by Consumer Demographics




(a) Income - Low




(b) Income - High




(c) Minority - Low




(d) Minority - High




(e) Education - Low




(f) Education - High




(g) Age - Low



(h) Age - High



(i) Location - Rural



(j) Location - Urban

Table E.4: Robustness of Results to Alternative Policies

DEPENDENT: CV/SPEND	TYPE	PRODUCT	RAMSEY	PIGOU
AGE	-6.0234 (1.7155)	-3.7956 (1.7853)	-6.7260 (4.4153)	0.2724 (2.4097)
EDUCATION	-3.2326 (0.6302)	-7.5975 (0.6468)	-14.0334 (1.774)	-8.4576 (0.8933)
INCOME	0.4230 (0.5688)	-3.7976 (0.6915)	-5.8918 (1.6426)	-4.0895 (0.738)
MINORITY	5.0919 (0.3712)	1.7433 (0.3382)	2.0507 (1.1219)	-0.7192 (0.5286)
R^2	0.5497	0.6616	0.4665	0.6123
N	454	454	454	454

CHANGE IN EQUILIBRIUM VARIABLES	TYPE	PRODUCT	RAMSEY	PIGOU
Change in Tax Revenue (\$M)	-9.10	-24.68	0.00	-18.32
Change in Tax Revenue (%)	-3.44	-8.80	0.00	-6.69
Change in Ethanol Consumption (%)	-14.02	-8.46	-33.61	-6.69
Ethanol Shadow Value (\$)	2.17	16.14	0.00	15.46
Change in Upstream Profits (\$M)	-8.42	1.02	-66.18	-9.48
Change in Upstream Profits (%)	-6.93	0.91	-36.91	-7.73
- Diageo	-6.20	-6.05	-44.42	-17.04
- Beam	-18.02	20.47	-28.08	27.21
- Jacquin	-0.35	61.89	10.78	81.74
- Bacardi	-4.56	-16.66	-48.75	-15.56
- Sazerac	-3.33	47.80	5.58	58.77
- Brown Forman	-18.20	-27.54	-53.79	-37.86
Compensating Variation (\$M)	-15.34	-9.47	-95.08	-22.12
CV as Percent of Spend (%)	2.49	1.54	15.44	3.59
Share of residents better off (%)	6.17	16.68	0.00	0.42
Change in Retail Price (%)	1.20	-0.91	11.11	0.51
- Low Income	-0.4	-2.4	9.3	-0.7
- High Income	1.9	0.7	12.8	2.1
Objective	Tax Revenue	Tax Revenue	Welfare	Welfare
Number of Markups	9	312	312	312
Constrained?	N	N	Y	Y

Notes: In the top panel we regress compensating variation divided by total liquor expenditure in market l on market characteristics. A negative value for mean compensating variation indicates the average consumer in market l is better off under the alternative policy considered. Robust standard errors reported in between parentheses. All regressions include a constant (not reported). AGE is $\log(\text{average age})$ in the location. EDUCATION, INCOME, and MINORITY are defined in Table 5. In the bottom panel we compare current policy to each alternative policy. For example, retail prices are on average 0.91% lower under the 30% single markup than they would be under product-specific markups (i.e., the “Product” column). “Number of Markups” refers to the number of markups chosen by the *PLCB*. A policy is “Constrained” when the *PLCB* chooses its markup policy subject to an aggregate budget constraint. “Product” and “Ramsey” policies are described in the text. “Type” corresponds to the equilibrium when the *PLCB* chooses markups for bottle size (3) and spirit type (6) to maximize tax revenue. “Ramsey” corresponds to the equilibrium when the *PLCB* chooses product markups to maximize consumer surplus subject to generating at least as much tax revenue as under the single pricing rule. “Pigou” corresponds to the Ramsey equilibrium with an ethanol-adjusted government budget constraint.