The Cost of Curbing Externalities with Market Power: Alcohol Regulations and Tax Alternatives

Christopher Conlon

Department of Economics
NYU Stern School of Business and NBER
cconlon@stern.nyu.edu

Nirupama Rao

Business Economics and Public Policy
University of Michigan Ross School of Business
nirurao@umich.edu

November 13, 2023*

Abstract

Products with negative externalities are often subject to regulations that limit competition. The single-product case may suggest that it is irrelevant for aggregate welfare whether output is restricted via corrective taxes or limiting competition. However, when products are differentiated curbing consumption through market power can be costly. Firms with market power may not only reduce total quantity, but distort the purchase decisions of inframarginal consumers. We examine a common regulation known as post-and-hold (PH) used by a dozen states for the sale of alcoholic beverages. Theoretically, PH eliminates competitive incentives among wholesalers selling identical products. We assemble unique data on distilled spirits from Connecticut,including matched manufacturer and wholesaler prices, to evaluate the welfare consequences of PH. For similar levels of ethanol consumption, PH leads to substantially lower consumer welfare (and government revenue) compared to excise, sales or Ramsey taxes by distorting consumption choices away from high-quality/premium brands and towards low-quality brands. Replacing PH with volumetric or ethanol-based taxes could reduce consumption by over 9% without reducing consumer surplus, and increase tax revenues by over 300%.

Keywords: Excise Tax, Pigouvian Tax, Tax Efficiency, Regulation, Vertical Restraints.

JEL Classification Numbers: H21, H23, H25, L12, L13, L42, L51, L66, L81, K21.

*This replaces a previous draft circulated under the title "The Price of Liquor is Too Damn High". The authors thank seminar participants at Stanford, Columbia, Yale, Ohio State, Harvard, Kellogg, Wharton, NYU Law, NYU Stern, Drexel, Boston College, Dartmouth, INSEAD, UBC, Wisconsin, Cornell; and conference participants at Utah WBEC, FTC Microeconomics, and IIOC. We have benefited from many useful discussions from: Michael Riordan, Kate Ho, John Asker, Katja Seim, JF Houde, Alan Auerbach, Daniel Shaviro, Robin Lee, Chris Nosko, Juan Carlos Suaŕez Serrato, Rob Porter, Judy Chevalier, Jacob Burgdorf, Julie Mortimer and Adam Cole. Any remaining errors are our own. Researchers' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researchers and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1. Introduction

The manufacture, distribution, and selling of alcoholic beverages are big business in the United States, with sales approaching $250 billion in 2021. Alcohol markets are also subject to an unusual degree of government intervention. Federal, state, and even local governments levy excise taxes on alcohol, raising more than $18.3 billion annually. Beyond industry-specific taxation, the sale and distribution of alcohol are also tightly regulated. A common state regulation is *post and hold* (PH), which governs wholesale alcohol pricing in 12 states – more than a third of states where alcohol is not sold by a state-run monopoly. These regulations discourage competition among wholesalers, leading to higher prices and lower output.

The PH law we examine requires wholesalers to "post" a uniform price schedule to a state regulator, and "hold" that price schedule for 30 days. Prior to sales taking place, wholesalers are offered a "lookback" period during which they are allowed to match but not undercut competitor prices. Theoretically, we show that PH softens competition and facilitates supra-competitive pricing in the wholesale market. Specifically, the unique iterated weak dominant Nash equilibria of the PH pricing game leads to wholesale prices as high as a single product monopolist would charge.[1] Empirically, we show that PH leads to unambiguously higher prices, particularly for more inelastically-demanded (higher-quality) products, and that if PH were replaced with simple tax instruments we could both reduce alcohol consumption and increase consumer surplus. Understanding these policies is particularly relevant now, given the evolving legal standing of these regulations and the growing interest among state governments in modifying alcohol regulations and increasing alcohol taxes.[2]

At first glance, outsourcing price increases to private firms might seem like an attractive way to limit consumption of alcohol and the associated negative externalities.[3] Intuition from the single-product case suggests that it is irrelevant from a total welfare perspective whether we limit harmful consumption via a Pigouvian tax or by restricting supply through increased market power (perhaps from lax merger approval, weaker antitrust enforcement, or market designs like PH).[4] Indeed, this argument is made by proponents of the "Green Antitrust" movement for allowing consolidation

[1]Thus even the effects on small retailers are ambiguous, as they face uniform but elevated wholesale prices. We provide some evidence of these harms in Appendix D.

[2]The recent split between the Second Circuit court's ruling in *Connecticut Fine Wine & Spirits, LLC v. Seagull* and the prior Ninth Circuit ruling in *Costco Wholesale Corp. v. Maleng* created an opening for the Supreme Court to decide the legal standing of post and hold and related regulations, but the Court denied the request to review the Second Circuit ruling in 2019, leaving the legal parameters of alcohol regulations that restrict competition uncertain. Even more recently, the Federal Trade Commission (FTC), Department of Justice Antitrust Division, and U.S. Treasury Department recently issued a joint report on competition in alcoholic beverage distribution that included a section on post-and-hold (US Department of Treasury, 2022).

[3]The primary rationale given for PH regulations was to address the negative externalities (Greenfield et al., 2009) of alcohol by raising prices to limit consumption (Saffer and Gehrsitz, 2016). A secondary rationale, is to protect small retailers from chain store rivals and PH can be thought of as a stronger form of the Robinson-Patman Act. In Appendix D we show some regression evidence that suggests this too may be ineffective.

[4]Levy et al. (2021) discuss public health externalities regarding the FTC investigation into the merger of cigarette maker Altria and leading e-cigarette (vape) manufacturer Juul.

(and sometimes coordination) among fossil fuel companies, and restricting "excessive competition" has been a key feature of market design in the legalization of marijuana.[5] The interaction of market power and taxes is also a concern in attempts to address the "internalities" of sweetened beverages (Allcott et al., 2019; Dubois et al., 2020; O'Connell and Smith, 2022).

However, the intuition from the single-product case fails when products are differentiated. Put simply, we can think about alcoholic beverages as a bundle of two characteristics: ethanol and branding/quality. A social planner concerned only with limiting the negative externalities would levy a Pigouvian tax on ethanol alone. A firm with market power recognizes that if consumers value both characteristics, it is optimal to "tax" both relative to their elasticities, leading to higher prices on products that consumers value for non-ethanol characteristics. This means that conditional on purchase, firms trade off the desired distortion in ethanol consumption against distortions in the choice of brand (even for *inframarginal consumers*). For example, the cheapest plastic-bottle vodka and the most expensive Scotch might contain equal amounts of ethanol, but differ vastly when it comes to consumer perceptions of quality or willingness to pay. These products will bear the same Pigouvian tax but firms with market power might set very different markups. From a social welfare perspective, this leads not only to markups on premium products that are *too high* but markups on low-end products that are *too low*.

Under PH, this means that consumers who substitute from premium products to inexpensive ones may consume similar amounts of ethanol but be worse off. We show that taxes – even simple tax instruments such as a single-rate sales tax or volumetric tax – can maintain the same aggregate ethanol consumption as PH while increasing consumer surplus. Consumer surplus gains stem from flattening the difference between price and marginal cost across products with the same ethanol content, allowing consumers to shift away from low-priced value brands and towards premium products, leaving most substantially better off. The obvious additional benefit of using taxes instead of market power to limit consumption is that revenue collected also increases threefold.

To assess the welfare implications of PH and tax alternatives, we assemble new, unique data from the Connecticut Department of Consumer Protection and private data sources. These data track the monthly prices of spirits products at the manufacturer, wholesaler, and retailer level, and measure shipments from manufacturers to wholesalers in Connecticut from August 2007 to June 2013. Using these data, we show that retail spirits prices are higher in Connecticut than elsewhere, particularly for premium products, and that spirits consumption in Connecticut is skewed towards lower-quality products despite it being one of the wealthiest states in the country. Following wholesale prices at the product level over time reveals that wholesalers price in parallel, as we would expect given the incentives created by PH, with monthly prices moving in nearly perfect

[5]See Hollenbeck and Giroldo (2021); Thomas (2019) on entry restrictions in marijuana markets; Hollenbeck and Uetake (2021) on market the interaction between taxes and market power in marijuana; and Hansen et al. (2020) for analysis of a (Pigouvian) "potency tax". For Green Antitrust see Kingston (2011) and Linklaters (2020) in favor and Schinkel and Treuren (2020) against.

lock-step.

We combine the price and quantity data and estimate a model of demand for spirits at the wholesale level that allows for correlated preferences among product categories such as gin or vodka, and heterogeneous preferences that vary with income over prices, package size, and overall demand. In addition to matching aggregate purchases, we also match moments based on observed wholesaler markups and individual purchases by income. Our estimates show that lower-end products, which are consumed more heavily by lower-income households, feature both more elastic demands and more substitution to the outside option. The price sensitivity among consumers of low-end products that we estimate arises from matching the observed wholesaler markups.

We use our demand estimates to compare the welfare effects of the existing PH system to a competitive wholesale tier under alternative taxes. We consider: an *ad valorem* sales tax, an ethanol tax like the U.S. federal government imposes, a volumetric tax like most states currently employ, and a minimum price per unit of ethanol. These counterfactuals make clear that PH imposes steep welfare costs by distorting inframarginal purchase decisions. The state could, for example, reduce ethanol consumption by more than 9% without reducing consumer surplus if it replaced PH with volumetric taxes. Meanwhile, revenue from alcohol taxes would more than triple. If revenues scaled similarly across PH states, this would amount to an additional $1B in tax revenue. Ethanol price floors, on the other hand, could reduce ethanol consumption by a quarter without reducing consumer surplus.

Our counterfactuals also yield interesting insights into the effectiveness of different tax instruments. Because we focus exclusively on distilled spirits, there is little distinction between taxes on volume and taxes on ethanol content (most products are around 40% alcohol by volume).[6] We define the frontier of the consumer surplus-tax revenue trade-off and consumer surplus-ethanol consumption (negative externality) trade-off by considering product-specific (Ramsey-style) taxes. We find that conventional *ad valorem* taxes are reasonably close to the frontier that trades off consumer surplus against additional tax revenue, albeit at significantly higher levels of ethanol consumption than the Ramsey-like alternative. Likewise, we find that a price floor per unit of ethanol is quite close to maximizing consumer surplus per unit of ethanol consumed, though not particularly effective at raising revenue. This provides a new interpretation of the objectives behind the minimum ethanol unit price enacted in Scotland (Griffith et al., 2022).

Our main results assume that absent the PH system, the wholesale tier would become perfectly competitive. An obvious criticism is that it is not costless for wholesalers to distribute products. We find that allowing for wholesalers to incur additional marginal costs of $1/L or $2/L (compared to existing price-cost margins of around $3/L) would not significantly affect our welfare results, but would reduce the amount of additional tax revenue that could be collected. (A per-liter tax

[6] If we included beer and wine in addition to distilled spirits this would likely not be the case (see Griffith et al. (2019)) though our previous work Conlon et al. (2024) suggests these may be distinct groups of customers in the United States.

and a per-liter markup or distribution cost are effectively isomorphic, the welfare gains come from having firms *not* price to the own-elasticity). Likewise, an additional concern might be that profit-maximizing manufacturers could respond to a competitive wholesale tier by increasing prices and thus "undoing" some of the benefits of increased competition and higher taxes (a similar concern was raised in Miravete et al. (2020)). We find that allowing manufacturers/distillers to adjust prices increases their profits by as much as 30%, and slightly reduces the additional tax revenue that can be raised with little impact on welfare.

While households across the income distribution gain from replacing PH with Pigouvian or volumetric taxes, the gains are concentrated among high-income households that are more likely to consume premium products, since premium products face the highest markups under PH and thus benefit most from its repeal. Replacing PH with volumetric taxes that hold fixed overall ethanol consumption would raise consumer surplus by 7.4—9.5% among high-income households (those earning $70K or more), while households earning below $35K would see surplus rise by only 2.6—3.5%. The regressivity of benefits from repealing PH means that if we want to leave each of the five income groups we model unharmed, we can reduce overall ethanol consumption by only 3% (instead of 9% when we consider consumer surplus in aggregate).

While it may be somewhat unsurprising that replacing PH and its peculiar incentives with taxes yields efficiency gains, the equilibrium model we estimate provides two key sets of insights. First, estimating demand parameters by matching observed markups reveals which product prices to raise to most efficiently reduce alcohol consumption. Because diversion to the outside good and own-price elasticities are highest for low-end products, unlike PH pricing, taxes act to raise these prices most to curb consumption. To the degree problem drinkers disproportionately consume low quality products (evidence suggests they do), raising these prices may be even more effective than our estimates suggest at curbing the externality. Second, understanding the price sensitivity and allocation of consumers across produces uncovers the distributional impacts of alternative tax policies that differentially change the relative prices of high- and low-end products. Much of the benefits of replacing PH with volumetric taxes accrue to the highest-income households.

Our analysis further provides a full assessment of the consequences of PH pricing for consumers at a time when its legal standing is precarious given recent circuit court decision splits. States such as Washington have reacted to this uncertainty by preemptively abandoning the PH system; others, including Connecticut, have considered repealing PH. Our findings suggest that PH is a costly way to achieve the objective of constraining alcohol consumption, and other policies could more effectively curb consumption with benefits for both consumers and government coffers. Moreover, we can achieve these gains with simple, commonly-employed tax instruments such as taxes on volume or ethanol content. More broadly, our findings suggest that using market power to address negative externalities is inefficient, especially when products are differentiated.

2. Alcohol Regulations and Taxes in the US

2.1. State regulations regarding alcohol beverages

While the federal government imposes substantial taxes on alcoholic beverages, the regulation of alcoholic beverage markets is almost wholly the purview of state governments.[7] Nearly all states that allow alcohol to be sold by private firms have instituted a *three-tier* system of distribution, in which the manufacture, distribution, and sale of alcoholic beverages are vertically separated by law. A common feature of nearly all systems is that retail firms (bars, restaurants, supermarkets, and liquor stores) must purchase alcoholic beverages from an in-state wholesaler.

In 18 states, known as *control states*, the state directly operates the wholesale distribution or retail tier, and in some cases does both. In some states, the state monopoly applies to all alcoholic beverages; in others, just to distilled spirits, but not wine or beer.[8] Recent empirical work has focused on these control states and on understanding the behavior and welfare consequences of state-run monopolies. Miravete et al. (2020) study Pennsylvania's policy of setting a uniform markup (of over 50%) on all products, and Miravete et al. (2018) shows this uniform markup is set above the revenue-maximizing level. Seim and Waldfogel (2013) show that Pennsylvania locates more stores in rural areas and fewer stores in urban areas than a profit-maximizing firm would choose. Other studies have examined how both quantity and prices rose when Washington State privatized its state monopoly. Different authors have offered competing explanations: Illanes and Moshary (2020) explain this phenomenon with increases in product variety, while Seo (2019) focused on increased convenience and one-stop shopping.

The majority of states are like Connecticut, where private businesses own and operate the wholesale and retail tiers. These *license states* often have ownership restrictions that restrict not only cross-tier ownership and cross-state shipping, but also a variety of other practices.[9] For example, welfare effects of both exclusive territories and exclusive dealing in the beer industry have been studied in Sass and Saurman (1993); Sass (2005); Asker (2016). In contrast, spirits wholesalers in Connecticut (and many other states) have a "duty to deal", and must supply all licensed retailers at posted prices.

The distribution of distilled spirits in Connecticut involves a substantial amount of *common agency*. The same product is often sold by as many as four statewide wholesalers. Wholesalers distribute products from multiple competing distiller/manufacturers, and do not divide markets

[7] The 21st Amendment ended Prohibition by turning the power to regulate the import, distribution and transportation of alcoholic beverages within their borders over to the states, largely exempting their regulations from scrutiny under the Commerce and the Import-Export Clauses of the U.S. Constitution. Since then numerous Supreme Court cases have eroded state control over alcohol policy, as the Court has held that state control of alcohol is subject to federal power under the Commerce Clause, the First Amendment and the Supremacy Clause, among others.

[8] A few control states, for example Maine and Vermont, maintain a state monopoly on the distribution and sale of spirits but contract with private firms for retail operations (including pricing).

[9] License states may also impose other restrictions, such as which days alcohol beverages can be sold; whether supermarkets can sell spirits, wine, or beer; and the number of retail licenses a single chain retailer can hold.

geographically. In other words, the *market structure* bears many of the hallmarks of competition, but the *market outcomes* in Connecticut appear anything but competitive. We attribute this to the regulatory environment known as *post and hold*, which governs how wholesalers set prices, and which we describe in the subsequent section.

2.1.1. Legal Environment of Post and Hold

Under PH, manufacturers and wholesalers must offer the same uniform prices to all purchasers, and quantity discounts are prohibited.[10] This is implemented by requiring manufacturers and wholesalers to provide the regulator with a price list for the following period (usually a month). In Connecticut, prices must be posted by the 12th day of the preceding month, and cannot be changed until the next posting period. However, some PH states, including Connecticut, also allow a *lookback* period, during which prices can be amended —but only downwards, and not below the lowest competitor price for the same item from the initial round.[11] During this period, wholesale firms are able to observe the prices of all competitors. In Connecticut, the lookback period lasts for four days after prices are posted. Many states, including Connecticut, also employ a formula that maps posted wholesale prices onto minimum retail prices. This limits retailers from pricing below cost (with limited exceptions to clear excess inventory).[12]

Our analysis focuses on the implications of PH for wholesale prices in Connecticut. Unlike manufacturers who sell only their own products, wholesalers frequently sell overlapping sets of products. While only Diageo manufactures *Smirnoff* vodka, multiple wholesalers distribute *Smirnoff* in any given month. As such we expect that Diageo has market power in the sale of *Smirnoff*; less obvious is why wholesalers located only dozens of miles apart selling identical bottles of *Smirnoff* statewide appear to enjoy even greater market power.

The legal status of PH laws has been challenged in several court cases, with different circuit courts drawing different conclusions as to whether Section 1 of the Sherman Act preempts alcohol-pricing statutes. In a Supreme Court case, *California Retail Liquor Dealers Ass'n v. Midcal Aluminum, Inc* (1980), the court ruled that the wholesale pricing system in California was in violation of the Sherman Act. The California system at the time resembled PH, but with the additional restriction that retail prices were effectively set via a resale price maintenance agreement by wholesale distributors.[13] The court's ruling established a two-part test for determining when

[10]It is worth noting that the *Robinson-Patman Act of 1936* prevents distributors from charging different prices to competing retailers. In practice, Robinson-Patman cases are rare, and offering the same menu of quantity discounts to retailers appears to be sufficient. PH can be viewed as a much stronger version of Robinson-Patman.

[11]In practice manufacturers are not able to amend their prices as each manufacturer is the only purveyor of its brands. As such there is no lower price from a competitor they could amend to match.

[12]There is a long history of policymakers being concerned about retailers using alcoholic beverages as "loss leaders." Some states (not including Connecticut) allow a limited number of "post offs," in which retailers can price below the most recent wholesale price in order to clear inventory. See `https://www.cga.ct.gov/2000/rpt/2000-R-0175.htm` for a list of various state regulations.

[13]It is worth pointing out that prior to the *Leegin* decision in 2007, minimum resale price maintenance was a *per se* violation in the United States.

state actions are immune to antitrust scrutiny: 1. a law must *clearly articulate* a valid state interest (such as temperance) 2. the policy must be *actively supervised* by the state. The PH system was directly challenged in *Battipaglia v. New York State Liquor Authority (1984)*. In this case, Judge Henry Friendly wrote for the Second Circuit:

> New York wholesalers can fulfill all of their obligations under the statute without either conspiring to fix prices or engaging in "conscious parallel" pricing. So, even more clearly, the New York law does not place "irresistible pressure on a private party to violate the antitrust laws in order to comply" with it. It requires only that, having announced a price independently chosen by him, the wholesaler should stay with it for a month.

A more recent challenge in the state of Washington found essentially the opposite. In *Costco v. Maleng (2008)*, the Ninth Circuit's appellate decision affirmed that "the post-and-hold scheme is a hybrid restraint of trade that is not saved by the state immunity doctrine of the Twenty-first Amendment."

In 2019, the Second Circuit (which comprises Connecticut, New York, and Vermont) upheld Connecticut's PH statute —departing from the conclusions of the Fourth and Ninth Circuits, which had previously struck down similar PH provisions in Maryland and Washington state as violations of the Sherman Act. This circuit split could prompt the Supreme Court, which has so far declined to weigh in, to resolve the issue. The Second Circuit majority opinion (perhaps incorrectly) focused on the lack of communication between wholesalers:

> Nothing about this arrangement requires, anticipates, or incents communication or collaboration among the competing wholesalers. Quite to the contrary: A post-and-hold law like Connecticut's leaves a wholesaler little reason to make contact with a competitor. The separate, unilateral acts by each wholesaler of posting and matching instead are what gives rise to any synchronicity of pricing.

The Second Circuit's dissenting opinion sharply criticized the majority's reasoning:[14]

> allow[ing] de facto state-sanctioned cartels of alcohol wholesalers to impose artificially high prices on consumers and retailers across all three states in our Circuit...The problem with Connecticut's law is not that it affirmatively compels wholesalers to collude in order to fix prices, but that it provides no incentive – or ability – for wholesalers to compete on price.

As we illustrate with our theoretical model in Section 3, both parties are partially correct. Connecticut's PH system leads to supra-competitive wholesale prices in a one-shot game via unilateral incentives, without requiring any communication or repeated cooperation among the parties.

[14]We should disclose that we were not engaged or compensated by any parties in the Connecticut case (or any other case). However, previous versions of this paper were cited by the briefs of several parties, including the theoretical result that PH could lead to prices as high as the collusive prices in a static unilateral effects framework.

2.2. Taxes on Distilled Spirits

Federal, state, and even some municipal governments levy their own excise taxes on distilled spirits. The overwhelming majority of these taxes take the form of specific taxes, which are a fixed dollar amount per unit (either volume or alcohol content), though in most states the general sales tax also applies to alcohol purchases.[15]

Federal taxes are remitted by the distiller/manufacturer, or upon import.[16] At the federal level, distilled spirits are generally taxed at $13.50 per proof-gallon, where a proof-gallon is one liquid gallon that is 50 percent alcohol. Most spirits are bottled at 80-proof or 40% alcohol by volume (ABV), and incur $2.85/L in federal taxes. Flavored spirits (generally 60-proof) incur lower taxes, and overproof spirits (often over 100 proof) pay higher taxes per liter.

Most state excise taxes, on the other hand, are volumetric, meaning they do not vary by alcohol content, and are remitted by the wholesaler. Connecticut's specific tax on spirits was raised from $4.50 per gallon ($1.18 per liter) to $5.40 per gallon ($1.42 per liter) on July 1, 2011. We use the timing of this tax increase as an instrument in our analysis. Like most states, Connecticut includes alcohol products in its general retail sales tax base. Connecticut also increased its general sales tax rate from 6% to 6.35% when it raised its excise tax on alcohol.

As a share of the overall retail price, these excise taxes can be large, particularly for the least expensive products. For example, a 1.75L bottle of 80-proof vodka in Connecticut (after 2011) includes $7.48 in combined state and federal taxes. At the low end of the spectrum, a 1.75L plastic bottle of *Dubra Vodka* (one of the best-selling and least expensive products) typically sells for $11.99 at retail; taxes therefore account for greater than 60% of the price. On the other end of the spectrum, a 750mL bottle of premium vodka (*Grey Goose* or *Belvedere*) or Scotch whisky (*Johnnie Walker Black*) might retail for over $40.00, of which only $3.21 (about 8%) would go to taxes.

3. A Theoretical Model of Post and Hold

Our theoretical model shows that the post-and-hold system used by Connecticut functions like a "price matching game." This eliminates the incentive to cut prices and increase market share. Even when multiple firms sell identical products, the iterated weak-dominant strategy is to set the monopoly price and then match any competitor price in the second stage. This will lead to higher prices when compared to competitive wholesale markets. We consider both a simple single-product example in Section 3.1 and also more realistic example with multi-product firms in Section 3.2.

[15] This applies largely to license states. In control states it is hard to differentiate retailer markups from *ad valorem* taxes.

[16] Imported spirits may also be subjected to additional *ad valorem* tariffs. In October 2019, President Trump imposed a 25% tariff on Scotch Whisky imports, which was later suspended for five years in June 2021 by the Biden administration.

3.1. PH with a Single Homogenous Good

Consider the following two-stage game among wholesale firms (designed to resemble the actual PH process in Connecticut described in Section 2). In the first stage, each wholesaler submits a uniform price to the regulator. Then, the regulator distributes a list of all prices to the same wholesale firms. During the second stage, firms are allowed to revise their prices with two caveats: a) prices can only be revised downwards from the first stage price, and b) prices cannot be revised below the lowest competitors' price for that item. Only after this second stage is demand realized. To start, we focus on the case of a single product:

1. **Price Posting:** Each wholesale firm $f \in \mathcal{F}$ submits an initial price p_0^f to the regulator.
2. **Lookback:** Firms observe all initial prices and may choose any price $p^f \in [\underline{p_0}, p_0^f]$ where $\underline{p_0} = \min_g \{p_0^g\}$ (the lowest initial price among all competitors).
3. **Sales take place:** Only after all prices are amended do sales take place.

Suppose that consumer demand is described by $Q(P)$, where P is the "market price," and firms charging p^f face demand:[17]

$$q^f(p^f, p^{-f}) = \begin{cases} 0 & \text{if } p^f > \min_g p^g; \\ \frac{Q(P)}{\sum_g \mathbb{I}[p^f = \min_g p^g]} & \text{if } p^f = \min_g p^g. \end{cases} \tag{1}$$

If each firm has constant marginal cost mc^f, then in the second stage, firms solve:

$$p^{f*} = \arg\max_{p^f \in [\underline{p_0}, p_0^f]} \pi^f = (p^f - mc^f) \cdot q^f(p^f, p^{-f})$$

which admits the dominant strategy:

$$p^{f*} = \max\{mc^f, \underline{p_0}\}$$

In the second stage, firms match the lowest price from the first stage $\underline{p_0}$ as long as it is above marginal cost. Now consider the first-stage game under the additional assumption of symmetric marginal costs $mc^f = mc$.[18] Given the dominant strategy in the second stage, a (symmetric) subgame perfect Nash equilibrium choice for p_0^f is:

$$p_0^f \in [mc^f, p_m^f]. \tag{2}$$

[17]This is just homogenous goods Bertrand so that firms charging above the "market price" sell zero units and other firms split the market evenly. Later, we consider the case when there is both intra-brand and inter-brand competition.

[18]In the Appendix, we consider the case of heterogeneous marginal costs. In this case, we order the firms by marginal costs and must also check each "limit price" the highest possible price (below the monopoly price) for each possible number of firms. In the case where costs are "sufficiently similar" and demand is "well-behaved" we can rule out most cases of limit pricing.

One possible (symmetric) equilibrium is the monopoly pricing equilibrium. That is, all firms set $p_0^f = p_m$. Here there is no incentive to deviate. In the second stage, all firms split the monopoly profits (symmetric costs rule out limit pricing). Cutting prices in the first stage merely reduces the size of the profits *without any change to the division.* Any upward deviation in the first stage has no effect because it doesn't change $\underline{p_0}$.

Another possible equilibrium is marginal cost pricing. Here there is no incentive to cut one's price and earn negative profits. Also, no single firm can raise its price and increase $\underline{p_0}$ as long as at least one firm continues to set $p_0^f = mc$. There are a continuum of (symmetric) equilibria in between.

While it might appear to be ambiguous as to which price is played in the initial period, there are several reasons to think that the monopoly price is the most likely. First, this is obviously the most profitable equilibrium for all of the firms involved; that is, the monopoly pricing equilibrium Pareto dominates all others. However, Pareto dominance is often unsatisfying as a refinement because it need not imply stability. Therefore, we also show that the monopoly price is the only equilibrium to survive iterated weak dominance.

Proposition 1. *In the absence of limit pricing (or under symmetric marginal costs $mc^f = mc \, \forall f$), the unique equilibrium of the single-period game under iterated weak dominance is the monopoly price: $\sigma(p_0^f, p^f) = (p_m^f, \underline{p_0})$ where $\underline{p_0} = \min_f p_0^f$. (Proof in Appendix A.1).*

An iterated weak dominant strategy is for firms to set their first-stage prices at their perceived monopoly price $p_m^f(mc^f)$; and in the second stage, match the lowest of the prices from the first stage (as long as price exceeds marginal cost) $p^f = \max\{mc^f, \underline{p_m}\}$. While we could extend the analysis to repeated games, because the monopoly price attains the maximum profits in the one-shot game, such analysis would be superfluous here.[19]

3.2. PH with Heterogeneous Costs and Multiproduct Firms

Consider the a multi-product wholesale firm $f \in \mathcal{F}$, which chooses prices for all products they sell $j \in \mathcal{J}_f$. A key feature of our market is that multiple wholesalers sell identical products (ie: *Smirnoff Vodka 750mL*). Following the single-product example in Section 3.1, an iterated weak-dominant strategy is for f to set the initial price p_j^f as if it can do so unilaterally, and then simply to match the lowest competitor price on that product in the second stage (assuming it exceeds marginal cost).

We relax the assumption that firms setting equal prices $p_j^f = p_j^g$ split the market equally and instead allow firms to split the market for each product in a known proportion $\gamma_j^f \perp Q_j(\mathbf{P})$. Now, firm f's sales of product j are given by $q_j^f(\mathbf{P}) = \gamma_j^f \cdot Q_j(\mathbf{P})$, where $Q_j(\mathbf{P})$ represents the total

[19] A more challenging extension would be to think about a different game where prices are locked in for 30 days at a time, but firms do not have a "lookback period". In such a game the monopoly price need not be the unique iterated weak-dominant equilibrium.

demand for product j, and $\mathbf{P}$ represents the vector of prices for all products available in that period.[20]

We write the profits of firm f (if all sellers charge the "market price" P_j) as:

$$\pi_f(\mathbf{P}) = \gamma_j^f \cdot Q_j(\mathbf{P}) \cdot (P_j - mc_j^f) + \sum_{k \in \mathcal{J}_f \setminus \{j\}} \gamma_k^f \cdot Q_k(\mathbf{P}) \cdot (P_k - mc_k^f) \tag{3}$$

If each firm f which sells j could unilaterally set the price, the first order condition of (3) with respect to P_j, and divided by $\gamma_j^f > 0$ would be:

$$\left[Q_j + \frac{\partial Q_j}{\partial P_j}(P_j - mc_j^f) \right] + \sum_{k \in \mathcal{J}_f} \frac{\gamma_k^f}{\gamma_j^f} \cdot \left[\frac{\partial Q_k}{\partial P_j}(P_k - mc_k^f) \right] \geq 0 \tag{4}$$

This is meant to reflect the FOC that governs the initial choice of price for j by f in the first stage. In the second stage, firms should still match the lowest-priced seller (as long as the price exceeds marginal cost). This means that (4) holds with equality for at least one firm (the initial lowest-priced seller), and with inequality (> 0) for the others.

What we would like to do is characterize the equilibrium of these second-stage prices, and identify which firm $f \in \mathcal{F}$ is the price-setter in the first stage. In the data, we observe second-stage prices (which are nearly always identical across wholesalers) but do not observe initial prices.[21] We can rewrite (4) to set marginal revenue equal to marginal cost as if firm f could unilaterally choose p_j:

$$p_j^f \left(1 + 1/\epsilon_{jj}(\mathbf{P})\right) = mc_j^f + \sum_{k \in \mathcal{J}_f \setminus \{j\}} \frac{\gamma_k^f}{\gamma_j^f} \cdot D_{jk}(\mathbf{P}) \cdot (p_k - mc_k^f) \tag{5}$$

Own-price elasticities ϵ_{jj} will vary across products based on the characteristics of those products (including but not exclusively ethanol content) and the demographics of the consumers who purchase them, with less elastic demands leading to higher markups. The right-hand side of (5) represents the full *opportunity cost* of selling j. In addition to the marginal cost mc_j, when customers leave j as the price rises, some fraction (the diversion ratio) $D_{jk} = \frac{\partial Q_k}{\partial P_j} / \left| \frac{\partial Q_j}{\partial P_j} \right|$ switch to k, with margins $p_k - mc_k^f$, and firm f will capture a fraction γ_k^f (as compared to γ_j^f of the customers of j).[22]

This is important because the firm with the lowest opportunity cost will choose the lowest price

[20] We still assume that firms which set $p_f > p_g$ sell zero units. The substantive restriction is that γ_j^f is constant and does not depend on prices. In practice, this allows us to estimate γ_j^f from our shipment data.

[21] Also notice that if a firm reduced its price in the first stage to p'_f so that $\underline{p} \leq p'_f < p_f$, this would have no effect on the market price in the second stage. This is the non-uniqueness of subgame perfect equilibria in (2), whereas the second-stage equilibrium is unique as long as the price-setting firm for each product j doesn't play a weakly-dominated strategy.

[22] See **?** for a more detailed explanation of diversion ratios.

p_j^f, and the other firms will simply match this price. In our empirical example, we observe (or can estimate) all of the objects in the bracketed expression from (5) and thus can determine which firm f is the "price setter" for product j. Taking this to data requires the additional assumption that mc_j does not vary by firm. In practice, the wholesalers' marginal costs are determined primarily by: uniform (by law) manufacturer prices and state excise taxes, both of which we observe.[23]

$$\kappa_{jk} \equiv \frac{\gamma_k^f}{\gamma_j^f}, \text{ such that } f = \arg\min_{f':\gamma_j^{f'}>0} \left[mc_j^{f'} + \sum_{k\in\mathcal{J}_{f'}\setminus\{j\}} \frac{\gamma_k^{f'}}{\gamma_j^{f'}} \cdot D_{jk} \cdot (p_k - mc_k) \right] \tag{6}$$

$$p_j = \frac{1}{1+1/\epsilon_{jj}} \cdot \left[mc_j + \sum_{k\in\mathcal{J}\setminus\{j\}} \kappa_{jk} \cdot D_{jk} \cdot (p_k - mc_k) \right] \tag{7}$$

Once we know which firm "sets the price" for each product j, we can re-write (4) in matrix form as (where $\odot$ denotes the Hadamard product):

$$\mathbf{q}(\mathbf{p}) = \big(\mathcal{H}(\kappa) \odot \Delta(\mathbf{p})\big) \cdot (\mathbf{p} - \mathbf{mc}) \tag{8}$$

where the elements of the matrix $\Delta_{(j,k)} = \frac{\partial Q_j}{\partial P_k}$, and the elements of the vector $\mathbf{mc}$ correspond to mc_j^f for the lowest *opportunity cost* firm from (5). Here, the *ownership matrix* has entries $\mathcal{H}_{(j,k)} = \kappa_{jk} = \frac{\gamma_k^f}{\gamma_j^f}$, which can be interpreted as *profit weights* or how the firm setting the price of j treats \$1 of (market-level) profit from k relative to \$1 of (market-level) profit from j. The profit weights depend on the relative share of the market controlled by f for products j and k. Following a long literature in industrial organization, we can solve the linear system in (8) for the (additive) markups:[24]

$$\boldsymbol{\eta} \equiv (\mathbf{p} - \mathbf{mc}) = \big(\mathcal{H}(\kappa) \odot \Delta(\mathbf{p})\big)^{-1} \mathbf{q}(\mathbf{p}). \tag{9}$$

Even though multiple firms sell identical products in a two-stage game with price matching, we can still recover a mapping from consumer demand for products $(\mathbf{q}(\mathbf{p}), \Delta(\mathbf{p}))$ and price cost margins $(\mathbf{p} - \mathbf{mc})$ using only second-stage prices by constructing the "ownership matrix" of lowest opportunity-cost firms on a product-by-product basis. Our only additional requirement is the assumption that the pivotal firm f for each product j does not play a weakly-dominated strategy.[25]

[23] This seems reasonable because Connecticut is a small and most of the wholesalers are located within a very small geographic region near the center of the state. Allowing for some homogenous (across firms and products) wholesaler cost is a straightforward extension which we consider later on.

[24] See other examples from the IO literature going back to Bresnahan (1987) and Nevo (2001, 2000) for mergers, Villas-Boas (2007) for double marginalization, Miller and Weinberg (2017); Miller et al. (2021) for coordinated effects, and Backus et al. (2021a,b) for partial (common) ownership.

[25] Notably, we don't need to make assumptions about the costs (or off-equilibrium beliefs) of the sellers who are not the lowest opportunity cost wholesaler.

When we take our model to the data in the subsequent section we define wholesale prices as the sum of the additive markup η_{jt} from (9) and the marginal cost (the manufacturer price p^m_{jt}, excise tax τ_{jt}, and any additional marginal cost of wholesaling wc_{jt}):

$$p^w_{jt} = \underbrace{p^m_{jt} + \tau_{jt} + wc_{jt}}_{mc_{jt}} + \eta_{jt}. \tag{10}$$

4. Some Descriptive Evidence

In this section, we present several stylized facts and patterns in the data consistent with the theory in Section 3. We show that: (1) Prices are higher in PH states than in other license states. When comparing Connecticut (our PH state) and Massachusetts (a nearby non-PH license state): (2) prices are higher in Connecticut; (3) relative prices are higher for "premium" products; (4) relative shares are lower for "premium" products. Finally, (5) when multiple wholesalers offer a product, prices largely move in lockstep.[26]

4.1. Cross State Evidence from Retail Prices

Our first set of stylized facts come from the NielsenIQ Retail Scanner Dataset (through the Kilts Center at Chicago Booth). These data report weekly unit sales and total revenue for each product (a unique UPC) for a set of retail stores that voluntarily share their data with NielsenIQ. We use the data from 2013 (the final year in our administrative dataset), and compute a volume-weighted average price for each product for the entire year.[27]

To compare prices we construct an index that measures how the average retail price for a fixed set of products varies across states. Using the 250 best-selling products nationwide, we construct the index value for each state:

$$PI^x = \frac{\sum_{j=1}^{250} p^x_j q^{US}_j}{\sum_{j=1}^{250} q^{US}_j} \tag{11}$$

where q^{US}_j is the retail quantity measured in liters of product j sold nationwide and p^x_j is the per-liter retail price of product j in state x. Figure 1 plots index values for control states, license states with PH, and license states without PH regulations. Dark bars on the left indicate the state excise tax for the national bundle in each license state. Retail prices are always inclusive of the excise tax (but not the general sales tax). We do not separate out excise taxes for control states.

[26] The Appendix also extends panel data analysis by Cooper and Wright (2012) to show that aggregate sales of alcoholic beverages are lower under PH, and employment in the retail sector is also lower under PH.

[27] While coverage across states in the NielsenIQ data for supermarkets is excellent, coverage for liquor stores is imperfect. This is because some control state monopolies don't share data with NielsenIQ at all. In some license states (such as California), supermarkets are allowed to sell distilled spirits leading to good coverage, while in others only standalone liquor stores can sell spirits (including New York, New Jersey, and Connecticut). Other license states such as Rhode Island and Delaware where Nielsen records fewer than 1,000 sales are excluded from the analysis.

Figure 1 illustrates two key facts. First, PH states feature some of the highest prices. In fact, PH states outrank nearly all other license states with one notable exception being Texas, which has a different and unusual market structure. Second, price differences are not fully explained by differences in tax rates. PH states have fairly typical tax burdens, ranking roughly in the middle of the distribution of taxes, but are uniformly in the upper third of the price distribution.

A simple way to think about what would happen if we eliminated PH in Connecticut might be to consider another license state as a counterfactual. For example, Illinois has prices that are approximately \$3 per liter lower, while having tax rates that are roughly double those we see in Connecticut. A more obvious comparison for Connecticut is the neighboring state of Massachusetts which eliminated PH in 1998.[28] The two states are demographically similar, and are likely to have similar local wages and transportation costs. Moreover much of Connecticut is either in the Boston media market, or the shared Hartford-CT/Springfield-MA media market, so we might expect that preferences for distilled spirits might be similar in the two states.[29] However, as Figure 1 suggests, prices are around \$1.90 per liter lower in Massachusetts, while excise taxes are only \$0.35 per liter lower.

In Figure 2 we plot the average retail price per liter in Connecticut against the average retail price per liter in Massachussets for each brand of vodka in the NielsenIQ data. We focus on vodka because it represents around 45% of the sales volume in each state. Different bottle sizes are indicated by color, and within brand, there is a substantial discount in the per-liter price for larger (1.75L) bottles. If the prices were identical in both states, all points would lie along the 45 degree line shown in solid black. Instead, prices in Connecticut generally exceed prices in Massachusetts. Moreover, the price premium is larger for more expensive products. We can see that budget brand *Popov* is priced similarly in the two states. Meanwhile Smirnoff, the most popular brand, is subject to a sizable Connecticut premium, and Belvedere, a high-end brand, is subject to an even larger premium. The best-fit line, $P_{CT} = 0.723 + 1.073 \cdot P_{MA}$, indicates that on average Connecticut consumers pay approximately \$1.45 per liter more for discount vodka, \$2.18 per liter more for mid-tier vodka, and \$3.64 per liter more for premium vodka.[30] (Recall, the tax difference is only a flat \$0.35 per liter).

[28] There has been some confusion in the literature as to whether Massachusetts is a PH state. Cooper and Wright (2012) report that Massachusetts ended PH in 1998 while Saffer and Gehrsitz (2016) draw their data regarding PH laws from the NIAAA's catalogue of wholesale pricing restrictions (`https://alcoholpolicy.niaaa.nih.gov/apis-policy-topics/wholesale-pricing-practices-and-restrictions/3`) which describes Massachusetts as a PH state. To clarify the status of the PH statute in Massachusetts we contacted the Massachusetts Alcoholic Beverages Control Commission. The General Counsel of the Massachusetts Alcoholic Beverages Control Commission explained that "The US District Court ruled the post and hold provision to be unconstitutional, so while it remains 'on the books,' it is not enforced so licensees do not need to post and hold (although they are still required to post prices). The case on point is Canterbury Liquors & Pantry v. Sullivan, 16 F.Supp.2d 41 (D.Mass.1998), as well as a Massachusetts Appeals Court case recognizing the District Court's ruling [in] Whitehall Company Limited v. Merrimack Valley Distributing Co., 56 Mass. App. Ct. 853 (2002)." As such, we follow Cooper and Wright (2012) and treat Massachusetts as a non-PH state after 1998.

[29] The remainder of southern Connecticut is in the New York media market.

[30] Here we've defined discount, mid-tier, and premium vodkas as \$10, \$20, \$40 per liter respectively.

One important distortion of the PH policy that we might expect: Firms with market power charge relatively higher markups on more expensive products, and thus influence the set of products consumers purchase. Again we use Massachusetts as our comparison. In Figure 3, we categorize vodkas based on the *national average* price per liter, and plot the share of sales (by volume) in each price band for each state. The idea is that the national average price captures some objective measure of "quality."[31] The upper panel describes purchase shares by volume for 750mL products, while the lower panel describes 1.75L products. The purchase patterns in Figure 3 show that relative to their Massachusetts neighbors, consumers in Connecticut are more likely to purchase products from the two lowest "quality" groups, and much less likely to purchase products from the two highest "quality" groups.[32] Again, this is purely descriptive, and it may be that preferences for vodka in plastic bottles are higher and preferences for *Grey Goose* are lower in Connecticut for other idiosyncratic reasons.[33] We provide an alternative comparison based on CDFs in Appendix C.

4.2. Administrative Data from Connecticut

Our main dataset is meant to capture the universe of distilled spirits sales at the *wholesale level* in the state of Connecticut from July 2007 through July 2013. This dataset has been collected and compiled by us (the authors), and has not been previously analyzed.

The first data source is the monthly price postings from Connecticut's Department of Consumer Protection (DCP). The PH system necessitates that all *wholesalers* submit a full price list for all products they sell.[34] A similar regulation requires that the manufacturer/distillers (firms like Bacardi, Diageo, Jim Beam, etc.) post prices each month.[35] This means that we see monthly product-level pricing for both the *manufacturer* tier and the *wholesale* tier.

There are several challenges related to data construction. The first is that the format of price filings is irregular. While some firms provide spreadsheets, others provide printed PDF reports, and many provide scans of faxed-in price lists. The second challenge is that a single product such as *Johnnie Walker Red* is sold by a single manufacturer (Diageo) but by as many as four wholesalers, and there is no product identifier that links the product between manufacturer and wholesaler or across wholesalers. This means that all of the matching of products and assignment to a unique product identifier must be done primarily by hand. A third challenge is that reporting of product flavors can be inconsistent: we might see shipments of one flavor (Cherry) but price postings only for another flavor (Orange). Within a brand-size-proof combination, we consolidate multiple flavors

[31] Alternatively, we could think about this as a measure of "expected prices" that is purged of local demand or preference shocks.

[32] A similar pattern holds for 1L bottles, but we exclude these from the analysis since 1L bottles are primarily purchased by bars and restaurants and account for only 4% of retail liquor store sales in Massachussets and Connecticut.

[33] Another possibility is that consumers in Connecticut drive to Massachusetts to save $9 on Grey Goose, but not to save $0.50 on Popov.

[34] Recall that the legislation prohibits quantity discounts, so firms are restricted to *uniform* prices.

[35] Each manufacturer/distiller is the sole seller for each of the brands they produce, unlike the wholesale tier which is categorized by a high degree of *common agency*.

so that *750mL Smirnoff Vodka (Flavored)* is a unique product, but "Orange" or "Cherry" is not.

The most serious limitation of the price-posting data is that we usually don't observe both: (a) the initial price postings; and (b) the amended or revised price postings. In some cases we see only the initial price posting, and some handwritten (or faxed) amendments. In others, we see only initial price postings and don't know whether prices were amended or not. And finally, in other cases, we observe only a list of amendments to prices and no price postings at all. In practice "amended prices" tend to overwrite "initial prices" in the DCP database. When in doubt, we treat price postings as if they are (second-stage) "as amended." This requires some careful data cleaning, and filling prices backwards and forwards when there are gaps.[36] One limitation is that we don't have two separate sets of prices we would need to analyze the two stages of the price-posting process. We can offer anecdotal evidence that when firms amend prices they are required to list the competitor whose price they "match", and this is verified by the DCP. However, one advantage of the model in Section 3.2 is that we only require the second-stage price from the lowest-opportunity cost seller in order to estimate demand and supply.

The second data source tracks shipments of distilled spirits from manufacturer/distiller/importers to wholesalers. These data were obtained from the Distilled Spirits Council of the United States (DISCUS). The DISCUS data track shipments from member manufacturers, generally the largest distillers, to wholesalers for each product.[37] These distillers constitute 78% of total shipments of distilled spirits (by volume) in the state of Connecticut.[38]

A key aspect of the DISCUS data is that it contains all shipments (of covered brands) to the state of Connecticut. This includes products that ultimately end up in bars in restaurants, and those sold in retail liquor stores. Another advantage of the DISCUS data is that we see total shipments not only by product, but to each wholesaler. This lets us estimate the γ_j^f parameters from our theoretical model in (6) directly from the shipment data. The primary disadvantage is that for less popular products, shipments can be lumpy with only a handful of shipments per year. For this reason, we focus our analysis primarily at the *quarterly* level of observation, and for the least popular products (one shipment per year or less, around 6% of total sales) we have to apply some further smoothing. For the 21.9% of products not included in our DISCUS sample, rather than exclude them from the analysis, we impute shipments using the NielsenIQ Retail Scanner data totals from 34 stores in Connecticut. We describe the construction of the quantity data in detail in the Appendix.

[Write this!]

[36] We discuss this in detail in our Data Appendix. Some manufacturers tend to post only the prices of products whose price changed from the previous month, which requires some care in constructing the full sequence of prices.

[37] DISCUS members include: Bacardi U.S.A., Inc., Beam Inc., Brown-Forman Corporation, Campari America, Constellation Brands, Inc., Diageo, Florida Caribbean Distillers, Luxco, Inc., Moet Hennessy USA, Patron Spirits Company, Pernod Ricard USA, Remy Cointreau USA, Inc., Sidney Frank Importing Co., Inc. and Suntory USA Inc.

[38] Some of the largest non-DISCUS members include: Heaven Hill Distillery and Ketel One Vodka.

Table 1 reports summary statistics for the restricted sample of 735 products by category and bottle size.[39] Products are brand-flavor-proof-size combinations, such as *Smirnoff Vodka 750mL* or *Tanqueray Gin 1L*. Vodka is the largest product category, accounting for 208 products, and 44.8% of all spirits liters sold. While a plurality of products are 750mL, it is 1.75L products that account for 56.8% of sales volume. Most products are 80-proof (40% alcohol by volume) and as such proof averages near 80 for most categories and bottle sizes, with some exceptions.[40]

Table 1 also reports the average price and average price-cost margin (or additive markup) net of any taxes: $(p_j - mc_j)$ at each tier of the distribution chain: Manufacturer/Distiller, Wholesaler, and Retailer. Table 2 reports similar information except with the average Lerner markup $L = \frac{p_j - mc_j}{p_j}$ instead of the additive markup, and broken out by manufacturer/distiller instead of by size and category. To produce meaningful summary measures across differently-sized products, product prices and margins are measured in per-liter terms, and all means are weighted by liters sold. Our data are unusual because we observe prices at the manufacturer p^m, wholesaler p^w, and retailer p^r level, as well as the excise taxes τ_j paid by wholesalers. This means we directly observe input costs except at the manufacturer level.[41] The largest manufacturer, Diageo, sells 155 products and accounts for 32.7% of sales by volume, and enjoys the higehst Lerner markups (around 30% on average).

The most important takeaway from Tables 1 and 2 is that the wholesale tier is significantly more profitable than other tiers. A "typical" product retails for slightly under \$20 per liter, with a breakdown of: \$3.97/L of wholesaler margin, \$2.98/L of manufacturer margin, \$2.71/L of retailer margin, and \$1.43 in state and \$2.85 in federal taxes.[42] Moreover, prices (per liter) and markups tend to be higher (for all tiers) on 750mL products than on the more popular (and less expensive) 1.75L products. Our counterfactuals will focus on the case where we make the wholesale tier more competitive and instead use taxes to constrain ethanol consumption and address negative externalities.

4.3. Wholesaler Pricing Behavior

Our main focus is the pricing behavior and market power of the wholesale tier. As many as four wholesalers sell identical products, yet each charges a substantial markup above the manufacturer

[39] We restrict the sample using the following criteria: (1) we only consider products in the top 750 (99.9% of sales volume); (2) only products whose average wholesale price is below \$60/L (mostly excluding rare Scotch Whisky); (3) we exclude Cordials and Liqueurs (e.g. Triple Sec, Baileys, Kahlua) which are generally 20% alcohol by volume or less and possibly complements rather than substitutes for distilled spirits; (4) we exclude Cognacs (e.g. Hennesey and Courvoisier) because these products contain vintage/age statements and are nearly impossible to match across data sources.

[40] Some popular gins and imported Scotch Whiskies are over-proof. Most flavored vodkas are 60 proof, and flavored rums can be as low as 42 proof (e.g. *Malibu Coconut Rum*).

[41] Manufacturer marginal costs are backed out of the first order conditions using our demand estimates and following the procedure described in Appendix B.1. Retail prices come from the NielsenIQ Scanner Dataset for Connecticut, and are available only for select retail stores, while manufacturer and wholesaler prices are statewide.

[42] The remainder being production costs.

price, and identical prices as each other. There are several innocuous possibilities, including the fact that wholesaling activities are costly to produce, provide valuable ancillary services, or that wholesale firms are substantially differentiated in ways we cannot observe.

Absent the PH system, a simple way to think about a counterfactual would be if wholesale markups were competed away so $p_j^w = p_j^m + \tau_j$ (manufacturer price plus excise taxes). We plot the wholesale prices and manufacturer prices in Figure 4. Rather than plot the 45 degree line, we plot the zero markup line: $p_j^w = p_j^m + \tau_j$. We see that (after accounting for taxes) wholesaler price-cost margins are larger on more expensive products, with products like *Grey Goose* and *Johnnie Walker Black* having very high price-cost margins. High wholesale price-cost margins are not exclusive to the most expensive products; the mid-priced product *Smirnoff Vodka* (the overall best-seller) also has a high margin, though other popular yet inexpensive products such as *Dubra Vodka* have small markups.

Figure 5 tracks the wholesale (case) prices of up to four different wholesale firms in addition to the manufacturer price for four popular spirits products: Stolichnaya Vodka (1000mL), Tullamore Dew Irish Whiskey (1750mL), Dewars White Label (750mL), and Johnnie Walker Black (1750mL). For each product, the prices set by the different wholesalers move in near lockstep with one another. While one innocuous explanation might be that this synchronous movement simply reflects changes in input prices, the manufacturer prices plotted alongside the wholesale prices do not support this reasoning. Manufacturer prices change only rarely while wholesale prices move more frequently and together. Instead it appears that wholesalers are pricing in parallel, which is consistent with the price-matching incentives created by PH.

Occasional price deviations are short-lived and typically involve only one of three to four wholesalers selling a product. When this happens, we interpret these deviations as cases where initial price postings rather than "amended" price postings are recorded. In the case of Johnnie Walker Black (1.75L), monthly wholesale prices oscillate between two price points, but for Eder we observe only the higher of the two prices.[43] When there is dispersion in our posted wholesale prices, in nearly 80% of such cases the cause is a single wholesaler with a higher recorded price. For this reason, in our econometric model, we assume that all firms play the iterated weak dominant strategy of matching the lower price in the second-stage. Moreover, because our econometric model looks at prices and quantities at the quarterly level rather than the monthly level, we end up smoothing out some of this higher frequency price variation.

5. Econometric Model of Demand and Supply

In much of the industrial organization literature, the goal of econometric estimates of supply and demand is to estimate own- and cross-elasticities in a setting with endogenous prices, and then use

[43] For some of the months in question, we are able to confirm the dates on the submitted prices are consistent with "initial" prices. A likely explanation is that "amended" prices were submitted via fax or were not properly digitized.

first-order conditions to recover markups and marginal costs.[44] We observe both wholesaler and manufacturer prices directly. Instead, we use these prices and wholesale shipments from Section 4.2 and the unique iterated weak-dominant equilibrium of the price-posting game from Section 3 to inform the elasticities in our demand system.

We rely on the estimated system of demand to explain how consumers will adjust purchase patterns under counterfactual pricing where the wholesale tier has incentives to compete (rather than incentives not to compete). In a sense, we are asking the demand system to do less than the usual case, but we still need it to evaluate counterfactual welfare.

5.1. Demand Specification

Our model for consumer demand assumes that in each period t (quarter), a consumer i makes a discrete choice to purchase a single product j, or chooses not to make a purchase. We define a "product" as a brand-flavor-proof-size combination (e.g., *750mL of Smirnoff Flavored Vodka at 60 proof*). We standardize the purchase volume at one liter and maintain the fiction that a consumer can purchase one liter of any product (irrespective of size) at the per-liter price.[45]

We estimate *derived wholesale demand* using the prices and quantities at the *wholesale* level, and abstract away from retailers. This allows us to capture statewide demand for spirits at bars and restaurants as well as liquor stores. As we document in our prior work (Conlon and Rao, 2020), the retail-pricing decision at liquor stores can be complicated by nominal rigidities around prices ending in 0.99. The main limitation of this approach is that our calculation of Marshallian "consumer surplus" combines both retailers (bars, restaurants, and liquor stores) and final consumers.[46] As we document in Figure 5, we rarely see price dispersion among wholesalers, but when we do we take the *minimum wholesale price*, and assume that all consumers face the same "market price."

We assume that consumer demand follows the random coefficients nested logit model (Brenkers and Verboven, 2006; Grigolon and Verboven, 2013). This model combines the random coefficients logit demand model of Berry et al. (1995) with a nested logit structure on the error term ε_{ijt}. The nesting structure is important because we want to allow for more substitution within a product category (Gin, Rum, Tequila, North American Whiskey, Irish/Scotch Whisky, and Vodka) than across categories.[47] The degree to which consumers substitute within the nest is governed by the parameter ρ, with $\rho = 0$ representing the plain (IIA) logit model, and $\rho = 1$ representing the case where all consumers substitute within the same category.

[44] See for example Nevo (2001) or Backus et al. (2021a) for RTE cereal, Villas-Boas (2007) for yogurt, or Miller and Weinberg (2017) for beer.

[45] Similar assumptions are common in the literature. For example, Nevo (2001); Backus et al. (2021a) assume that consumers purchase a single serving of RTE cereal at the per-serving price.

[46] Margins of retail liquor stores are small compared to those of wholesalers in Table 1

[47] We made this assumption in our original draft, and it has since been adopted in other studies of distilled spirits Miravete et al. (2018) and beer Miller and Weinberg (2017).

The utility of consumer i for product j in market t is given by:

$$u_{ijt} = \beta_i \, x_{jt} + \alpha_i \, p_{jt} + \xi_{b(j)} + \xi_t + \Delta\xi_{jt} + \varepsilon_{ijt}(\rho) \tag{12}$$

Here p_{jt} represents the minimum wholesale per-liter price, and x_{jt} represents additional product characteristics (bottle size, proof), and $(\xi_{b(j)}, \xi_t)$ represent brand and time fixed-effects respectively.[48] We define the individual purchase probability $s_{ijt} = Pr(u_{ijt} > u_{ij't} \mid \alpha_i, \beta_i)$ for all $j \neq j'$. The aggregate market share is given by:

$$s_{jt}(\boldsymbol{\xi_t}; \theta_1, \theta_2) = \int s_{ijt}(\alpha_{it}, \beta_{it}, \boldsymbol{\xi_t}; \theta_1, \theta_2) f(\alpha_i, \beta_i \mid y_i, \theta_2) \, g(y_i) \, \partial\alpha_i \partial\beta_i \partial y_i \tag{13}$$

We allow consumers to have heterogeneous preferences for product characteristics that are determined by observed demographics y_i (income) and unobserved characteristics ν_i (a vector of standard normal draws). We also require that the price coefficient α_i is lognormally distributed, so that all consumers have downward-sloping demand curves. In our main specification we discretize income into five quintiles, and allow each quintile to have a separate set of parameters so that:[49]

$$\begin{pmatrix} \ln \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \overline{\alpha} \\ \theta_1 \end{pmatrix} + \Sigma \cdot \nu_i + \sum_k \Pi_k \cdot \mathbb{I}\{y_i \in \text{bin}_k\} \tag{14}$$

Following Conlon and Gortmaker (2020), we partition the parameters into those that enter the problem linearly θ_1, and those that enter the problem non-linearly or pertain to the endogenous objects $\theta_2 = [\rho, \overline{\alpha}, \Sigma, \Pi]$. Each of the parameters in θ_2 requires at least one instrument for identification. For the vector $\boldsymbol{\xi_t}(\theta_1, \theta_2)$, which sets (13) equal to the observed shares, we can construct conditional moment restrictions of the form $\mathbb{E}[\Delta\xi_{jt} \mid z_{jt}^D] = 0$.

5.2. Specification of Supply Moments

While it is possible to estimate parameters and recover markups from the demand side alone, the original BLP papers (Berry et al., 1995, 1999) found it valuable to impose additional moments from the first-order conditions of firms. As shown in Conlon and Gortmaker (2020), a correctly-specified supply side can aid in the estimation of the θ_2 parameters.

We have already derived an expression for the additive wholesale markups $\eta_{jt} = p_{jt}^w - mc_{jt}$ (in matrix form) in (9) and (10) for the PH game. We specify the marginal costs in four parts: the manufacturer prices p_{jt}^m, state excise taxes τ_{jt} (all measured per-liter), labor and other marginal

[48] We let $\boldsymbol{\xi_t} = [\xi_t, \xi_{b(j)}, \Delta\xi_{jt} \, \forall j]$, the stacked vector of fixed effects and demand shocks for each market t.

[49] As a robustness test, we estimate a model that treats income y_i as continuous and estimates a single interaction for each element of α_i and β_i.

costs incurred by wholesalers wc_{jt}, and an unobserved cost shock ω_{jt}:

$$p_{jt}^w - \eta_{jt}\left(\mathcal{H}_t(\kappa), \theta_2\right) = mc_{jt} \equiv p_{jt}^m + \tau_{jt} + wc_{jt} + \omega_{jt}. \tag{15}$$

Normally we would specify $mc_{jt} = h(\mathrm{x}_{jt}; \theta_3) + \omega_{jt}$ and estimate the parameters of the marginal cost function θ_3 with some instruments $\mathbb{E}[\omega_{jt}|z_{jt}^s] = 0$. However, because we observe the manufacturer prices for each product p_{jt}^m, and excise taxes τ_{jt}, we know not only the determinants of the marginal cost, but the coefficients as well. This allows us to construct moments of the form:[50]

$$\mathbb{E}[\omega_{jt}] = 0, \text{ with } \omega_{jt} = \left(p_{jt}^w - p_{jt}^m - \tau_{jt} - wc_{jt}\right) - \eta_{jt}\left(\mathcal{H}_t(\kappa), \theta_2\right). \tag{16}$$

This provides additional over-identifying restrictions on the parameters in θ_2 (most importantly the price sensitivity α) by setting the observed price cost margins in the data $p_{jt}^w - p_{jt}^m - \tau_{jt}$ as close as possible to those implied by the demand model $\eta_{jt}\left(\mathcal{H}_t(\kappa), \theta_2\right)$. This is particularly helpful because we have a small number of markets (we observe statewide wholesale shipments), and spirits are becoming more popular and more expensive over time (including after the tax hike in July 2011). Absent these additional restrictions, we would estimate demand curves that are substantially less elastic, and would imply much larger wholesale markups than those we see in the data.

Some may worry about the "endogeneity of p_{jt}^m" or that manufacturers may choose prices p_{jt}^m with $\Delta\xi_{jt}$ (the demand shock) in mind. This may certainly be the case, but it is not problematic in (16) because we fix the coefficients on (p_{jt}^m, τ_{jt}) at unity, rather than estimate them.[51]

This approach is not without its drawbacks. The observed manufacturer price and excise taxes are a lower bound on the wholesaler marginal cost. An alternative might instead impose the inequality $\mathbb{E}[\omega_{jt}] \geq 0$.[52] We test sensitivity to adding a strictly positive (and fixed) wholesaling cost $wc_{jt} \in \{\$0, \$0.5, \$1, \$2\}$ per liter. Larger wholesaling costs lead to slightly more elastic demand (and worse fit), while not imposing the supply moments at all leads to less elastic demand. Our baseline specification implicitly assumes that the wholesaler incurs no additional cost to take delivery from manufacturer/distillers, store the products, and deliver them to bars, restaurants, and liquor stores.

The second drawback is that we are imposing system of first order conditions from the PH game in (4), and thus cannot test them. There is a long (and growing) literature on testing conduct in differentiated-products settings (Bresnahan, 1987; Villas-Boas, 2007; Berry and Haile,

[50]We have experimented with additional moments of the form $\mathbb{E}[\omega_{jt} \mid p_{jt}^m] = 0$ (with a polynomial series in p_{jt}^m) and $\mathbb{E}[\omega_{jt}\, \tau_{jt}] = 0$, which appear to have little impact on either parameter values or standard errors. The polynomial allows our first stage estimates to capture the fact that additive markups are larger on more expensive products as seen in Figure 2. This is less important for second-stage estimates, because the marginal components are known, and the expected Jacobian for (15) reduces from $\mathbb{E}\left[\frac{\partial \omega_{jt}}{\partial \theta_2} \mid z_{jt}^s\right]$ to $\mathbb{E}\left[\frac{\partial \eta_{jt}}{\partial \theta_2} \mid z_{jt}^s\right]$. When we follow the recipe in Conlon and Gortmaker (2020); Berry et al. (1999) we use the model itself to compute $\mathbb{E}[\eta_{jt} \mid z_{jt}]$ and its derivatives.

[51]This is an old solution to the endogeneity problem, and the basis for the Anderson and Rubin (1949) test. Concerns about a relationship between ξ_{jt} and p_{jt}^m would suggest not including p_{jt}^m in z_{jt}^D the demand-side instruments.

[52]In practice we find that this leads to $\omega_{jt} = 0$ for all (j, t). That is, the unconstrained model would set $wc_{jt} = 0$ (or a negative value if allowed).

2014; Backus et al., 2021a; Duarte et al., 2021). Testing conduct amounts to detecting violations of the supply moment(s) in (16) for different choices of markups $\eta_{jt}\left(\mathcal{H}_t(\kappa),\theta_2\right)$. The bigger challenge here would be specifying the markups in the absence of PH, particularly when multiple wholesalers offer identical products at identical prices. This would require more data than ds

5.3. Estimation Details

We use the quarterly data on shipments from manufacturers to wholesalers (as described in Section 4.2) to construct market shares. We use the legal drinking-age population estimates from the NIAAA to construct an estimate for the potential market size. As in Berry et al. (1995), the unobservable demand shock $\Delta\xi_{jt}$ is chosen to equate the observed market share, with the market share predicted by the demand model (13). The model is governed by the parameters $\theta = [\theta_1, \theta_2]$ and defined by the conditional moment restrictions $\mathbb{E}[\Delta\xi_{jt}|z^d_{jt}]$ and the scalar supply moment $\mathbb{E}[\omega_{jt}] = 0$.[53] We augment the supply and demand moments with some additional moments described below.

In July 2011, the state of Connecticut increased the volumetric tax levied on wholesalers from $\tau_{jt} = \$1.18/L$ to $\tau_{jt} = \$1.43/L$. This provides both some useful variation in τ_{jt} that serves as an instrument for p_{jt} (in z^D_{jt}), and is consistent with a long literature in public finance exploiting changes in excise tax rates to instrument for prices (Randolph, 1995; Goolsbee, 1998; Gruber and Saez, 2002). It also enables us to compute a quasi-experimental estimate of the aggregate elasticity of demand for spirits $\varepsilon^{AGG} \approx -0.41$.[54] We require that our estimated system of demand and supply match this aggregate elasticity as an additional moment of the form:

$$\mathbb{E}_t\left[\varepsilon^{AGG} + 1 - \frac{\sum_j s_{jt}(\mathbf{p}\cdot 1.01;\theta)}{\sum_j s_{jt}(\mathbf{p};\theta)}\right] = 0. \tag{17}$$

In words, we require that when we increase the prices of all products by 1%, the total sales decrease by ε^{AGG} percent. We expect this moment to be most informative about the nesting parameter ρ, which governs the extent to which consumers respond to higher prices by substituting from one vodka to another vodka or to a product from another category (including the outside good).

In addition to the moments from (aggregate) supply and demand, we augment these with moments formed from the decisions of individual panelists in the NielsenIQ data. These micro-moments (Petrin, 2002; Berry et al., 2004) are constructed by evaluating interactions of product characteristics with consumer demographics (conditional on purchase). We employ the following

[53] We've experimented with additional supply moments of the form $\mathbb{E}[\omega_{jt} \mid p^m_{jt}, \tau_{jt}] = 0$, but they seem to have little to impact on the parameter estimates.

[54] We provide more details in the appendix on our estimated aggregate elasticity. A meta-analysis in Wagenaar et al. (2009) reported the mean elasticity for the spirits category to be $\varepsilon^{AGG} \approx -0.29$.

four types of micro-moments:

$$
\begin{aligned}
&\mathbb{P}\left[\text{Income}_i \in \text{bin}_k \mid \text{Purchase}\right] \\
&\mathbb{P}\left[\text{Income}_i \in \text{bin}_k \mid 750\text{mL}\right] \\
&\mathbb{P}\left[\text{Income}_i \in \text{bin}_k \mid 1750\text{mL}\right] \\
&\mathbb{E}\left[p^w_{jt} \mid \text{Income}_i \in \text{bin}_k \text{ and Purchase}\right]
\end{aligned}
\tag{18}
$$

That is, we match the average price paid (per liter) conditional on purchase for each of the five income "quintiles." We also match the probability that the buyer of a generic liter of spirits or a 750mL/1750mL bottle falls into each income "quintile." Each of these moments is meant to be informative about a particular parameter in Π (the interactions of consumer preferences with income "quintiles"). These are not necessarily the "ideal" micro-moments from Conlon and Gortmaker (2023), but we found that they can be reliably constructed from the NielsenIQ panelist data without making additional assumptions.[55]

We estimate a separate set of moments for each year in the panelist data from 2007-2013, in part because the set of panelists (and weights) differs by year.

We also need to estimate $g(y_i)$, the distribution of household income for the state of Connecticut. This appears both in the calculation of market shares (13) and the micro-moments (18). The NielsenIQ Panelist data doesn't report exact levels of household income, but rather reports it in discrete ranges. We assign income y_i into a set of discrete "quintile" bins: $\{<$ \$25k, \$25k-\$45k \$45k-\$70k, \$70k-\$100k, $\geq$ \$100k$\}$. Because Connecticut is a high-income state, and NielsenIQ top codes income at \$100k, 29% of households are in our top "quintile" while only 8.5% are in the bottom "quintile". As one might expect, household incomes decline during the Great Financial Crisis, then rise slowly over time. We examine the possibility of including other consumer demographics in y_i. We don't see enough Black or Hispanic households purchasing spirits in Connecticut to accurately estimate micro-moments on these sub-populations. The age of the head of household doesn't seem to vary in a meaningful way with any of the product characteristics in our data, and education is highly correlated with income.[56]

Estimation takes place in `PyBLP` (Conlon and Gortmaker, 2020). We use all four sets of moments: demand $\mathbb{E}[\Delta\xi_{jt}\, z^d_{jt}] = 0$; supply $\mathbb{E}[\omega_{jt}] = 0$; aggregate elasticity ε^{AGG}; and the micro-moments. For the demand side, we need to choose instruments z^d_{jt}. Here the obvious instruments are the excluded cost variables: p^m_{jt} (the manufacturer price); and τ_{jt} (the per-liter excise tax, which

[55] One issue pointed out in Conlon and Gortmaker (2023) is that micro-moments work best when they are *compatible*. In our case, we worry that the fraction of 1.75L bottles purchased by households in the Panelist dataset is significantly higher than the fraction of 1.75L bottles (by volume) in the shipment data. Thus the marginal distribution of $\mathbb{P}(x_{jt} = 1.75L)$ is not the same across the two datasets, and we instead use a moment that conditions on the purchase of a bottle size, rather than the expectation $\mathbb{E}[x_{jt} \cdot y_i \mid \text{ purchase }]$.

[56] See Conlon et al. (2024) for an in-depth examination of interaction between household demographics and purchases of sin goods. There we find households over 55 are more likely to be heavy consumers of distilled spirits.

changes in July 2011). In addition, we follow the recipe in Gandhi and Houde (2019) and construct instruments based on differences in exogenous product characteristics $d_{jkt} = |x_{jt} - x_{kt}|$. We interact these distances with dummies for each product category; and use the local variant that constructs $z_{jt}^{GH} = \sum_k \mathbb{I}\left[d_{jk,t} < a\right]$, where a is one standard deviation of $d_{jk,t}$. In words, these instruments convey, "How many other 750mL flavored vodkas are available?" or "How many other similar proof whiskeys are for sale?" These are meant to capture the "crowding" of the product space over time.[57] An important characteristic not typically in x_{jt}, but available in our study, is the manufacturer (upstream) price p_{jt}^m. This allows us to ask: "How many other 750mL Vodkas with manufacturer prices between $35-$39 per liter are available?".

The identification argument here is somewhat different from the "classic" BLP setup in Berry et al. (1995) and Nevo (2001), which both rely primarily on "characteristics of other goods" varying across markets in the aggregate moments $\mathbb{E}[z_{jt}^D \, \Delta\xi_{jt}] = 0$. In a sense, rather than rely on cross-market variation, our identification strategy more closely follows the "identification from micro-data" argument in Berry and Haile (2022). We rely largely on our auxiliary moments: matching average markups to recover the price sensitivity $\overline{\alpha}$; the levels of the micro-moments to recover the demographic interactions Π; the cross-market variation in the micro-moments to recover the unobserved heterogeneity Σ; and the aggregate elasticity ε^{AGG} to recover the nesting parameter ρ. In Appendix C.2, we show that larger values of ρ involve more substitution to other products within the same category (ie: other vodkas) and less substitution to the outside good, and how matching the aggregate elasticity pins this down. Ideally, to estimate ρ, we would have data on the fraction of consumers whose first and second choice products were both vodka or whiskey. Unfortunately, this is not something we can construct from the Nielsen Panelist data. In Appendix C.2, we also show that mis-specification in the wholesaler marginal cost (by varying wc_{jt}) does not qualitatively impact our results.

We estimate the parameters $\theta = [\theta_1, \theta_2]$ using two-step GMM, and then use these estimated parameters to construct a feasible approximation to the optimal instruments (Chamberlain, 1987) following the recipe in Conlon and Gortmaker (2020) and Berry et al. (1999).[58] We then re-estimate the problem a second time using two-step GMM, and report those estimates for both the full model and several restricted models.

5.4. Parameter Estimates

We report our estimated parameters for the full model in Table 3. The parameters themselves are not easily interpretable, though we can see some obvious patterns. At higher income levels, consumers become less price sensitive (as one might expect), but also the intercept for all spirits

[57] In the data we see more U.S. whiskey products entering and fewer flavored-rum products.

[58] With initial estimates of $\widehat{\theta}_2$, we solve (8) for $\hat{\boldsymbol{\eta}}$ at $\mathbf{mc} = \mathbf{p^m} + \tau$ and $\boldsymbol{\xi} = 0$ in order to recover $\hat{\mathbf{p}}$ and $\hat{\mathbf{s}}(\hat{\mathbf{p}})$ and then compute the relevant components of the Jacobian for the demand moments: $\mathbb{E}\left[\frac{\partial \xi_{jt}}{\partial \theta} \mid \mathbf{z_t}, \widehat{\theta}_2\right]$ and supply moments $\mathbb{E}\left[\frac{\partial \eta_{jt}}{\partial \theta} \mid \mathbf{z_t}, \widehat{\theta}_2\right]$.

products declines. This is important as it implies that higher-income consumers don't purchase all of the alcohol, but instead purchase similar quantities at higher prices. We see less of a discernible pattern for the large format size 1750mL, other than all consumers prefer it to the other two sizes (the more ubiquitous 750mL or somewhat rare 1L sizes). Our fixed effects are at the brand level (e.g., *Smirnoff Vodka 80-Proof*), so different sizes share the same ξ_j term, but may differ in the 1750mL dummy.

The model estimates that the average markups (under the PH system) have an IQR of $(20.5\%, 25.9\%)$, with a median of 23.3% across all products and markets. In the data, the IQR is $(18.8\%, 27.6\%)$ with a median of 23.3% suggesting that these are matched quite well, though our predictions are somewhat less dispersed. Other than the common ξ_t term, we don't have any demand parameters to rationalize the high-frequency wholesale price changes in Figure 5 beyond the residual $\Delta\xi_{jt}$; instead we end up matching the average markups.

In order to rationalize these markups, our model estimates an IQR for own-price elasticities between $(-5.77, -4.70)$ with a median of -5.11. We also report product-level own-price elasticities in Figure 6 for the final period in our data. We see that less expensive products tend to have more elastic demand, particularly for the larger 1.75L bottles.[59] Additionally Figure 6 reports the outside good diversion ratio, which tells us: Conditional on leaving a product, how likely is a consumer to switch to the no-purchase option? We see that diversion to the outside good declines steeply with prices. Taken together, these suggest that when we raise prices at the lower end of the price distribution, consumers are more likely to substitute away from drinking (more than 30% of switchers), while when we raise prices at the higher end, they are more likely to switch to another product (or to simply pay more for their preferred product), with fewer than 20% of switchers choosing the outside option. This is important for the welfare results of counterfactual tax policies, because raising prices at the lower end of the price distribution will be more effective at getting consumers to substitute away from ethanol consumption (but also raise less revenue).

Our estimated own-price elasticities tend to be a little more elastic than in previous studies that do not impose the supply-side restriction (Miravete et al., 2020, 2018).[60] However, we still obtain relatively inelastic demand at the aggregate level. We estimate that a 1% increase in the price of all products would lead to a reduction in demand of $\widehat{\varepsilon}^{AGG} = -0.34$ (the targeted value was $\varepsilon^{AGG} = -0.41$). Part of the challenge is matching both the level of markups and the aggregate elasticity using a single nesting parameter $\widehat{\rho} = 0.47$. In an ideal world, we might have quasi-experimental estimates of aggregate elasticities at the category level rather than just the overall level; however, we observe only a single (uniform across categories) tax change in our data.

Perhaps the best way to validate our demand model is to examine the predicted substitution

[59]For comparison, the plain logit model imposes that more expensive products will have *more elastic* demand, providing evidence that modeling heterogeneity is important here.

[60]If we did not try to match the level of markups, we would estimate significantly *less elastic* demand. Increasing the wholesaling cost wc_{jt} leads to smaller markups and even *more elastic* demand.

patterns. For several top products, we compute the diversion ratio from that product to its closest substitutes and report both the name and diversion ratios in Table 4. For the most part, products appear to compete with similarly-priced products within the same category. For example, Dubra Vodka (1.75L), the least expensive product in our sample, appears to compete most closely with the other discount vodka brands (Popov, Sobieski, Gray's Peak, and Wolfschmidt) as well as Smirnoff vodka (a mid-range vodka and the best selling product overall). Belvedere (a super premium vodka) appears to compete with Grey Goose, Absolut, and Ketel One. Woodford Reserve, a premium bourbon, competes primarily with Maker's Mark and Jack Daniels, the two best selling American whiskeys. Because of the nesting structure, we see that Captain Morgan's competes primarily with Bacardi Rum, and Beefeater Gin competes largely with other gins (as well as Smirnoff vodka). We also see that consumers largely substitute from 1.75L bottles to 1.75L bottles or 750mL bottles to 750mL and 1L bottles.

6. Welfare Under Counterfactual Policies

Our welfare analysis focuses on what would happen if PH were replaced with a competitive market for wholesale distribution in all products, and existing state volumetric excise taxes were replaced with a single tax instrument. This is motivated by the fact that the distilled spirits market in Connecticut has a high degree of *common agency* (multiple wholesalers distributing otherwise identical products); additionally, unlike beer distribution in many states, the market for distilled spirits in Connecticut does not have *franchise laws* that restrict wholesalers to *exclusive territories*.[61] But for the PH system, this is a market that we would otherwise expect to be quite competitive.

Because our demand estimates reflect derived demand at the *wholesale level* and abstract away from retail pricing, our notion of Marshallian consumer surplus corresponds to the joint welfare of both retailers (bars, restaurants, and liquor stores) as well as households. When we report distributional analyses, these are based on both the types of products purchased by households at different income levels in retail environments (liquor stores), and the estimated price sensitivities by income group, even though some of the surplus accrues to downstream firms.

The main motivation to limit the consumption of distilled spirits is the associated negative externalities. One serious concern is that the heaviest drinkers account for the bulk of the external damage (Griffith et al., 2019; Conlon et al., 2024). Because our analysis focuses largely on the wholesale tier, we are limited in our ability to model *who* does the drinking. Instead we treat the externality as if it were *atmospheric* (i.e., it depends on only the aggregate level of ethanol consumption). This would be problematic if we were concerned that there were larger negative externalities associated with drinking tequila rather than vodka, or that lost productivity was

[61] While most large products are sold by multiple wholesale firms who could compete until $P = MC$, several products have a single wholesale distributor. Even though that distributor might have market power in the distribution of that product, our counterfactuals treat distribution as perfectly competitive. The assumption is that a manufacturer would have an incentive to seek a second wholesale distributor in the counterfactual world, even though they don't in the PH world.

greater for households earning over \$100K. Rather than take a stand on the externality, we consider three policy targets: (a) keeping ethanol consumption fixed at the existing level; (b) increasing ethanol consumption by 10%; (c) reducing ethanol consumption by 10%. In our final exercise, we ask: How much can we reduce ethanol consumption without reducing consumer surplus?

6.1. Counterfactual Tax Instruments

We describe several simple tax instruments that we consider as alternatives to the PH system in Table 5: (a) a volumetric tax (similar to the one Connecticut and most license states use currently); (b) an ethanol specific tax (similar to the one used by the federal government); (c) an *ad-valorem* tax (similar to the general sales tax); (d) a price floor per unit of ethanol (similar to that enacted in Scotland and examined by Griffith et al. (2022)); (e) a product specific (Ramsey) tax that maximizes consumer surplus subject to either a revenue or aggregate ethanol constraint; and (f) a profit-maximizing monopoly (similar to the market structure in Maine). Under each of our policy alternatives, we do not change the baseline federal excise tax, which we treat as part of p_{jt}^m, but replace the existing state volumetric tax with the policy alternative.

To understand how these tax instruments differ, we choose the amount of each tax so that the level of aggregate ethanol consumption remains the same (after replacing PH with a perfectly competitive wholesale tier). We then examine how the counterfactual prices compare to those observed under the PH system in Figure 7. One advantage of our approach is that we observe manufacturer prices, and do not need to infer them from a first order condition like (9). Product prices that lie below the black 45-degree line denote products that are less expensive under the alternative policy than under PH, while prices above the line are more expensive under the alternative policy. The most obvious difference is that very low marginal cost products are generally more expensive under the alternative policies. Profit-maximizing wholesalers tend to set very low markups on the least expensive and most elastically-demanded products (recall Figure 4). These are also the products where there is the most substitution to the outside good (Figure 6), and what makes PH ineffective as a tool to discourage ethanol consumption.

Beyond this are several notable patterns. First, minimum prices reduce aggregate ethanol consumption by raising the price of 80-proof products to roughly \$17 per liter (and by definition lower all other prices since these products now sell at $p_{jt}^w = p_{jt}^m$). Second, there is little difference between taxing volume and taxing ethanol content, since the bulk of products are around 80 proof (40% alcohol by volume). These taxes effectively add a fixed $\tau = \$5.45$ (per liter) or $\tau = \$13.62$ (per liter of ethanol) to each product, which leads to higher prices at the low end of the market, and lower prices at the high end of the market (because PH markups are generally increasing in marginal costs). Third, sales-tax instruments, whether in combination with the existing volumetric tax or otherwise, generally lead to higher prices for much of the quality distribution. In part this stems from the fact that the marginal costs (manufacturer prices p_{jt}^m) at the low end of the distribution are quite low – and raising those prices sufficiently requires very high sales-tax rates (78% without existing excise taxes, and 44.6% if we don't eliminate existing excise taxes), which exceed the typical markups under PH (around 23%). Because all price distributions yield the same amount of overall ethanol consumption, it should be clear that the volumetric and ethanol taxes (along with minimum unit prices) are likely to yield significant benefits to consumers (particularly those who prefer high-end products).

6.2. Welfare Results

As described above, governments tax alcohol with dual objectives: to curb alcohol consumption and to raise revenue for the state. We map out the welfare trade-offs for each tax instrument for a variety of different tax rates in Figure 8 and compute the percentage changes relative to the status quo policy (PH). In the left panel of Figure 8, we consider the trade-off between consumer surplus and tax revenue, while in the right panel we consider the trade-off between consumer surplus and ethanol consumption (external damage). For each tax instrument we denote the point on the curve that: leaves ethanol consumption unchanged ($*$), increases ethanol consumption by 10% ($+$), and decreases ethanol consumption by 10% ($\times$).

The frontier in the left panel of Figure 8 is defined by the Ramsey-Revenue scenario, which uses product specific taxes to maximize consumer surplus at each level of tax revenue. This traces out a curve from the perfectly competitive price (with no additional taxes) in which consumer surplus increases by 56% (and ethanol consumption by 88%); to the monopoly price, which achieves the highest possible revenue increase of 440% (but reduces consumer surplus and ethanol consumption by 36%).[62] The uniform sales tax gets surprisingly close to this frontier (both with and without the existing volumetric taxes). This would make sense if upstream (manufacturer/distiller) prices reflect heterogeneous elasticities. As in Figure 7, taxing volume or taxing ethanol content yields nearly identical results, but would be less effective at raising revenue than sales taxes, and are inside our frontier. Likewise, the Ramsey-Ethanol problem (which chooses product-specific taxes to maximize consumer surplus at each level of ethanol consumption, without concern for revenue) and the minimum ethanol unit price are highly similar to each other – but relatively ineffective at raising tax revenue, because they both leave products with manufacturer prices over \$20/L effectively untaxed. We denote the existing PH policy at the origin of the graph, where the only source of revenue is the \$1.43/L volumetric tax. This policy is not only far from the frontier, but also dominated by all of the alternative tax instruments. One major limitation of the PH policy is that revenue is captured as wholesaler profits rather than as tax revenue. For this reason, we also report $PH + PS$, which includes wholesaler profits in the tax revenue calculation. One interpretation is that the state of Connecticut could extract all wholesaler profits via a lump-sum tax or auction. Even in this extreme scenario, we can see that simple tax instruments (sales taxes,

[62] The state of Maine auctioned a 10-year lease to the monopoly wholesaler in 2014, but the auction generated substantially less revenue than expected.

volumetric taxes, or ethanol taxes) can be designed to raise more revenue while simultaneously increasing consumer surplus.

The right panel of Figure 8 considers the trade-off between overall ethanol consumption (the source of the negative externality) and consumer surplus. Here the frontier is defined by Ramsey-Ethanol, which maximizes consumer surplus at each level of ethanol consumption by setting product-specific tax rates. As was the case in Figure 7, the minimum ethanol unit price is now remarkably close to the frontier. The existing PH system is dominated by simple taxes on volume or ethanol content, which allow for higher levels of consumer surplus at each level of ethanol consumption. In the same vein, the existing PH system does generate higher levels of consumer surplus for each level of ethanol consumption than the sales taxes. This is because raising prices at the low end of the distribution enough to discourage consumption requires extremely high sales-tax rates, leading to even higher prices at the high end of the distribution (as seen in Figure 7). This distorts inframarginal decisions.

Many of our counterfactual policies change the relative prices of products when compared to the existing PH system; taxes on volume, ethanol content, or minimum ethanol unit prices all increase the prices at the low end of the distribution and reduce prices at the high end, while sales taxes do the opposite. One concern is that changing relative prices might have distributional consequences for consumers. Increasing the prices of the least expensive products and reducing the prices of the most expensive products may increase overall consumer surplus but still harm the most price-sensitive (and lowest-income) consumers. In order to investigate this possibility, we break out the change in consumer surplus for each of our five income bins in Table 6. In the first panel, we report the welfare implications of alternative tax policies that hold ethanol consumption fixed. Here we see that taxes on ethanol content or volume not only increase tax revenues by over 270% and increase overall consumer surplus by over 6.4%; they also increase consumer surplus for consumers at all income levels. It is true, though, that households earning $70k or more see a larger increase in consumer surplus (7.4% or more) than those in the lowest income bin (4.2% or less). In the second panel, we consider counterfactual taxes that reduce the overall level of ethanol consumption by 10%. Now volumetric taxes reduce consumer surplus for each income group. Taxes based on ethanol content increase aggregate consumer surplus slightly (by 0.6%) by increasing the consumer surplus to the highest income bin (by 2.9%) and reducing it by as much as 5.5% for the lowest income bin (recall that nearly 30% of households in Connecticut earn more than $100k per year). However, tax revenues increase by over 300% when compared to the existing PH system.

One approach might be to transfer some of the additional tax revenue in order to hold harmless the lowest-income groups (such as by reducing taxes on wage income or expanding the EITC) while still reducing the overall level of ethanol consumption in Connecticut. However, such transfers may be complicated by political considerations. Instead, we ask: how much can we reduce aggregate ethanol consumption (a) without reducing aggregate consumer surplus, and (b) without reducing the consumer surplus of any income group? We focus on a volumetric tax, not because it is "optimal" in any sense, but because it performs relatively well in Figure 8 and is what most license states (including Connecticut) currently employ. We report these results in Table 7 in the columns labeled "Base." We find that we can reduce overall ethanol consumption by as much as 9.3% without reducing aggregate consumer surplus and while increasing tax revenue by 306%. However, the consumer surplus of all income groups declines except for households earning over $100k per year, because the volumetric tax increases the prices at the low end of the distribution by levying a $6.45/L tax on distilled spirits (compared to the existing $1.43/L tax). The $6.45/L tax is still below the PH markup for premium products, so that the highest-income households are left better off.

The third panel of Table 7 shows that if we replace the PH system with a $5.78/L volumetric tax, we can reduce overall ethanol consumption by 3.2%, and increase aggregate consumer surplus by 4.23%, without reducing the consumer surplus of the lowest-income consumers (and meanwhile increasing the consumer surplus of the high-income consumers by 5.7% or more). Such a policy would leave nearly everyone (except wholesalers) better off, as even the profits of the upstream manufacturers/distillers would increase by nearly 15% as consumers substitute to premium brands (with higher manufacturer margins). We should caution that this need not constitute a "Pareto Improvement" because we are reporting only the aggregate consumer surplus for each income group, which is not the same as the consumer surplus for each individual. For example, if individuals have idiosyncratically high preferences for inexpensive plastic-bottle vodkas, they may still be left worse off by policies that make their favorite products more expensive.

6.3. Robustness to Key Assumptions

Our baseline counterfactuals assume that the post-and-hold system can be replaced with a perfectly competitive wholesale tier that sets $p_{jt}^w = p_{jt}^m + wc_{jt} + \tau_{jt}$. We implicitly assume it is costless to operate the wholesale tier and set $wc_{jt} = 0$ from (15). As an alternative, we set $wc_{jt} = \$1/\text{L}$ and include these welfare results as an additional set of columns in Table 7.[63] Because these additional per-liter wholesaling costs are isomorphic to volumetric taxes, the resulting prices and quantities (and thus welfare) are unchanged. The only effect is that it leads to a one-for-one reduction in the tax revenue collected. We repeat our full analyses (including for other taxes) at a variety of different levels of wc_{jt} in Appendix C.

Additionally, our counterfactuals have held upstream (manufacturer/distiller) prices fixed. This makes sense if distillers/manufacturers are pricing regionally and can't price discriminate across wholesalers in Connecticut and those in Massachusetts, New York, or New Jersey. In the second scenario, we allow for multi-product distillers/manufacturers (e.g. Bacardi, Diageo) to adjust prices. This scenario requires estimates not only of *manufacturer prices*, which we observe, but also of

[63]Recall from Table 1 that the average wholesale price-cost margin is around $3/L. If we increase wc_{jt} too much, some wholesalers start to lose money on every transaction.

manufacturer marginal costs which we do not.[64] However, we are able to back those $\mathbf{c_m}$ out of the manufacturer first-order condition:[65]

$$\mathbf{p^m} - \mathbf{c_m} = \left[\mathcal{H}_m \odot \left(\frac{\partial\, \mathbf{p^w}}{\partial\, \mathbf{p^m}}(\mathbf{p^w}, \kappa)\right)^T \Delta(\mathbf{p^w})\right]^{-1} \mathbf{s}(\mathbf{p^w}). \tag{19}$$

This requires knowledge of the manufacturer ownership matrix $\mathcal{H}_m$ (which we observe) and the manufacturer-wholesale price pass-through matrix, which we can estimate from the demand system.[66] In the counterfactual world, with competitive wholesaling, the pass-through matrix reduces to the identity matrix plus any ad valorem taxes $\frac{\partial\, \mathbf{p}}{\partial\, \mathbf{m}} = I_J \cdot (1 + \tau_r)$, while the effective marginal cost becomes the production cost $\mathbf{c_m} + \tau$, where τ are any per-unit taxes.

Once we recover $\mathbf{c_m}$, we can re-solve (19) for the optimal manufacturer prices $\mathbf{p^m}(\mathbf{c_m}+\tau)$ at each proposed level of taxes. This allows us to add a third set of columns to Table 7, which incorporates an endogenous response by the manufacturers/distillers. For each of the three scenarios, this reduces the corresponding tax rate by around 40-60 cents per liter. It has only a limited effect on how much we can reduce ethanol consumption without reducing consumer surplus (either in aggregate or for all income groups), and minor distributional consequences. Manufacturers capture some of the additional surplus, largely by slightly increasing markups on certain premium products, but this acts largely as a transfer rather than significantly changing the resulting equilibrium. The main effect is that it modestly reduces the amount of additional tax revenue collected (from an increase of 307% to 270%, holding aggregate CS fixed) and substantially enhances manufacturer profits from an increase of 9.6% to an increase of 28.6%. Unlike the PH wholesaling system which significantly distorts the relative prices of products, because manufacturers already have market power in setting upstream prices, allowing these manufacturers to re-optimize after eliminating PH functions mostly like a transfer from tax revenue to manufacturer profits.

6.4. Why does PH perform so poorly?

In theory, one might have expected the post-and-hold system to perform better. Firms with market power have the ability to choose prices more flexibly than simple tax instruments such as volumetric or sales taxes. Moreover, they can (and do) choose prices with knowledge of own- and cross-price elasticities. Indeed, it has been known since Ramsey (1927) that there is a duality between the optimal tax problem and the monopoly problem, and that the monopolist minimizes deadweight loss for a particular level of revenue. This raises the question, how does the PH problem solved by

[64] We use our estimated manufacturer costs in Table 1 and Table 2.

[65] See Appendix B.1 or the Appendix to Miller and Weinberg (2017) for a derivation. With some additional modifications, we could follow the latter and attempt to parameterize the problem in (19) and interpolate between no manufacturer response and the fully flexible manufacturer response. We find that the two outcomes are not far enough apart for this to matter.

[66] Our average product-level own pass-through rate is 1.3 which is *overshifted*, but consistent with reduced form estimates in our prior work Conlon and Rao (2020).

wholesale firms differ from the social planner's problem?

In Appendix A.3, we solve the constrained optimization problem of a social planner who chooses prices to maximize social surplus subject to a minimum value of revenue (with Lagrange multiplier λ_r) and a maximum value of ethanol consumption (with Lagrange multiplier λ_e). To simplify things we assume the externality is atmospheric with the ethanol content of each product given by e_j. This produces the following FOC:

$$p_j = \frac{|\epsilon_{jj}|}{|\epsilon_{jj}| - \frac{\lambda_r}{1+\lambda_r}} \left[mc_j + \frac{\lambda_e}{1+\lambda_r} e_j + \sum_{k \neq j} D_{jk} \left(p_k - mc_k - \frac{\lambda_e}{1+\lambda_r} e_k \right) \right]. \tag{20}$$

The first term functions like an inverse elasticity markup rule, where $\frac{\lambda_r}{1+\lambda_r} = \theta$ behaves like a conduct parameter with $\theta = 0$ corresponding to the perfectly competitive solution and $\theta = 1$ corresponding to the monopoly solution. We compare this to the solution to the PH problem from (7):

$$p_j = \frac{|\epsilon_{jj}|}{|\epsilon_{jj}| - 1} \cdot \left[mc_j + \sum_{k \in \mathcal{J}_f \setminus \{j\}} \kappa_{jk} \cdot D_{jk} \cdot (p_k - mc_k) \right].$$

The main difference is that the PH first-order conditions effectively set $\lambda_e = 0$ and $\lambda_r \to \infty$ as in the monopoly problem. An additional wedge arises because the *opportunity cost* (shown in brackets) depends on the term $\kappa_{jk} = \gamma_k^f / \gamma_j^f$, which measures the relative market shares of k and j for the pivotal seller of j. The Ramsey solution would set $\kappa_{jk} = 1$ for all (j, k), whereas the PH solution will set $\gamma_k^f = 0$ for products not distributed by the pivotal seller of j, and may set $\gamma_k^f / \gamma_j^f > 1$ for others. In practice, the former tends to dominate, so that the PH solution tends to understate the effective diversion ratios $\kappa_{jk} \cdot D_{jk}$ instead of D_{jk}.[67] This is particularly true for premium products, for which diversion to other brands is larger (because diversion to the outside good D_{j0} is smaller, as in Figure 6). The result is that PH applies a larger markup ($\theta = 1$ or $\lambda_r \to \infty$) to a smaller marginal cost than the planner, which distorts not only the price levels, but the relative prices as well.

7. Discussion

We show that the post-and-hold system employed by Connecticut is not effective at discouraging consumption of ethanol or raising tax revenues when compared to simple, commonly-used tax instruments. Indeed, it is possible to reduce overall ethanol consumption (and associated externalities) by more than 9%, without reducing consumer surplus, and while increasing tax revenues by

[67] Because not all wholesalers distribute all products $\mathcal{J}_f \subset \mathcal{J}$.

over 300% (or around $180 million per year).

Our results shed additional light on previous studies of alcoholic beverages because we are able to trace out a wide range of policy instruments over a variety of different values. As an example, we show that the minimum ethanol unit price adopted by Scotland (and analyzed by Griffith et al. (2022)) is very similar to the solution of a social planner who wishes to maximize consumer surplus subject to an upper bound on aggregate ethanol consumption. While this policy is effective at limiting consumption, it is ineffective at raising tax revenues, which perhaps explains why it has not been more widely adopted. Likewise, we show that a uniform sales-tax rate does a relatively good job approximating the the problem of a social planner who maximizes consumer surplus subject to a revenue constraint. However, while the uniform sales tax is able to generate similar levels of consumer surplus and tax revenue as the "Ramsey" planner, it does so at significantly higher levels of ethanol consumption (and hence negative externalities). This helps to reconcile our results with prior studies of uniform markup rules (which operate like sales or *ad valorem* taxes) set by the state-run monopolist in Pennsylvania in Miravete et al. (2018, 2020). Where we do depart from those studies is that we find that lower-end products have more elastic demand and *lower markups* under the PH system, and are more substitutable to the outside option, whereas the studies examining the Pennsylvania state monopolist find that lower-end products have *less elastic demand* and *higher markups*. There are multiple ways to explain this seeming discrepancy, including that the relative prices in Connecticut and Pennsylvania are quite different exactly because profit-maximizing firms do not set a uniform markup.

Our findings are driven by our unusually complete data. Our ability to combine wholesaler prices with upstream (manufacturer/distiller) input prices allows us to measure the wholesale markups of profit-maximizing firms, showing that they generally increase with input price (see Figure 4). Matching these markups – along with micro-moments that reveal lower-income consumers generally pay lower prices but tend to consume somewhat less rather than more alcohol than high-income consumers[68] – yields demand estimates that directly inform our counterfactual policies. Our estimates indicate that the least expensive products tend to have more elastic demand and are more substitutable to the outside option (see Figure 6). By raising the prices of these products, and reducing the prices of premium products, we are able to undo the distortion in relative prices caused by the PH system, increase consumer surplus, and decrease ethanol consumption.

The seemingly "free lunch" arises because firms with market power may face substantially different incentives than a social planner. When products are differentiated, relying on firms with market power to provide "second-best" regulation of externalities may be far from optimal. In our context, consumers care about product quality, and firms with market power set the effective "tax" on product quality too high and the effective "tax" on externalities too low, and significantly distort the choices of inframarginal consumers.

[68]We confirm this pattern in our other work (Conlon et al., 2024)

These results should serve as a cautionary tale to policymakers who wish to outsource the mitigation of negative externalities to private firms. They can also be applied to the broader context beyond distilled spirits. As states have legalized other sin goods such as marijuana, they have limited competition by placing significant restrictions on entry (Thomas, 2019) – or levied *ad valorem* taxes at different parts of the supply chain (Hansen et al., 2022) that may not perform as well as Pigouvian taxes (Hansen et al., 2020) in addressing negative externalities while generating tax revenue. Restricting competition, particularly when products are differentiated, may not perform as well as policymakers hope.

References

ALLCOTT, H., B. B. LOCKWOOD, AND D. TAUBINSKY (2019): "Regressive Sin Taxes, with an Application to the Optimal Soda Tax*," *The Quarterly Journal of Economics*, 134, 1557–1626.

ANDERSON, T. W. AND H. RUBIN (1949): "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *The Annals of Mathematical Statistics*, 20, 46–63.

ASKER, J. (2016): "Diagnosing Foreclosure due to Exclusive Dealing," *The Journal of Industrial Economics*, 64, 375–410.

BACKUS, M., C. CONLON, AND M. SINKINSON (2021a): "Common Ownership and Competition in the Ready-to-Eat Cereal Industry," Working Paper 28350, National Bureau of Economic Research.

——— (2021b): "Common Ownership in America: 1980–2017," *American Economic Journal: Microeconomics*, 13, 273–308.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

——— (1999): "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy," *American Economic Review*, 89, 400–430.

——— (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 112, 68–105.

BERRY, S. T. AND P. A. HAILE (2014): "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 82, 1749–1797.

——— (2022): "Nonparametric Identification of Differentiated Products Demand Using Micro Data," Tech. Rep. arXiv:2204.06637, arXiv.

BRENKERS, R. AND F. VERBOVEN (2006): "Liberalizing A Distribution System: The European Car Market," *Journal of the European Economic Association, MIT Press*, 4, 216–251.

BRESNAHAN, T. F. (1987): "Competition and Collusion in the American Automobile Industry: The 1955 Price War," *The Journal of Industrial Economics*, 35, 457–482.

CHAMBERLAIN, G. (1987): "Asymptotic efficiency in estimation with conditional moment restrictions," *Journal of Econometrics*, 34, 305–334.

CHEN, Y. AND M. H. RIORDAN (2008): "Price-Increasing Competition," *The RAND Journal of Economics*, 39, 1042–1058.

CONLON, C. AND J. GORTMAKER (2020): "Best practices for differentiated products demand estimation with PyBLP," *The RAND Journal of Economics*, 51, 1108–1161.

——— (2023): "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP," .

Conlon, C., N. L. Rao, and Y. Wang (2024): "Who Pays Sin Taxes? Understanding the Overlapping Burdens of Corrective Taxes," *Review of Economics and Statistics.*

Conlon, C. T. and N. L. Rao (2020): "Discrete Prices and the Incidence and Efficiency of Excise Taxes," *American Economic Journal: Economic Policy*, 12, 111–143.

Cooper, J. C. and J. D. Wright (2012): "Alcohol, Antitrust, and the 21st Amendment: An Empirical Examination of Post and Hold Laws," *International Review of Law and Economics*, 32, 379–392.

Dixit, A. K. (1985): "Handbook of Public Economics," New York: North-Holland, vol. 1, 313–374, section: Tax Policy in Open Economies.

Duarte, M., L. Magnolfi, S. Mikkel, and C. Sullivan (2021): "Testing Firm Conduct," Working Paper.

Dubois, P., R. Griffith, and M. O'Connell (2020): "How Well Targeted Are Soda Taxes?" *American Economic Review*, 110, 3661–3704.

Gandhi, A. and J.-F. Houde (2019): "Measuring Substitution Patterns in Differentiated-Products Industries," Working Paper 26375, National Bureau of Economic Research.

Goolsbee, A. (1998): "Investment Tax Incentives, Prices, and the Supply of Capital Goods," *The Quarterly Journal of Economics,*, 113, 121–148.

Greenfield, T. K., Y. Ye, 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*, 6(12), 3205–24.

Griffith, R., M. O'Connell, and K. Smith (2019): "Tax design in the alcohol market," *Journal of Public Economics*, 172, 20 – 35.

Griffith, R., M. O'Connell, and K. Smith (2022): "Price Floors and Externality Correction," *The Economic Journal*, 132, 2273–2289.

Grigolon, L. and F. Verboven (2013): "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation," *Review of Economics and Statistics.*

Gruber, J. and E. Saez (2002): "The elasticity of taxable income: evidence and implications," *Journal of Public Economics*, 84, 1–32.

Hansen, B., K. Miller, B. Seo, and C. Weber (2020): "Taxing the potency of sin goods: evidence from recreational cannabis and liquor markets," *National Tax Journal*, 73, 511–544.

Hansen, B., K. Miller, and C. Weber (2022): "Vertical integration and production inefficiency in the presence of a gross receipts tax," *Journal of Public Economics*, 212, 104693.

Hollenbeck, B. and R. Z. Giroldo (2021): "Winning Big: Scale and Success in Retail Entrepreneurship," *Marketing Science.*

Hollenbeck, B. and K. Uetake (2021): "Taxation and market power in the legal marijuana industry," *The RAND Journal of Economics*, 52, 559–595.

Illanes, G. and S. Moshary (2020): "Market Structure and Product Assortment: Evidence from a Natural Experiment in Liquor Licensure," Working Paper 27016, National Bureau of Economic Research.

Jaffe, S. and E. G. Weyl (2013): "The First-Order Approach to Merger Analysis," *American Economic Journal: Microeconomics*, 5, 188–218, publisher: American Economic Association.

Kingston, S. (2011): *Greening EU Competition Law and Policy*, Antitrust and Competition Law, Cambridge: Cambridge University Press.

Kopczuk, W. (2003): "A note on optimal taxation in the presence of externalities," *Economics Letters*, 80, 81–86, publisher: Oxford University Press.

Levy, D. T., L. M. Sánchez-Romero, C. E. Douglas, and D. T. Sweanor (2021): "An Analysis of the Altria-Juul Labs Deal: Antitrust and Population Health Implications," *Journal of Competition Law & Economics*, 17, 458–492.

Linklaters (2020): "92 percent of businesses call for changes to competition rules to boost climate change collaboration," .

Miller, N. H., G. Sheu, and M. C. Weinberg (2021): "Oligopolistic Price Leadership and Mergers: The United States Beer Industry," *American Economic Review*, 111, 3123–3159.

Miller, N. H. and M. C. Weinberg (2017): "Understanding the Price Effects of the MillerCoors Joint Venture," *Econometrica*, 85, 1763–1791.

Miravete, E. J., K. Seim, and J. Thurk (2018): "Market Power and the Laffer Curve," *Econometrica*, 86, 1651–1687.

——— (2020): "One Markup to Rule Them All: Taxation by Liquor Pricing Regulation," *American Economic Journal: Microeconomics*, 12, 1–41.

Nevo, A. (2000): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *The RAND Journal of Economics*, 31, 395–421.

——— (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 307–342.

O'Connell, M. and K. Smith (2022): "Optimal Sin Taxation and Market Power," Working Paper.

Oum, T. H. and M. W. Tretheway (1988): "Ramsey Pricing in the Presence of Externality Costs," *Journal of Transport Economics and Policy*, 22, 307–317.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110, 705–729, publisher: The University of Chicago Press.

Ramsey, F. P. (1927): "A Contribution to the Theory of Taxation," *The Economic Journal*, 37, 47–61.

Randolph, W. C. (1995): "Dynamic Income, Progressive Taxes, and the Timing of Charitable Contributions," *Journal of Political Economy*, 103, 709–738.

Saffer, H. and M. Gehrsitz (2016): "The effect of post-and-hold laws on alcohol consumption," *International Review of Law and Economics*, 45, 63–71.

Sandmo, A. (1975): "Optimal Taxation in the Presence of Externalities," *The Swedish Journal of Economics*, 77, 86–98.

Sass, T. (2005): "The competitive effects of exclusive dealing: Evidence from the U.S. beer industry," *International Journal of Industrial Organization*, 23, 203–225.

Sass, T. and D. S. Saurman (1993): "Mandated Exclusive Territories and Economic Efficiency: An Empirical Analysis of the Malt-Beverage Industry," *Journal of Law and Economics*, 36, 153–77.

Schinkel, M. P. and L. Treuren (2020): "Green Antitrust: (More) Friendly Fire in the Fight against Climate Change," SSRN Scholarly Paper ID 3749147, Social Science Research Network, Rochester, NY.

Seim, K. and J. Waldfogel (2013): "Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board's Entry Decisions," *American Economic Review*, 103, 831–862.

Seo, B. (2019): "Firm Scope and the Value of One-Stop Shopping in Washington State's Deregulated Liquor Market," SSRN Scholarly Paper ID 2863094, Social Science Research Network, Rochester, NY.

Slater, M. E. and H. R. Alpert (2022): "Publications | National Institute on Alcohol Abuse and Alcoholism | Surveillance Report #119," .

Thomas, D. (2019): "License Quotas and the Inefficient Regulation of Sin Goods: Evidence From the Washington Recreational Marijuana Market," SSRN Scholarly Paper ID 3312960, Social Science Research Network, Rochester, NY.

US Department of Treasury (2022): "Competition in the Markets for Beer, Wine, and Spirits," Tech. rep.

Villas-Boas, S. B. (2007): "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data," *The Review of Economic Studies*, 74, 625–652, publisher: [Oxford University Press, Review of Economic Studies, Ltd.].

Wagenaar, A. C., M. D. Livingston, and S. S. Staras (2015): "Effects of a 2009 Illinois Alcohol Tax Increase on Fatal Motor Vehicle Crashes," *American Journal of Public Health*, 105, 1880–1885.

Wagenaar, A. C., M. J. Salois, and K. A. Komro (2009): "Effects of beverage alcohol price and tax levels on drinking: a meta-analysis of 1003 estimates from 112 studies," *Addiction (Abingdon, England)*, 104, 179–190.

Table 1: Summary Statistics: Wholesale and Manufacturer Price Connecticut Q3 2007 - Q2 2013

	# Obs	Share	Proof	% Flavored	Manufacturer Price	Manufacturer Margin	Wholesaler Price	Wholesaler Margin	Retailer Price	Retailer Margin
Gin	59	7.4	87.07	0.02	11.15	2.99	16.21	3.79	18.72	2.34
Rum	147	17.5	73.63	0.21	10.17	2.54	15.08	3.65	17.60	2.52
Tequila	92	4.9	80.04	0.00	15.17	4.14	22.05	5.60	28.51	4.70
Vodka	208	44.8	79.19	0.15	10.73	2.68	15.42	3.42	18.05	2.54
NA Whiskey	127	15.2	81.80	0.00	11.59	3.11	17.41	4.54	20.08	2.76
UK Whiskey	102	10.2	80.79	0.00	18.36	4.31	25.04	5.41	28.15	3.12
750mL	310	20.1	79.05	0.18	16.44	4.17	23.57	5.85	28.32	4.74
1L	174	23.2	79.32	0.12	13.80	3.71	19.92	4.85	24.85	4.35
1.75L	251	56.7	79.55	0.08	9.32	2.26	13.53	2.94	14.91	1.36
All	735	100.0	79.40	0.11	11.79	2.98	17.03	3.97	19.82	2.71

Note: The table above describes manufacturer, wholesale and retail prices and margins for 735 of 1,365 products (used in our estimation procedure) by category and size. The number of products corresponds to brand-size combinations, such as Smirnoff Vodka-750mL or Tanqueray Gin-1L. All averages are weighted by total liters sold. *Share* describes the share of total liters sold. The average *Proof* and percentage *Flavored* is reported. The average prices and margins are reported on a *per liter* basis.
The *Manufacturer Margin* is the difference between the manufacturer price the estimated manufacturer marginal cost from the demand and supply model (net of federal excise taxes). All other columns in this table are observed rather than estimated.
Retail Margin is the difference between the retail price and the wholesale price.
Wholesale Margin is the difference between the wholesale and manufacturer price plus state excise tax.
Federal alcohol excise taxes of $2.85 per liter of 80 proof spirits are levied on manufacturers. Connecticut state alcohol taxes, which are remitted by wholesalers, were raised from $1.18 to $1.42 per liter regardless of proof in July 2011.
Source: Harmonized Price and Quantity Data (top 750 products, average price under $60 per liter).

Table 2: Manufacturer Summary

	# Obs	Share	750mL	1L	1.75L	Manufacturer Price	Manufacturer Lerner	Wholesaler Price	Wholesaler Lerner	Retailer Price	Retailer Lerner
Diageo	155	32.7	0.16	0.21	0.63	11.75	0.30	17.00	0.23	19.26	0.11
Bacardi	48	14.2	0.21	0.34	0.45	14.30	0.23	20.03	0.23	22.63	0.11
Pernod	68	14.2	0.20	0.33	0.47	15.03	0.24	20.74	0.21	23.96	0.13
Jim Beam	102	8.3	0.18	0.23	0.59	9.59	0.26	14.55	0.24	17.56	0.14
Brown Forman	32	5.2	0.23	0.30	0.47	14.83	0.28	22.49	0.28	26.01	0.13
Skyy	26	2.9	0.27	0.06	0.67	11.18	0.20	16.00	0.21	18.58	0.13
Constellation Brands	6	2.8	0.19	0.11	0.71	7.43	0.26	12.45	0.29	14.37	0.13
Constellation	24	2.1	0.05	0.12	0.83	4.91	0.25	8.09	0.22	9.72	0.14
Star Industries	16	2.1	0.13	0.29	0.58	4.67	0.26	7.88	0.24	9.53	0.17
Imperial	6	2.1	0.19	0.10	0.71	5.72	0.26	9.16	0.23	12.16	0.24
MHW	44	2.0	0.41	0.16	0.43	11.68	0.22	16.97	0.23	20.77	0.17
Black Prince	7	2.0	0.10	0.29	0.62	3.97	0.28	5.93	0.11	7.15	0.17
Heaven Hill	21	1.5	0.18	0.05	0.77	6.65	0.21	10.12	0.20	12.05	0.15
White Rock	8	1.3	0.24	0.00	0.76	7.04	0.21	10.53	0.21	13.48	0.21
William Grant	17	1.3	0.22	0.12	0.65	10.40	0.24	16.02	0.25	18.62	0.11
Other	36	1.0	0.42	0.16	0.42	9.57	0.22	13.86	0.21	18.34	0.23
Remy-Cointreau	16	1.0	0.35	0.13	0.52	18.09	0.20	24.74	0.20	28.94	0.15
US Distributors	6	0.8	0.23	0.00	0.77	7.02	0.20	9.47	0.11	14.52	0.33
Sazerac	20	0.7	0.34	0.24	0.42	9.92	0.24	14.52	0.20	18.69	0.22
Moet Hennessy	10	0.6	0.41	0.37	0.22	24.35	0.21	31.02	0.17	37.43	0.17
LuxCo	19	0.6	0.17	0.38	0.45	7.13	0.24	10.97	0.23	14.60	0.23
MS Walker	10	0.2	0.08	0.48	0.45	5.33	0.22	7.32	0.09	10.99	0.25
McCormick	7	0.2	0.07	0.56	0.37	5.09	0.28	7.59	0.17	11.96	0.23
Proximo	3	0.1	1.00	0.00	0.00	20.02	0.24	29.27	0.26	37.64	0.21
Duggans	2	0.1	0.00	0.26	0.74	8.00	0.20	12.93	0.28	15.30	0.15
Infinium	1	0.0	0.00	0.00	1.00	5.54	0.22	8.84	0.24	10.78	0.17
Castle Brands	1	0.0	1.00	0.00	0.00	11.86	0.21	16.79	0.21	22.03	0.23

Note: The table above reports product shares, average prices and Lerner markups by manufacturer for 735 of 1,365 products (used in our estimation procedure). The number of products corresponds to brand-size combinations, such as Smirnoff Vodka-750mL or Tanqueray Gin-1L. Average prices and Lerner markups are reported on a *per liter* basis. All averages are weighted by total liters sold.

Share describes the share of total liters sold by each manufacturer.

Manufacturer Markup is the difference between the manufacturer price the estimated manufacturer marginal cost from the demand and supply model (net of federal excise taxes) scaled by the estimated manufacturer marginal cost. All other columns in this table are observed rather than estimated.

Retail Lerner is the difference between the retail price and the wholesale price scaled by the retail price.

Wholesale Lerner is the difference between the wholesale and manufacturer price plus state excise tax scaled by the wholesale price.

Federal alcohol excise taxes of $2.85 per liter of 80-proof spirits are levied on manufacturers. Connecticut state alcohol taxes, which are remitted by wholesalers, were raised from $1.18 to $1.42 per liter regardless of proof in July 2011.

Source: Harmonized Price and Quantity Data (top 750 products, average under $60 per liter).

Table 3: Parameter Estimates: Full Model

Π	Const	Price	1750mL
Below \$25k	2.928	-0.260	0.543
	(0.233)	(0.056)	(0.075)
\$25k-\$45k	0.184	-0.170	0.536
	(0.236)	(0.054)	(0.083)
\$45k-\$70k	0.000	-0.179	0.980
	(0.000)	(0.053)	(0.093)
\$70k-\$100k	-0.452	-0.496	0.608
	(0.227)	(0.051)	(0.079)
Above \$100k	-1.777	-1.543	0.145
	(0.234)	(0.047)	(0.055)
Σ^2			
Price	0.000	0.697	0.695
	(0.107)	(0.028)	(0.048)
1750mL	0.000	0.695	1.167
	(0.086)	(0.048)	(0.236)
Nesting Parameter ρ		0.423	
		(0.026)	
Fixed Effects		Brand+Quarter	
Model Predictions	25%	50%	75%
Own Elasticity: $\frac{\partial \log q_j}{\partial \log p_j}$	-5.839	-5.162	-4.733
Aggregate Elasticity: $\frac{\partial \log Q}{\partial \log P}$	-0.333	-0.329	-0.322
Own Pass-Through: $\frac{\partial p_j}{\partial c_j}$	1.256	1.284	1.320
Observed Wholesale Markup (PH)	0.188	0.233	0.276
Predicted Wholesale Markup (PH)	0.205	0.231	0.259

Note: The table above reports parameter estimates from our RCNL model. The price coefficient is lognormally distributed so that $\alpha_i = -e^{\pi_k^p + \Sigma \cdot \nu_i}$ is always negative and more negative for values of π_k^p closer to zero. High-income consumers $\pi^p = -2.04$ have smaller coefficients than low-income consumers -0.643 and are thus *less* price sensitive.

Own pass-through is change in equilibrium prices for product j (under PH) in response to a \$1.00 increase in the price of good j.

Aggregate elasticity is change in total spirits volume in response to a 1% price increase for all products.

Source: Harmonized Price and Quantity Data (top 750 products, average wholesale price below \$60 per liter), 24 quarterly periods. Authors' calculations.

Table 4: Best Substitutes: Diversion Ratios 2013 Q2

	Median Price	% Substitution		Median Price	% Substitution
Capt Morgan Spiced 1.75 L ($15.85)			Cuervo Gold 1.75 L ($18.33)		
Bacardi Superior Lt Dry Rum 1.75 L	12.52	13.07	Don Julio Silver 1.75 L	22.81	5.00
Bacardi Dark Rum 1.75 L	12.52	2.71	Cuervo Gold 1.0 L	21.32	3.82
Bacardi Superior Lt Dry Rum 1.0 L	15.03	2.44	Sauza Giro Tequila Gold 1.0 L	8.83	3.07
Smirnoff 1.75 L	11.85	2.36	Smirnoff 1.75 L	11.85	2.44
Lady Bligh Spiced V Island Rum 1.75 L	9.43	2.18	Absolut Vodka 1.75 L	15.94	2.06
Woodford 0.75 L ($34.55)			Beefeater Gin 1.75 L ($17.09)		
Jack Daniel Black Label 1.0 L	27.08	7.66	Tanqueray 1.75 L	17.09	12.80
Jack Daniel Black Label 1.75 L	21.85	4.91	Gordons 1.75 L	11.19	4.14
Jack Daniel Black Label 0.75 L	29.21	4.83	Seagrams Gin 1.75 L	10.23	2.85
Makers Mark 1.0 L	32.79	4.52	Bombay 1.75 L	21.95	2.27
Makers Mark 0.75 L	31.88	2.80	Smirnoff 1.75 L	11.85	2.27
Dubra Vdk Dom 80P 1.75 L ($5.88)			Belvedere Vodka 0.75 L ($30.55)		
Popov Vodka 1.75 L	7.66	7.56	Grey Goose 1.0 L	32.08	5.09
Smirnoff 1.75 L	11.85	3.15	Absolut Vodka 1.75 L	15.94	3.82
Sobieski Poland 1.75 L	9.09	3.14	Absolut Vodka 1.0 L	24.91	2.74
Grays Peak Vdk Dom 1.75 L	9.16	2.87	Smirnoff 1.75 L	11.85	2.43
Wolfschmidt 1.75 L	6.92	2.48	Grey Goose 0.75 L	39.88	2.22

Note: The table above reports diversion rates for five popular products. Per liter wholesale prices are reported for 2013Q2. We compute the diversion ratio for a small price change $D_{j\to k} = \frac{\partial q_k}{\partial q_j} / \left|\frac{\partial q_j}{\partial q_j}\right|$. Substitutes maintain the same product category but differ across products due to the nesting parameter and random coefficients incorporated in the RCNL model.
A plain logit would predict the best substitute as the product with the largest overall share: Smirnoff Vodka (80 Proof, 1.75L) with $s_{jt} = 1.2\%$ or 4.38% of "inside" sales.
Source: Authors' calculations

Table 5: Counterfactual Policies to Limit Ethanol Consumption

Policy	Product Prices
Sales Tax	$p_{jt} = mc_{jt} \cdot (1 + \tau_r)$
Volumetric Tax	$p_{jt} = mc_{jt} + \tau_v$
Ethanol Tax	$p_{jt} = mc_{jt} + \tau_e \cdot ABV_{jt}$
Sales+Volume Taxes	$p_{jt} = (mc_{jt} + \tau_0) \cdot (1 + \tau_r)$
Minimum Unit Price	$p_{jt} = \max\{mc_{jt}, \tau_u \cdot ABV_{jt}\}$
Ramsey-Revenue	$\mathbf{p}(\overline{R}) = \arg\max_{\mathbf{p} \geq \mathbf{mc}} CS(\mathbf{p})$ s.t. $(\mathbf{p} - \mathbf{mc}) \cdot \mathbf{q}(\mathbf{p}) > \overline{R}$
Ramsey-Ethanol	$\mathbf{p}(\overline{E}) = \arg\max_{\mathbf{p} \geq \mathbf{mc}} CS(\mathbf{p})$ s.t. $\sum_j e_j\, q_j \leq \overline{E}$
Monopoly	$\mathbf{p} = \arg\max_{\mathbf{p}} (\mathbf{p} - \mathbf{mc}) \cdot \mathbf{q}(\mathbf{p})$

Note: We examine eight policy alternatives to PH. In all counterfactuals PH pricing is replaced with taxes levied on a competitive wholesale market. *Sales* levies a single-rate sales tax (τ_r) on all spirits products to achieve the desired aggregate ethanol consumption level. Similarly, *Volume* and *Ethanol* model the impact of volumetric (τ_v) and ethanol-based (τ_e) taxes set to limit ethanol consumption. We also consider a policy which retains Connecticut's existing volumetric tax (τ_0) and layers on a sales tax. A *Minimum Price* enforces a floor based on ethanol content ($\tau_u \cdot ABV_{jt}\}$) but otherwise prices products competitively.
Finally, we examine the impacts of Ramsey prices where individual product prices are set to maximize consumer surplus while meeting different constraints. The first set of Ramsey prices are set to generate a required revenue (regardless of ethanol consumption). The second set of Ramsey prices are set to cap aggregate ethanol consumption (regardless of revenue generated).

Table 6: Distributional Impacts of Counterfactual Policies

	% Total Revenue	% Change in CS % Overall	Below $25k	$25k-$45k	$45k-$70k	$70k-$100k	Above $100k
No Change in Ethanol							
Ramsey (Ethanol)	-0.6	22.2	15.4	23.0	27.1	38.3	21.9
Minimum Price	12.7	21.6	12.9	21.5	24.3	34.5	22.5
Ethanol	272.2	7.4	4.2	7.2	7.5	11.7	7.9
Volume	279.1	6.4	2.6	5.3	5.2	9.5	7.4
Sales+Volume	365.7	-5.5	-1.8	-3.4	-3.6	-5.1	-7.1
Ramsey (Revenue)	375.7	-7.0	-4.7	-5.9	-6.2	-9.2	-7.7
Sales	386.6	-13.2	-4.5	-8.7	-8.5	-13.9	-16.6
-10% Ethanol							
Ramsey (Ethanol)	22.8	15.2	8.0	11.5	14.5	26.2	16.5
Minimum Price	47.1	14.5	5.3	10.1	11.2	21.6	17.2
Ethanol	301.2	0.6	-3.7	-3.2	-4.7	0.0	2.9
Volume	308.6	-0.5	-5.5	-5.3	-7.1	-2.3	2.3
Sales+Volume	396.3	-13.8	-10.2	-15.0	-16.7	-17.9	-14.1
Ramsey (Revenue)	406.6	-15.5	-13.8	-17.9	-19.8	-22.4	-14.6
Sales	407.1	-21.0	-12.6	-19.9	-21.1	-25.8	-23.2
+10% Ethanol							
Ramsey (Ethanol)	-13.7	28.1	21.4	33.3	38.3	48.2	26.4
Minimum Price	-12.7	27.7	19.6	32.2	36.6	45.8	26.9
Ethanol	239.7	13.8	11.6	17.8	20.0	23.2	12.4
Volume	245.9	13.0	10.3	16.1	17.9	21.3	12.1
Sales+Volume	327.1	2.6	6.4	8.4	10.0	8.1	-0.3
Ramsey (Revenue)	339.7	0.8	3.8	6.0	7.5	3.6	-1.6
Sales	358.4	-5.3	3.3	2.8	4.6	-1.6	-10.1

Note: The table above reports estimates of the impacts of the counterfactual policy alternatives described in Table 5 on tax revenue collected, overall consumer surplus and the distribution of consumer surplus across the five income bins. All effects are reported as percentage changes relative to the PH baseline. The top panel describes the impact of alternative policies that limit ethanol consumption to the same aggregate level as under PH while panels B and C report the effects of alternative policies that reduce and increase ethanol consumption by 10%, respectively. Revenue is calculated as the additional tax revenue raised by the state compared to the existing excise tax collections.
Source: Authors' calculations

Table 7: Reducing Overall Ethanol Consumption (Volumetric Taxes)

	No Change to Ethanol			No Change to Overall CS			No Change to CS by Income		
	Base	$wc = 1$	p^m	Base	$wc = 1$	p^m	Base	$wc = 1$	p^m
% Δ Ethanol	-0.00	-0.00	0.00	-9.30	-9.30	-8.67	-3.21	-3.21	-3.86
% Δ Tax Revenue	279.05	209.52	246.33	306.61	243.57	269.71	288.95	221.66	257.08
% Δ Manufacturer Profit	17.76	17.76	35.18	9.57	9.57	28.61	14.96	14.96	32.29
% Δ Total CS	6.40	6.40	5.93	0.01	0.01	0.02	4.23	4.23	3.33
% Δ CS by Income									
Below $25k	2.58	2.58	3.09	-4.96	-4.96	-3.93	0.01	0.01	0.00
$25k-$45k	5.34	5.34	5.42	-4.55	-4.55	-3.81	1.92	1.92	1.31
$45k-$70k	5.22	5.22	5.42	-6.29	-6.29	-5.33	1.22	1.22	0.61
$70k-$100k	9.55	9.55	9.90	-1.49	-1.49	-0.38	5.74	5.74	5.33
Above $100k	7.42	7.42	6.46	2.72	2.72	2.14	5.84	5.84	4.58
Tax per Liter	5.45	4.45	4.98	6.45	5.45	5.83	5.78	4.78	5.35

Note: The table above reports welfare estimates for the impacts of a counterfactual volumetric tax under three scenarios: (a) no change in overall ethanol consumption (b) minimizing ethanol consumption without reducing aggregate consumer surplus (c) minimizing ethanol consumption without reducing consumer surplus for any income bin.
In addition to the base model which sets the marginal cost of wholesaling equal to the manufacturer price and holds manufacturer prices fixed (Base), we allow for an additional $1 per liter wholesaling cost ($wc = 1$), or we allow manufacturers to endogenously set prices (p^m) in response to counterfactual taxes but with perfectly competitive wholesaling.
Manufacturer profits increase even when prices are held fixed because absent PH, consumers substitute to higher margin/quality products.
Existing volumetric taxes are $1.42/L.
Source: Authors' calculations

Figure 1: Price Indices by State, National Consumption Bundle (2013)

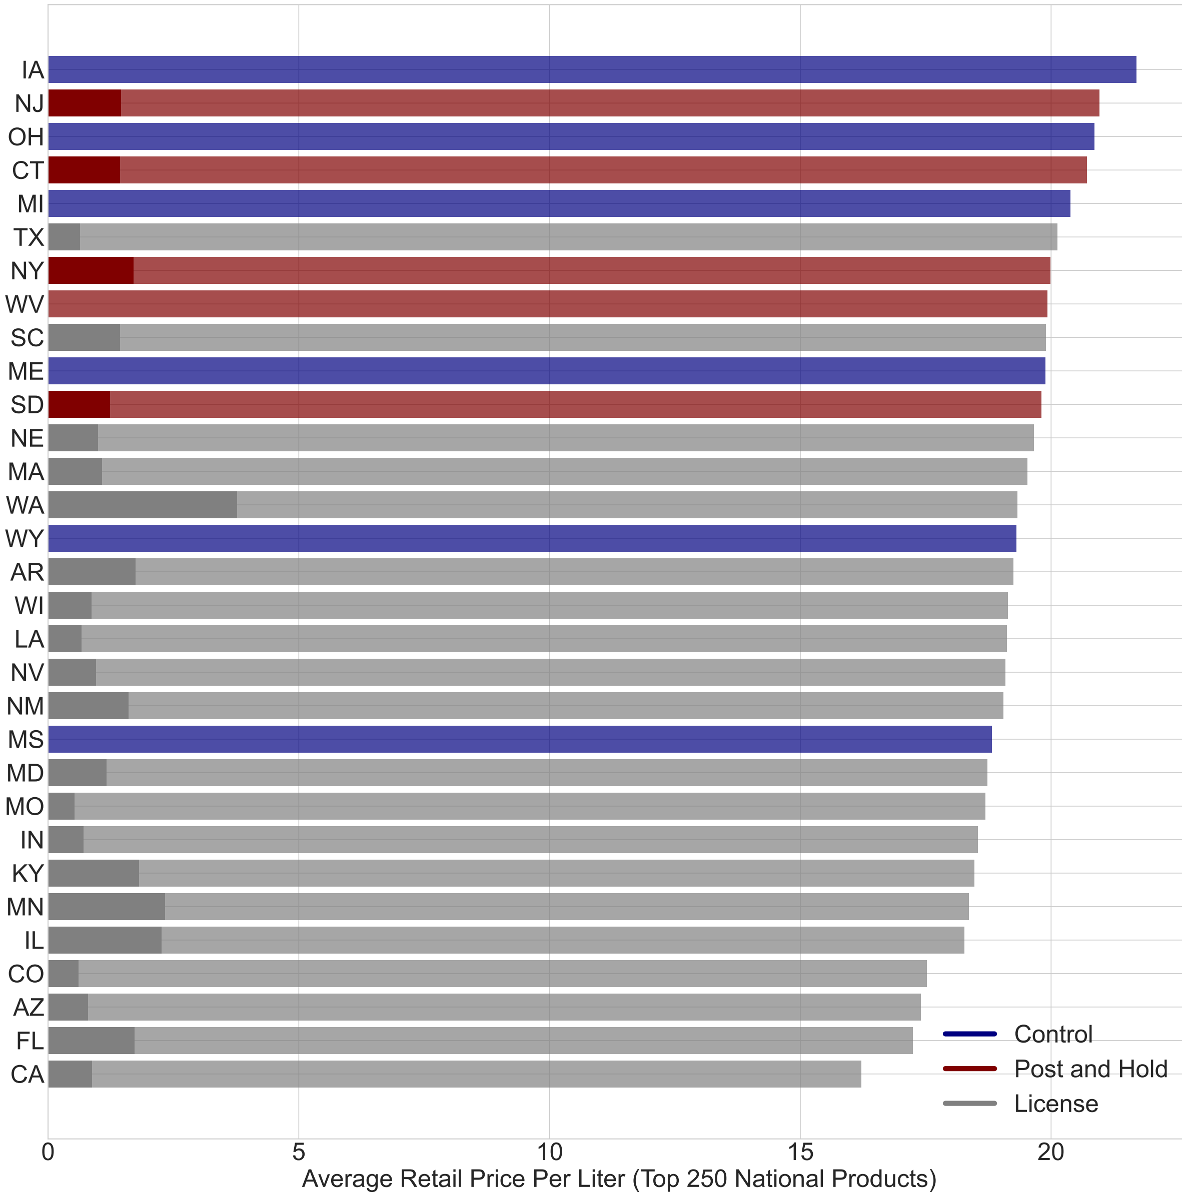

Note: The figure above plots the average retail price by state of the 250 best-selling products nation-wide. Retail prices in each state are weighted by the product's share within the top 250 national bundle by volume. As such, sales weights are constant across states so that the indices reflect only the differences in prices for the national bundle. License states such as Rhode Island and Delaware where we lack data describing sales of at least 1,000 products are excluded. Control states are shaded in blue, post and hold states in red and license states without post and hold regulations in grey. Darkly shaded bars on the left indicate state excise tax levied on the national bundle in license states (control states generally do not levy taxes on top of state markups).
Source: NielsenIQ Scanner Dataset.

Figure 2: Retail Prices for Vodka Products in Connecticut vs. Massachusetts (2013)

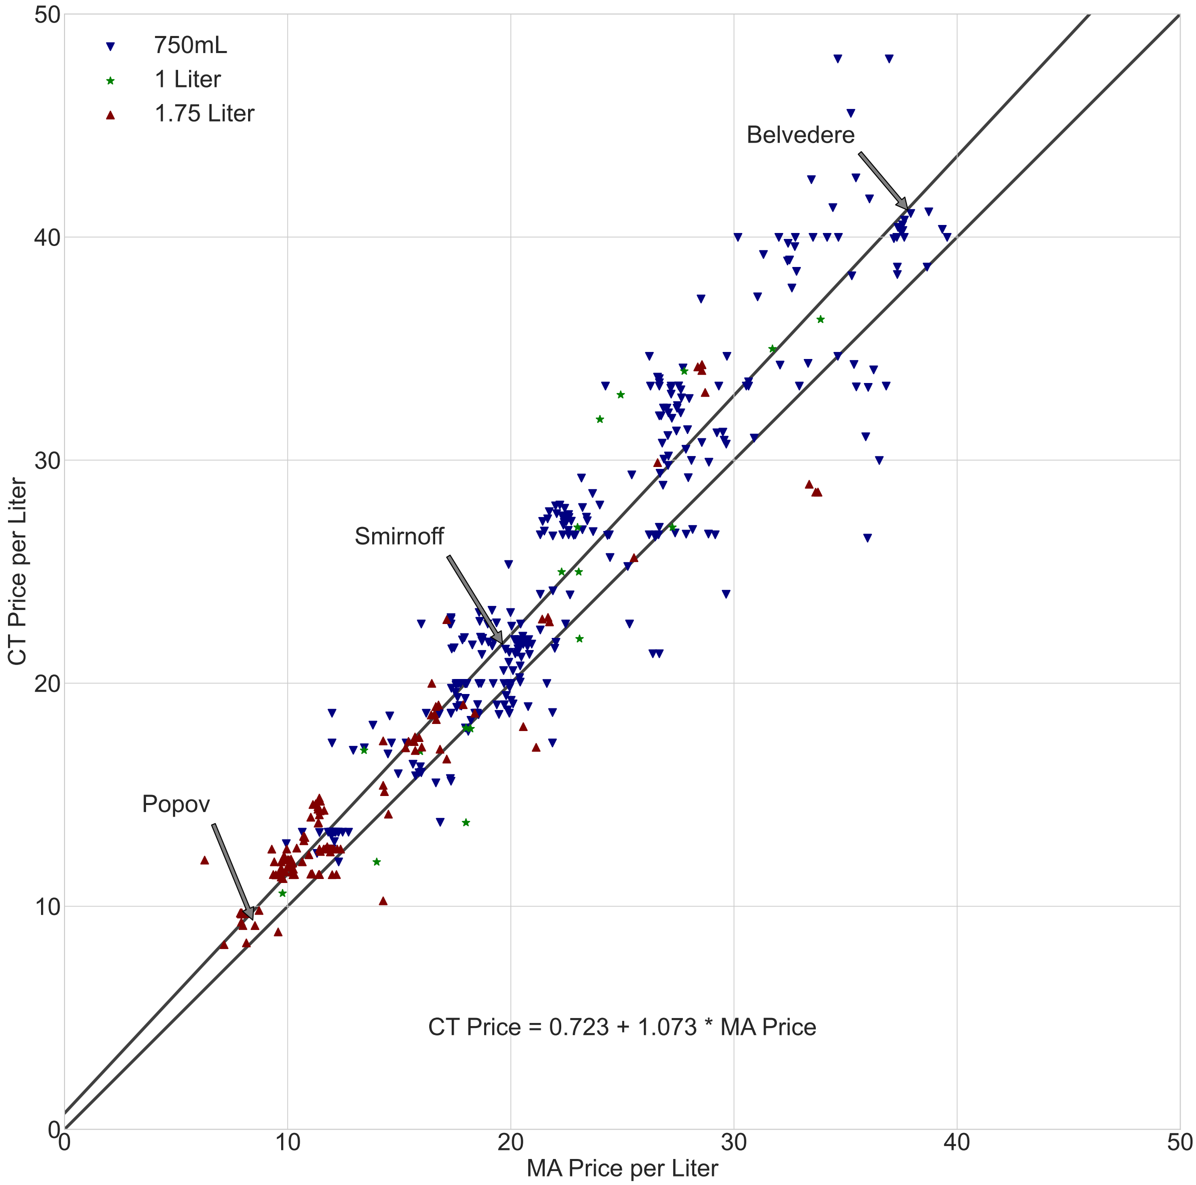

Note: The figure above compares the retail prices of individual products in Connecticut and the neighboring state of Massachusetts. Massachusetts prices are plotted on the x-axis and Connecticut prices are plotted on the y-axis with each dot representing brand-size combination, such as Smirnoff Vodka-750mL or Tanqueray Gin-1L. Prices are converted into dollars per liter and different colored markers denote 750mL (blue), 1000mL (green) and 1750mL (red) products. The dashed line plots the linear best fit and its coefficients are reported. The 45 degree line, corresponding to equal prices in Connecticut and Massachusetts, is shown as well.
Source: NielsenIQ Scanner Dataset.

Figure 3: Vodka Consumption in Connecticut and Massachusetts by National Price Per Liter (2013)

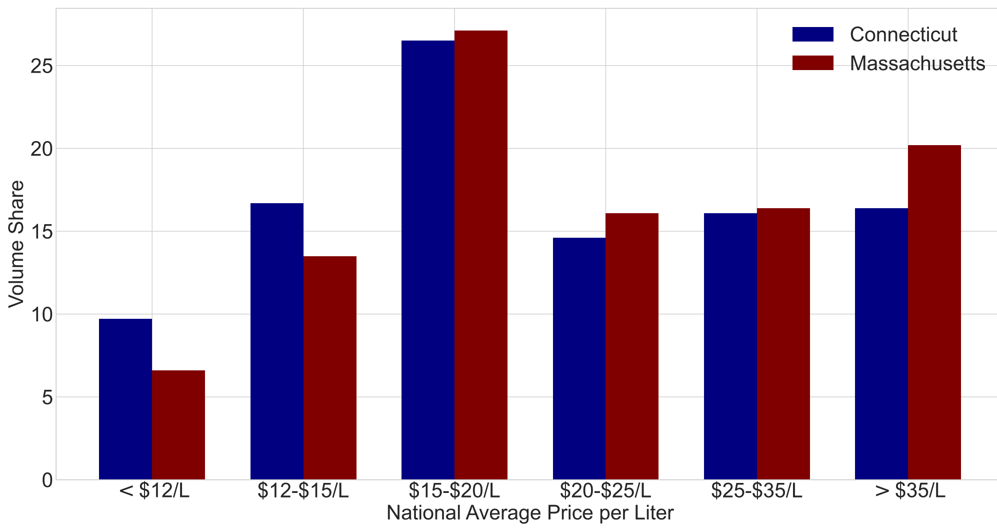

(a) 750mL Products

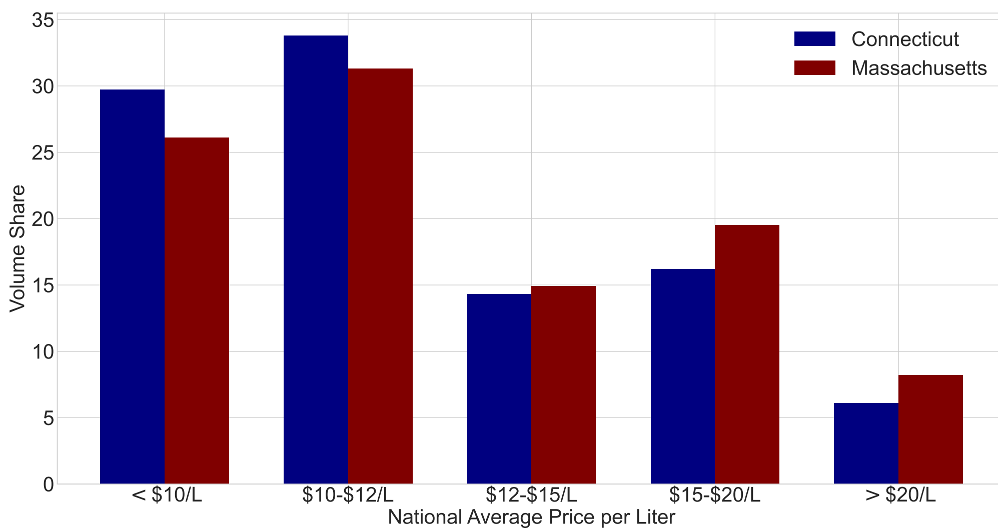

(b) 1.75L Products

Note: The charts above show the share of vodka consumption by volume in Connecticut (blue) and Massachusetts (red) for 750mL and 1.75L products by national price per liter category. A product's national price category is determined using the average price per liter across all Nielsen markets outside of Connecticut designated market areas. For products only sold in Connecticut or Massachusetts the state price is used in place of the national price to calculate price per liter. .
Source: NielsenIQ Scanner Data. All of 2013.

Figure 4: Manufacturer and Wholesale Prices Q2 2013

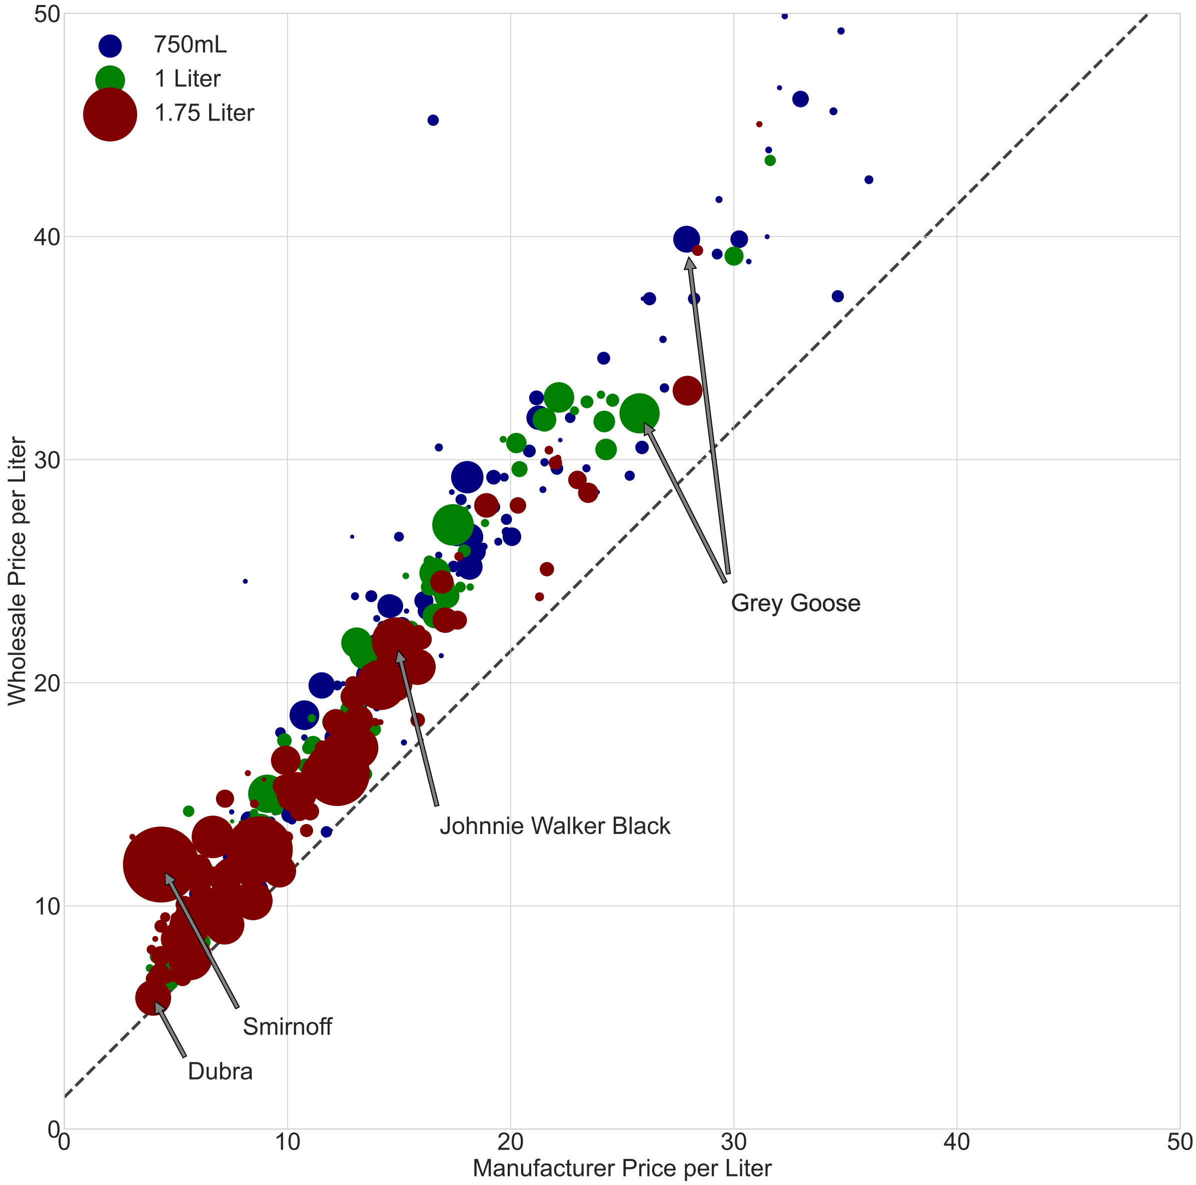

Note: The figure above plots the wholesale price against the manufacturer price, capturing how the ratio of wholesale to manufacturer price rises with manufacturer price. Prices are dollars per liter and different colored markers denote 750mL (blue), 1000mL (green) and 1750mL (red) products. Marker sizes are proportional to quarterly sales totals. The 45 degree line, corresponding to zero wholesale markup, is shown as well.
Source: Harmonized Price and Quantity Data. Period from 2013-04-01 to 2013-06-30.

Figure 5: Case Price by Wholesaler and Manufacturer Price, Four Top Selling Products

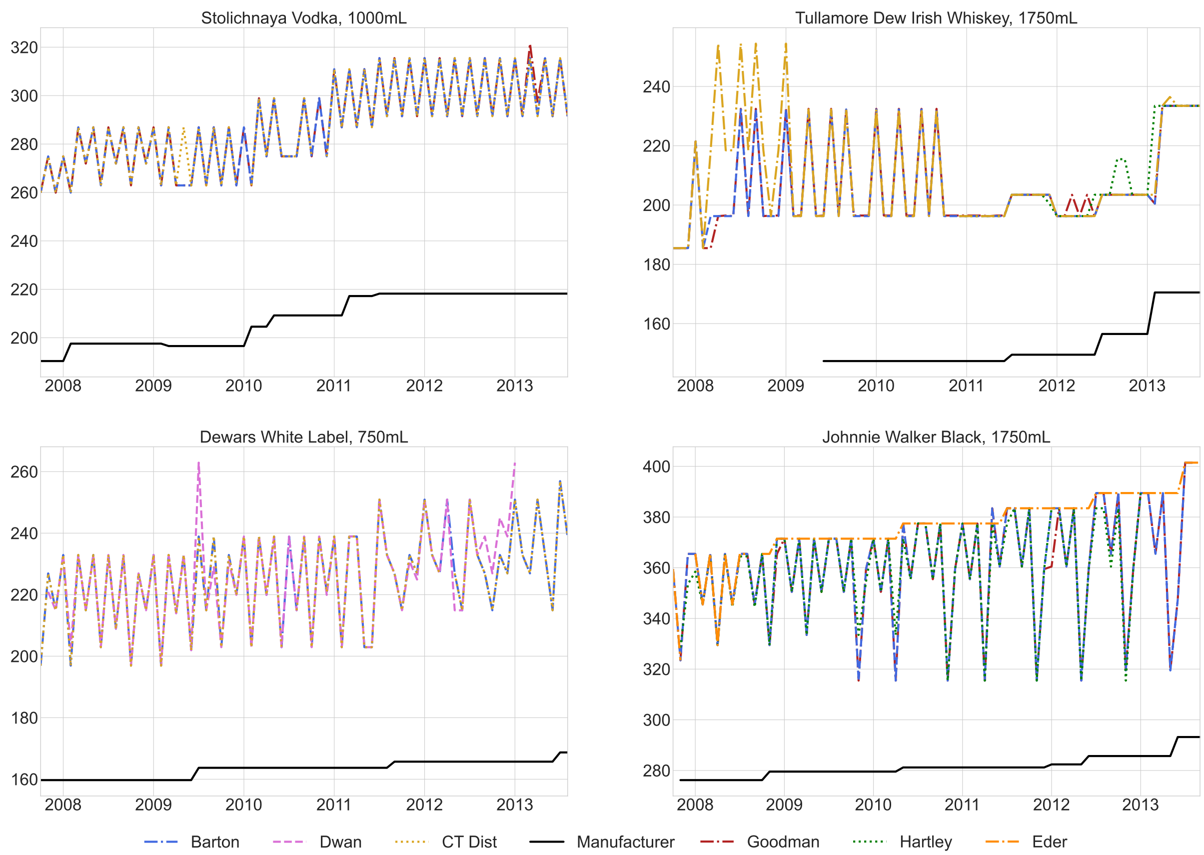

Note: The figure above plots monthly wholesale prices as well as the manufacturer price for four popular products between October 2007 and August 2013. Three wholesalers offer Stolichnaya Vodka, 1000mL (Goodman, Barton and CT Dist) and Dewars White Label, 750mL (Barton, CT Dist and Dwan), while four wholesalers sell Tullamore Dew, 1750mL (Barton, CT Dist, Goodman and Hartley) and Johnnie Walker Black, 1750mL (Barton, Eder, Goodman and Dwan) over the period. Prices offered by these distinct wholesalers overlap in the vast majority of months. While we might expect correlated wholesale price increases when manufacturer prices rise, which we observe, prices also exhibit considerable month-to-month changes between manufacturer price adjustments that happen in near lockstep across wholesalers.
Source: Harmonized Price and Quantity Data. Period from 2013-04-01 to 2013-06-30.

Figure 6: Estimated Own Elasticities and Diversion to the Outside Good

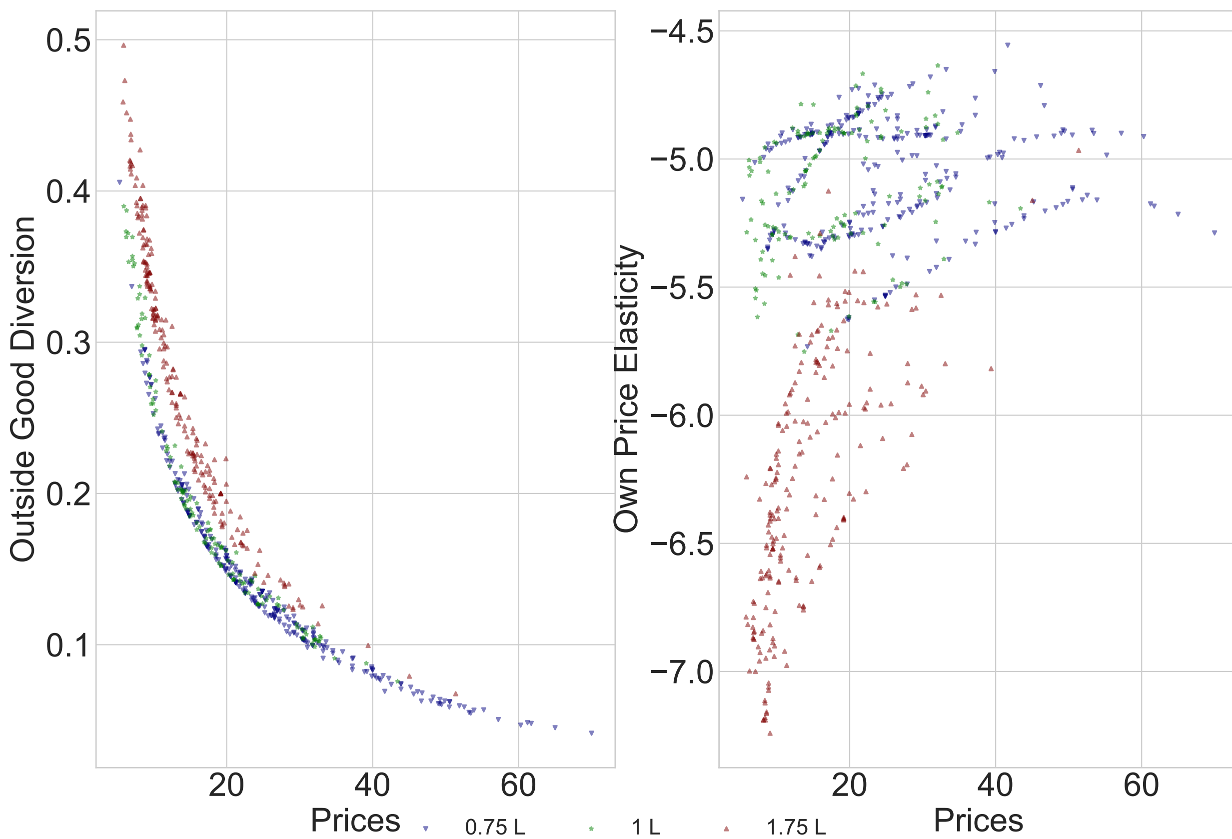

Note: The figure above plots estimated diversion ratios to the outside good (left panel) and own-price elasticities (right panel) against price where each observation is a product in 2013Q2. Diversion to the outside good is calculated as $D_{j\to 0} = \frac{\partial s_0}{\partial p_j} / \left|\frac{\partial s_j}{\partial p_j}\right|$ while own-price elasticities are given by: $e_{jj} = \frac{\partial s_j}{\partial p_j} \cdot \frac{p_j}{s_j}$. Low-price products have both higher diversion to the outside good and larger own-price elasticities, indicating that raising the prices of these products will most readily reduce aggregate ethanol consumption.

Figure 7: Prices Under PH vs. Other Policy Alternatives

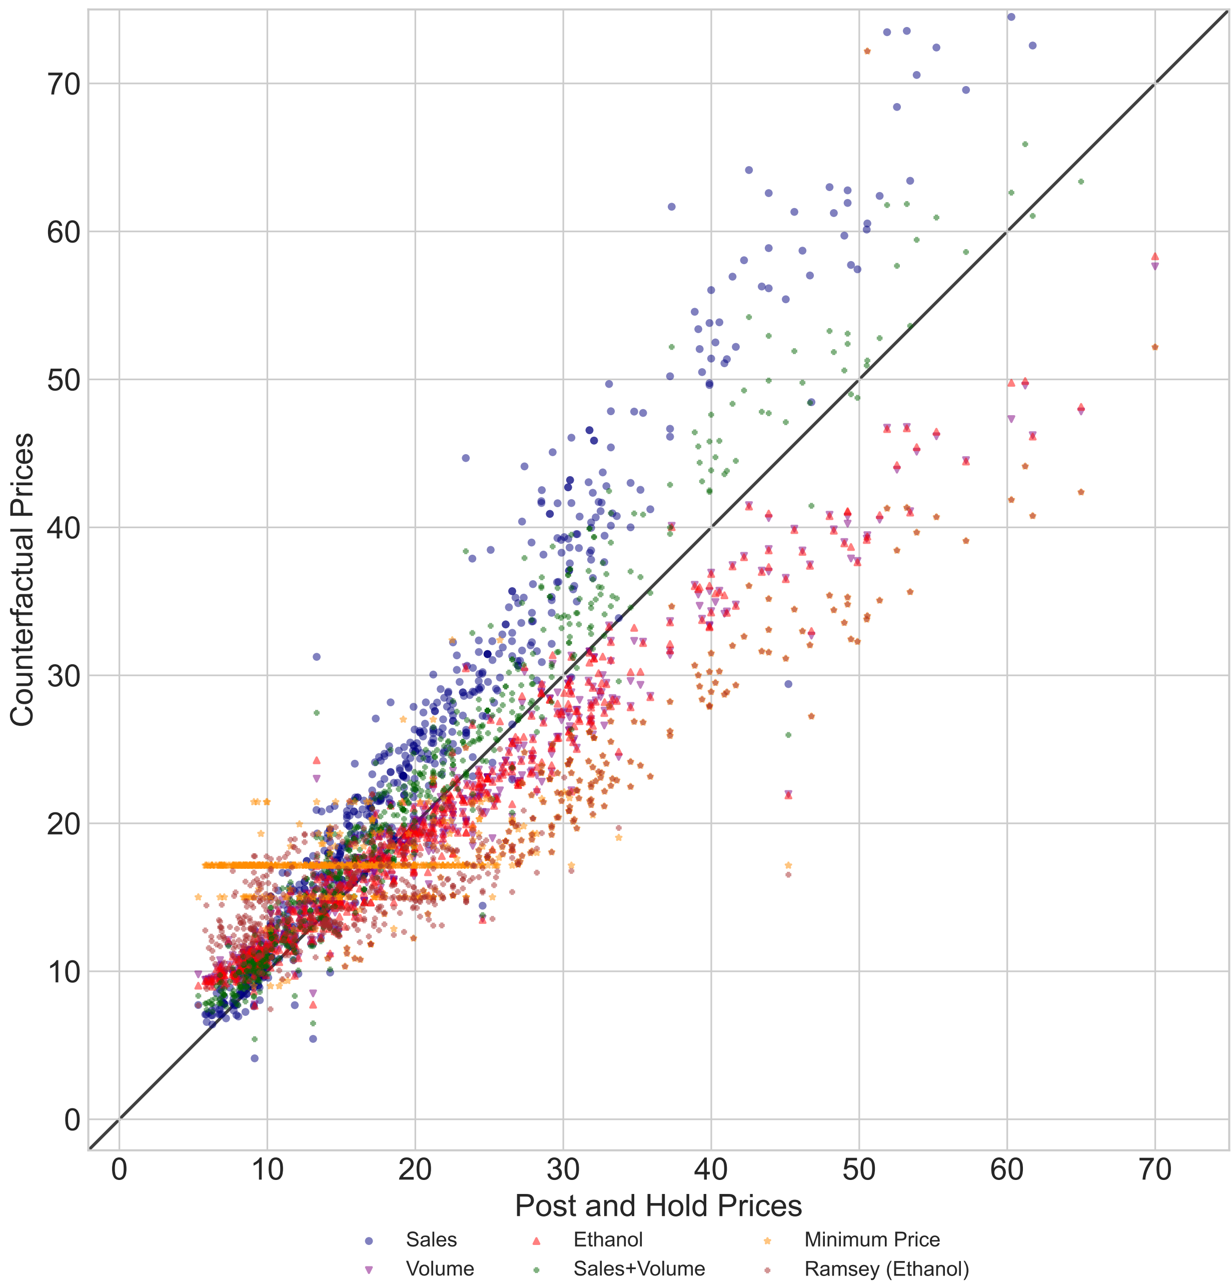

Note: The figure above plots product prices under PH against prices under our counterfactual policy alternatives. In each of our counterfactual scenarios we consider a tax rate that would keep the overall level of ethanol fixed at the status quo. Our taxes follow the definitions in Table 5, and are levied on a competitive market where wholesale price equals manufacturer price (with the exception of "Sales+Volume" which applies a sales tax to existing manufacturer prices and excise taxes). The solid black 45 degree line illustrates prices unchanged from PH.

Figure 8: Consumer Surplus vs. Tax Revenue and Ethanol Consumption Under Alternative Policies

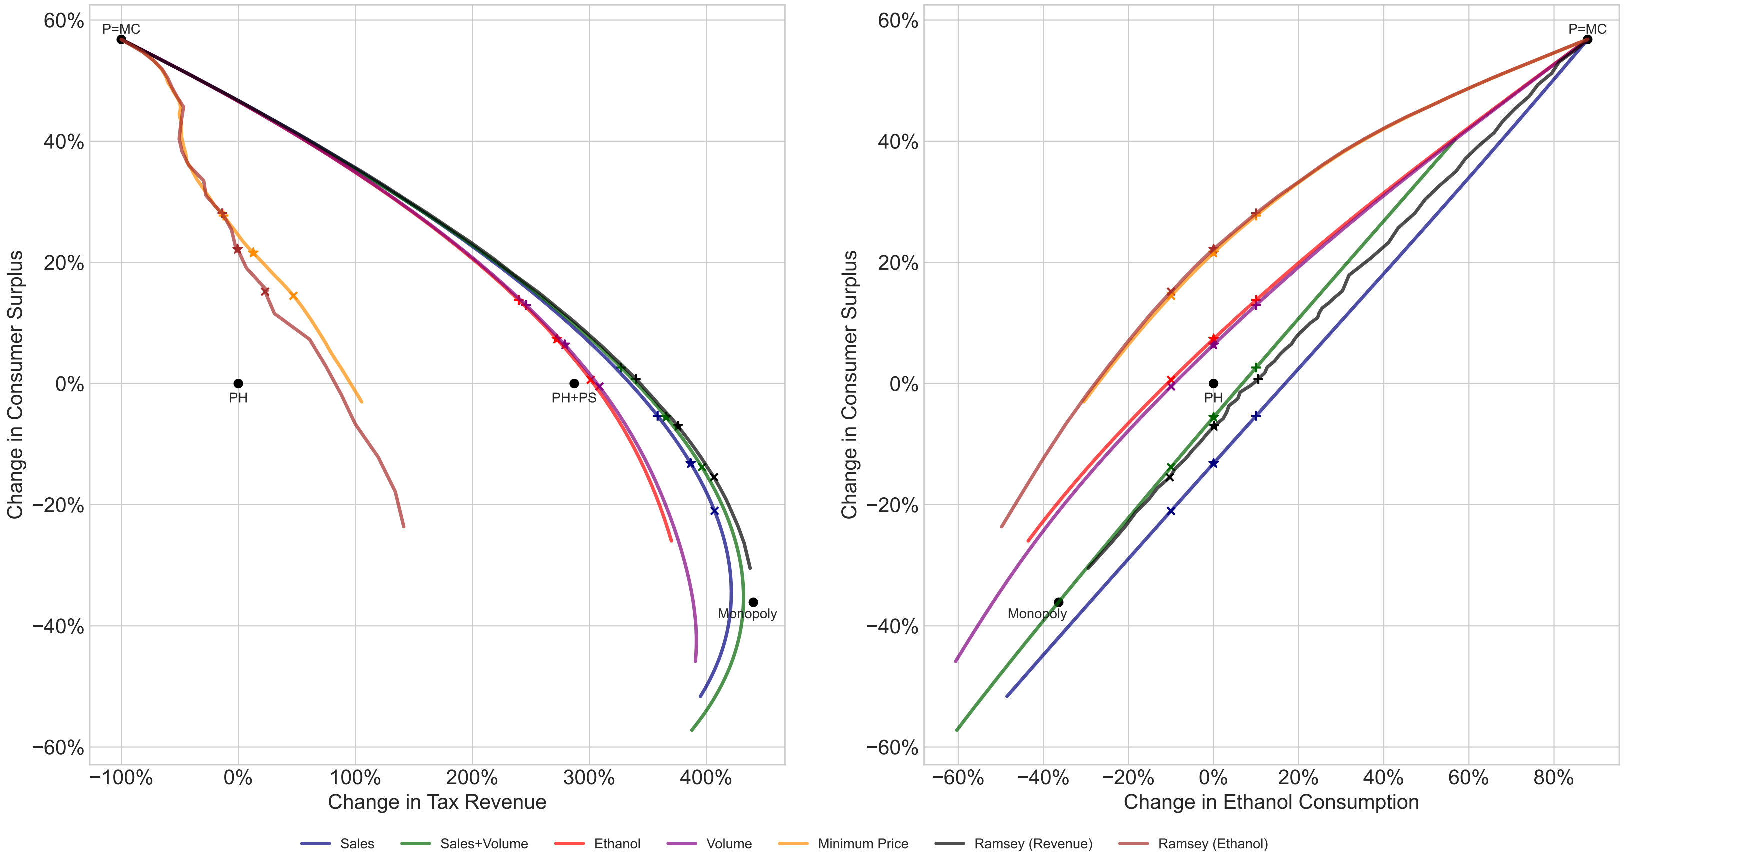

Note: The figure above plots the change tax revenue (left panel) and ethanol consumption (right panel) against the change in consumer surplus for each of the policy alternatives to PH detailed in Table 5. The frontiers trace the trade-off between consumer surplus and tax revenue or ethanol consumption for each policy instrument. Stars indicate an aggregate ethanol consumption level equal to total ethanol under PH, while x denotes 10% less and + denotes 10% more ethanol consumption (in the left panel higher ethanol consumption corresponds to less tax revenue). We also mark competitive prices without taxes (denoted by P = MC), and PH pricing. In the left panel we indicate the revenue generate by existing excise taxes under PH pricing as well as the sum of tax revenue and wholesale profits generated by PH (denoted by PH+PS).

A. Additional Theoretical Results

A.1. PH with a Single Product and Homogeneous Costs

We address this case in the main text and show that the first stage admits a dominant strategy of matching the lowest priced competitor so long as it is above your marginal cost.

Proof for Theorem 1
Consider a two-stage strategy of the form $\sigma_i(p_i^0, p_i^1)$. The second stage admits the unique dominant strategy where all players set $p_i^{1*} = \max\{c_i, \underline{p_i^0}\}$ where $\underline{p_i^0} = \min_i p_i^0$. For strategies of the form: $\sigma_i(p_i^0, \underline{p_i^0})$: $\sigma_i(p_i + \epsilon, \underline{p_i^0}) \geq \sigma_i(p_i, \underline{p_i^0})$ for $p_i \in [c_i, p_i^m)$. By induction the unique Nash Equilibrium to survive iterated weak dominance is $\sigma_i(p_i^m, \underline{p_i^0})$.

A.2. PH with a Single Product and Heterogeneous Costs

In the case of heterogeneous costs, the first stage becomes a bit more complicated. Begin by ordering the firms by marginal costs $c_1 \leq c_2 \cdots \leq c_N$. The market price $\hat{p}$ will be set by the lowest-cost firm (player 1). Other players play the iterated-weak-dominant-strategy $\sigma(p_i^0, p_i) = (p_i^m, \max\{\underline{p}^0, c_i\})$. Player 1 chooses p_i^0 to maximize the residual profit function:

$$\hat{p} = \underset{p_1^0 \in \{p_1^m, c_2, \ldots, c_n\}}{\arg\max} \quad \pi_i(p_1^0) = \frac{(p_1^0 - c_1) \cdot Q(p_1^0)}{\sum_k I[c_k < p_1^0]}$$

Player 1 can choose either to play its monopoly price and split the market evenly with the number of firms for which $c_i \leq p_1^m$, or it can set a lower price to reduce the number of firms who split the market. When the cost advantage of player 1 is small, we expect to see outcomes similar to the collusive outcome. As the cost advantage increases, it becomes more attractive for player 1 to engage in limit-pricing behavior. Because our wholesalers buy the same products from the upstream manufacturer/distillers in roughly similar quantities, we ignore the possibility of heterogeneous marginal costs in our empirical example. In practice, as long as the dispersion between heterogeneous costs is not too large, firms will not have an incentive to engage in limit-pricing.

A.2.1. PH with Heterogeneous Costs and Multiproduct Firms

We extend the single homogeneous good result to the case of heterogenous costs and multi-product firms, but continue to consider a single static Bertrand game. Now for each product j, the second stage admits the same form of a dominant strategy:

$$p_{ij}^* = \max\{c_{ij}, \underline{p_j^0}\} \quad \forall i, j$$

Firms now choose optimal strategies in first-stage prices, understanding what the outcome of the subgame will be, and facing both an *ad valorem* tax τ and a specific tax t:

$$\begin{aligned}\pi_i &= \max_{p_{ij}:j\in J_i} \sum_{j\in J_i} (p_{ij}(1-\tau) - c_{ij} - t)\cdot q_{ij} \\ \frac{\partial \pi_i}{\partial p_k} &= q_{ik}(1-\tau) + \sum_{j\in J_i} (p_{ij}(1-\tau) - c_{ij} - t)\cdot \frac{\partial q_{ij}}{\partial p_k} \quad \forall i \in \mathcal{I}_k \end{aligned} \tag{A.1}$$

The insight from the homogenous goods case is that firms will not all operate by setting their FOC to zero. The idea is that firms act as a monopolist when decreasing prices, but act as price-takers when increasing prices. For each firm $i \in \mathcal{I}_k$ (where $\mathcal{I}_k$ denotes the set of firms selling product k), only the weaker condition $\frac{\partial \pi_i}{\partial p_k} \geq 0$ holds, and it is not necessarily true that $\frac{\partial \pi_i}{\partial p_k} \leq 0$ for all $i \in \mathcal{I}_k$.

If firms have sufficiently similar marginal costs,[69] no firm will engage in limit pricing and there will be a constant division of the market on a product by product basis (depending on how many firms sell each product). Let λ_{ik} be the share that i sells of product k. Under a constant division, $\lambda_{ik} \perp p_k$, we can write $q_{ik} = \lambda_{ik} Q_k$ where Q_k is the market quantity demanded of product k, so that $\forall i = 1, \ldots, N$:

$$Q_k \lambda_{ik}(1-\tau) + (p_k(1-\tau) - c_{ik} - t)\cdot \frac{\partial Q_k}{\partial p_k}\lambda_{ik} + \sum_{j\in J_i} (p_j(1-\tau) - c_{ij} - t)\cdot \frac{\partial Q_j}{\partial p_k}\lambda_{ij} \geq 0$$

$$\underbrace{Q_k(1-\tau) + (p_k(1-\tau) - c_{ik} - t)\cdot \frac{\partial Q_k}{\partial p_k}}_{\text{Single Product Monopolist}} + \underbrace{\sum_{j\in J_i} (p_j(1-\tau) - c_{ij} - t)\cdot \frac{\partial Q_j}{\partial p_k}\frac{\lambda_{ij}}{\lambda_{ik}}}_{\text{Cannibalization}} \geq 0$$

For each product k, except in the knife-edge case, the first-order condition holds with equality for exactly one firm i. This establishes a least upper bound:

$$\underbrace{Q_k(1-\tau) + (p_k(1-\tau) - t)\cdot \frac{\partial Q_k}{\partial p_k}}_{\text{Marginal Revenue}} + \min_{i:k\in J_i} \underbrace{\left[-c_{ik}\frac{\partial Q_k}{\partial p_k} + \sum_{j\in J_i} (p_j(1-\tau) - c_{ij} - t)\cdot \frac{\partial Q_j}{\partial p_k}\frac{\lambda_{ij}}{\lambda_{ik}} \right]}_{\text{Opportunity Cost of Selling}} = 0 \tag{A.2}$$

Intuitively, the firm that sets the price of good k under PH is the firm for which the opportunity cost of selling k is the smallest, either because of a marginal cost advantage, or because it doesn't sell close substitutes. Given the derivatives of the profit function, the other firms would prefer to set a higher price, the price they would charge if they were a monopolist selling good k. This arises because just as in the single good case, firms can unilaterally reduce the amount of surplus (by cutting their first-stage price), but no firm can affect the division of the surplus (since all price cuts are matched in the second stage).[70]

The competitive equilibrium under PH results in prices at least as high as the lowest-opportunity-

[69] Formally we need that $c_{ik} \leq \underline{p}_k^0$ for all firms $i \in \mathcal{I}_k$

[70] Again this presumes that λ is fixed, and that firms do not engage in limit pricing to drive competitors out of the market.

cost single-product monopolist would have set, even though firms play a single period non-cooperative game, in which several firms distribute identical products. This also suggests a strategy we could observe in data. In the first stage, firms set their preferred "monopoly" price for each good, and in the second stage, firms update to match the lowest-opportunity-cost monopolist. In practice, we see very little updating in the second stage of the game, perhaps because the game is played month after month among the same players.

We can also do some simple comparative statics. Assume we increase the number of firms who sell product k. Normally this would lead to a decrease in price p_k. However, unless the entrant has a lower opportunity cost of selling than any firm in the existing market, prices would not decline, and we would expect the division of surplus λ_k to be reduced for the incumbents to accommodate the entrant. If this raises the opportunity cost of selling for the lowest-price firm, then more wholesale firms might counter-intuitively lead to higher prices.[71]

A.3. Comparing Markups Under PH and a Social Planner

In the main text we present the pricing rules of a PH wholesalers and a social planner maximizing social surplus while ensuring a minimum level of revenue and limiting external damage from the atmospheric externality of ethanol consumption. Below we derive the social planner's pricing rule ((20) and compare markups set under PH to markups a social planner would set whether she is ignoring or addressing the ethanol externality.

A.3.1. Social Planner's Pricing Rule

We consider the problem of a social planner who faces demand $\mathbf{Q}(\mathbf{p})$, and sets the prices $p_j \in \mathbf{p}$ of all products to maximize total surplus subject to two additional constraints: a minimum level of revenue $\overline{R}$, and a maximum level of externalities arising from ethanol consumption $\overline{E}$.

$$\begin{aligned} \max_{\mathbf{p}} \quad & CS(\mathbf{Q}(\mathbf{p})) - C(\mathbf{Q}(\mathbf{p})) \\ \text{subject to} \quad & \mathbf{p} \cdot \mathbf{Q}(\mathbf{p}) - C(\mathbf{Q}(\mathbf{p})) \geq \overline{R} \\ \text{and} \quad & E(\mathbf{Q}(\mathbf{p})) \leq \overline{E}. \end{aligned} \tag{A.3}$$

where the social benefit of consumption is the same as the private benefit defined as the sum of the areas under the demand curves: $CS(\mathbf{Q}(\mathbf{p})) = \sum_{k \in \mathcal{J}} \int_0^{Q_k} p_k(Q1, Q_2, ..., Q_{k-1}, Z_k, Q_{k+1}, ..., Q_n) dZ_k$. The cost of producing alcoholic beverages is captured by $C(\mathbf{Q}(\mathbf{p}))$. We can write the social planner's Lagrangian:

$$\mathcal{L}(\mathbf{p}) = CS(\mathbf{Q}(\mathbf{p})) - C(\mathbf{Q}(\mathbf{p})) + \lambda_r(\mathbf{p} \cdot \mathbf{Q}(\mathbf{p}) - C(\mathbf{Q}(\mathbf{p})) - \overline{R}) - \lambda_e(E(\mathbf{Q}(\mathbf{p})) - \overline{E}). \tag{A.4}$$

The Lagrange multiplier λ_r measures the social value of an additional dollar of revenue, while λ_e measures the shadow cost of an extra unit of external damage caused by alcohol consumption. This nests the well-known *Ramsey problem.* A common (though by no means necessary) assumption is that the externality is *atmospheric*, or that it depends only on total ethanol consumption and not the source of the ethanol nor the identity of the consumer, such that $E(\mathbf{Q}(\mathbf{p})) = \sum_j e_j \cdot Q_j(\mathbf{p})$.[72]

[71] This is different from the mechanism in other work on price-increasing competition such as Chen and Riordan (2008).

[72] This assumption would be violated if for example, if tequila generates more externalities per unit of ethanol than

The (interior solutions to the) first order conditions return the two constraints:

$$\frac{\partial \mathcal{L}}{\partial \lambda_r}: \quad \mathbf{p}\cdot\mathbf{Q}(\mathbf{p}) - C(\mathbf{Q}(\mathbf{p})) = \overline{R}$$
$$\frac{\partial \mathcal{L}}{\partial \lambda_e}: \quad E(\mathbf{Q}(\mathbf{p})) = \sum_{k\in\mathcal{J}} e_k \cdot Q_k(\mathbf{p}) = \overline{E}$$
$$\frac{\partial \mathcal{L}}{\partial p_j}: \quad \sum_{k\in\mathcal{J}}(p_k - mc_k)\frac{\partial Q_k}{\partial p_j} + \lambda_r\left(Q_j + \sum_{k\in\mathcal{J}}(p_k - mc_k)\frac{\partial Q_k}{\partial p_j}\right) - \lambda_e \sum_{k\in\mathcal{J}} e_k \frac{\partial Q_k}{\partial p_j} = 0.$$

Separating out product j, dividing through by $\frac{\partial Q_j}{\partial p_j}$ and re-writing the expression in terms of the diversion ratio, $D_{jk} = -\frac{\partial Q_k}{\partial p_j}/\frac{\partial Q_j}{\partial p_j}$[73] and own price elasticity $\epsilon_{jj} = \frac{\partial Q_j}{\partial p_j}\cdot\frac{p_j}{Q_j}$ gives:

$$(1+\lambda_r)(p_j - mc_j) - \lambda_r p_j \frac{1}{|\epsilon_{jj}|} - \lambda_e e_j - (1+\lambda_r)\sum_{k\neq j} D_{jk}(p_k - mc_k) + \lambda_e \sum_{k\neq j} D_{jk} e_k = 0.$$

which can be solved for p_j as the social planner's pricing rule:

$$p_j = \frac{|\epsilon_{jj}|}{|\epsilon_{jj}| - \frac{\lambda_r}{1+\lambda_r}}\left(mc_j + \frac{\lambda_e}{1+\lambda_r}e_j + \sum_{k\neq j} D_{jk}\left[p_k - mc_k - \frac{\lambda_e}{1+\lambda_r}e_k\right]\right)$$

or equation (20) in the main text.

The first term functions like the usual inverse elasticity rule Lerner markup with $\frac{\lambda_r}{1+\lambda_r} = \theta$ behaving like a conduct parameter where $\theta = 0$ corresponds to the perfectly competitive solution and $\theta = 1$ corresponds to the monopoly solution. The first two terms in parentheses, $mc_j + \frac{\lambda_e}{1+\lambda_r}e_j$, represent the effective marginal cost. When $\lambda_e > 0$, the marginal cost of production, mc_j, is augmented by the marginal external damage, e_j. The final term, $\sum_{k\neq j} D_{jk}\left[p_k - mc_k - \frac{\lambda_e}{1+\lambda_r}e_k\right]$, represents the *opportunity cost* of selling j, which is that fraction of consumers D_{jk} who switch to k as the price of j rises multiplied by the price less marginal cost (adjusted for the externality). Trading off these opportunity costs is a distinguishing feature of the multi-product Ramsey problem. Absent any revenue constraint, $\lambda_r = 0$, the first best solution to the planner's problem is to set prices at their Pigouvian rates $p_k = c_k + \lambda_e e_k$. More generally, for any revenue level and external damage (λ_r, λ_e) the Ramsey solution in (20) will maximize social surplus or minimize deadweight loss.

A.3.2. Comparing PH and the Planner's Pricing Rule Ignoring the Externality

As described in the main text, under PH the price for product j will be set by the wholesaler with the lowest opportunity cost of selling according to:

vodka or if 1750mL bottles generate more externalities *per liter* than 750mL bottles. Recent work by Griffith et al. (2019) shows that if consumer preferences across beer, wine and spirits are correlated with their marginal externality of alcohol consumption, taxes that vary across categories will more effectively address the external damage of alcohol consumption.

[73] Antitrust practitioners will recognize D_{jk} as the diversion ratio from j to k given by $D_{jk} = \frac{\partial q_k}{\partial p_j}/\left|\frac{\partial q_j}{\partial p_j}\right|$.

$$p_j = \frac{|\epsilon_{jj}|}{|\epsilon_{jj}| - 1} \cdot \left[mc_j + \sum_{k \in \mathcal{J} \backslash j} \kappa_{jk} \cdot D_{jk} \cdot (p_k - mc_k) \right] \tag{A.5}$$

where κ_{jk} is the ratio of the pivotal (lowest opportunity cost) firm's market share of product k relative to product j, which may be zero for many products in the full set $\mathcal{J}$.

In contrast, a social planner facing a minimum revenue constraint, $\overline{R}$, but ignoring the externality would price product j according to the Ramsey rule:

$$p_j = \frac{|\epsilon_{jj}|}{|\epsilon_{jj}| - \frac{\lambda_r}{1+\lambda_r}} \left[mc_j + \sum_{k \in \mathcal{J} \backslash j} D_{jk} (p_k - mc_k) \right] \tag{A.6}$$

where λ_r is the shadow value of revenue.

For both PH wholesalers and the social planner, prices will be higher on products with less elastic demand and higher on products where consumers react to higher prices by switching to higher margin products (and not to the outside option).

But as the main text discusses markups set by the social planner will differ from those set by wholesalers under PH in two key ways. First, the social planner's opportunity cost of selling any product will differ from the PH wholesaler. While the lowest opportunity cost wholesaler considers diversion to the other products it sells in accordance with $\kappa_{j,k}$, the social planner effectively sets $\kappa_{j,k} = 1$ for all substitutes. If a wholesaler controls a small share of the market for j and a large share for k it may be that $\kappa_{j,k} > 1$, though for many products we have $\kappa_{j,k} = 0$.

Because she accounts for the broadest set of opportunity costs, the social planner will raise even the same amount of revenue as the PH wholesalers in aggregate with less deadweight loss than PH. Second, since $\frac{\lambda_r}{(\lambda_r+1)} < 1$, except in the limit where $\lambda_r \to \infty$, the term multiplying the opportunity cost will be smaller under the social planner than under PH wholesalers. This is because the profit maximizing firms don't place any weight on consumer surplus like the planner does.

A.3.3. Planner's Problem and Decentralized Solution with Externality

If the social planner aims to also also limit ethanol consumption while raising revenue $\overline{R}$, the pricing rule will also account for external damage[74]:

$$p_j = \frac{|\epsilon_{jj}|}{|\epsilon_{jj}| - \frac{\lambda_r}{1+\lambda_r}} \left[mc_j + \sum_{k \neq j} D_{jk} [p_k - mc_k] + \frac{\lambda_e}{1+\lambda_r} \left(e_j - \sum_{k \neq j} D_{jk} e_k \right) \right] \tag{A.7}$$

The first terms are as above with the addition of a term that relates prices and product-specific externalities. The price of product j rises with its marginal damage (e_j) but declines if consumers readily shift to high marginal damage products (e_k).

Another way to frame this problem is Dixit (1985)'s "principle of targeting", which is further detailed by Sandmo (1975) and Oum and Tretheway (1988), and shown to be reasonably general

[74]We assume the more interesting case where the revenue resulting from addressing the externality alone would not raise $\overline{R}$ and thus $\lambda_r > 0$.

by Kopczuk (2003). In this framework correcting the externality and hitting revenue target are independent problems. A fiscal authority seeking to raise revenue $\overline{R}$ would set the Pigouvian tax equal to marginal damage $\lambda_e e_j$, and then set the remaining markup on product j according to a Ramsey pricing rule with a revenue requirement of $\overline{R}' = \overline{R} - R^P$ where R^P is the total revenue raised from Pigouvian taxes.

This delineation of the problem highlights how addressing the externality flattens markups across the consumer's perceived "quality" gradient. Two products with the same proof will carry more similar markups under this pricing regime than one where $\overline{R}$ is solely raised by a pricing rule like equation A.5 where the externality is not addressed. Because the prices resulting from equation A.7 raise prices exactly on those products that most contribute to the public health externalities of spirits consumption, these prices will most efficiently raise $\overline{R}$. While our policy experiments deviate from this formulation as we seek to hold ethanol consumption under PH fixed rather than raising the same revenue as the wholesalers in aggregate, the intuition that the state sets lower markups than PH wholesalers on products favored by consumers for characteristics besides ethanol content will carryover to all of our policy experiments.

For any given product it is not clear whether price would be lower under a social planner of PH wholesaler. Low-quality but high ethanol products like Dubra vodka will see higher prices under the social planner as the price is raised to reflect its external damage relative to a PH price that reflected only its high own-price and cross-price elasticity. Low-proof products like Malibu rum, which is 21% ABV, on the other hand, may see price reductions as their external damage is relatively modest.

A.3.4. Counterfactual Analysis

Instead of setting prices to maximize social surplus subject to revenue and ethanol constraints, our counterfactual analysis sets product-specific ("Ramsey") prices that maximize consumer surplus subject to a revenue or ethanol constraint, as well as the constraint that no product is sold below marginal cost:

$$\begin{gathered}\max_{\mathbf{p}\geq\mathbf{mc}} CS(\mathbf{Q}(\mathbf{p})) \\ \text{subject to } \mathbf{p}\cdot\mathbf{Q}(\mathbf{p}) - C(\mathbf{Q}(\mathbf{p})) \geq \overline{R} \\ \text{and } E(\mathbf{Q}(\mathbf{p})) \leq \overline{E}.\end{gathered}$$

This allows us to benchmark just how much better off the state could make consumers through alternative prices while raising a certain amount of revenue or achieving a specific aggregate ethanol reduction. Focusing on consumer welfare allows us to explicitly show these trade-offs. Because the state is setting product-specific taxes/prices on top of a competitive wholesale tier, the difference between price and cost summed across all products that is often considered producer surplus, is instead tax revenue in our counterfactuals. Figure 8 maps the trade-off between consumer surplus and tax revenue, including prices that maximize the combination of consumer surplus and tax revenue that comprises total surplus.

B. Empirical Implementation Details

B.1. Recovering Manufacturer Marginal Costs

This part builds on Jaffe and Weyl (2013) and Appendix E from Miller and Weinberg (2017) and almost exactly follows the implementation in Backus et al. (2021a); Conlon and Gortmaker (2020). The wrinkle here is that we observe the manufacturer prices $\mathbf{p^m}$ which simplify matters considerably, and we have the addition of the existing excise tax τ_0, which we show does not create any new issues.

We write the manufacturer's first order conditions as:

$$\mathbf{p^m} = \mathbf{mc^m} + \left(\mathcal{H}_M \odot \left(\frac{\partial \mathbf{p^w}}{\partial \mathbf{p^m}} \cdot \Omega(\mathbf{p^w})\right)\right)^{-1} \mathbf{s}(\mathbf{p^w}) \tag{B.1}$$

This requires that we estimate the pass-through matrix $\frac{\partial \mathbf{p^w}}{\partial \mathbf{p^m}}$.

In order to do so, we re-examine the wholealers' problem: a system of J first order conditions and J prices $\mathbf{p^w}$, with manufacturer prices $\mathbf{p^m}$ and wholesaling costs (including taxes) τ_0 serving as parameters:[75]

$$f(\mathbf{p^w}, \mathbf{p^m}, \tau_0) \equiv \mathbf{p^w} - \underbrace{(\mathbf{p^m} + \tau_0)}_{=\mathbf{mc^w}} - (\underbrace{\mathcal{H}_{PH}(\kappa) \odot \Omega(\mathbf{p^w})}_{\equiv \Delta(\mathbf{p^w})})^{-1} \mathbf{s}(\mathbf{p^w}) = 0 \tag{B.2}$$

Where $\Delta(\mathbf{p^w}) \equiv \mathcal{H}_{PH} \odot \Omega(\mathbf{p^w})$ is the PH augmented matrix of demand derivatives.

We differentiate the wholesalers' system of FOC's with respect to p_l, to get the $J \times J$ matrix with columns l given by:

$$\frac{\partial f(\mathbf{p^w}, \mathbf{p^m}, \tau_0)}{\partial p^w_\ell} \equiv e_\ell - \Delta^{-1}(\mathbf{p^w}) \left[\mathcal{H}_{PH} \odot \frac{\partial\, \Omega(\mathbf{p^w})}{\partial\, p^w_\ell}\right] \Delta^{-1}(\mathbf{p^w})\, \mathbf{s}(\mathbf{p^w}) - \Delta^{-1}(\mathbf{p^w})\, \frac{\partial \mathbf{s}(\mathbf{p^w})}{\partial p^w_\ell}. \tag{B.3}$$

The complicated piece is the demand Hessian: a $J \times J \times J$ tensor with elements (j, k, ℓ), $\frac{\partial^2 s_j}{\partial p^w_k \partial p^w_\ell} = \frac{\partial^2 \mathbf{s}}{\partial \mathbf{p^w} \partial p^w_\ell} = \frac{\partial\, \Omega(\mathbf{p^w})}{\partial\, p^w_\ell}$.

We can follow Jaffe and Weyl (2013) and apply the multivariate IFT. The multivariate IFT says that for some system of J nonlinear equations $f(\mathbf{p^w}, \mathbf{p^m}, \tau_0) = [F_1(\mathbf{p^w}, \mathbf{p^m}, \tau_0), \ldots, F_J(\mathbf{p^w}, \mathbf{p^m}, \tau_0)] = [0, \ldots, 0]$ with J endogenous variables $\mathbf{p^w}$ and J exogenous parameters $\mathbf{p^m}$.

$$\frac{\partial \mathbf{p^w}}{\partial \mathbf{p^m}} = -\begin{pmatrix} \frac{\partial F_1}{\partial p^w_1} & \cdots & \frac{\partial F_1}{\partial p^w_J} \\ \cdots & \cdots & \cdots \\ \frac{\partial F_J}{\partial p^w_1} & \cdots & \frac{\partial F_J}{\partial p^w_J} \end{pmatrix}^{-1} \cdot \underbrace{\begin{pmatrix} \frac{\partial F_1}{\partial p^m_k} \\ \cdots \\ \frac{\partial F_J}{\partial p^m_k} \end{pmatrix}}_{=-\mathbb{I}_J} \tag{PTR}$$

[75] Because the marginal costs are additively separable we can also define the system as $f(\mathbf{p}, 0, 0) + \mathbf{c} + \tau_0 = 0$.

Because the system of equations is additive in $\mathbf{p^m}$ and τ_0 this simplifies dramatically $\frac{\partial f(\mathbf{p^w},\mathbf{p^m},\tau_0)}{\partial \mathbf{p^m}} = -\mathbb{I}_J$. The pass-through matrix (PTR) is merely the inverse of the matrix whose columns are defined in (B.3).

Implmentation Notes:

1. PyBLP method `compute_passthrough()` will deliver (PTR) (this is very time consuming).
2. PyBLP method `compute_demand_jacobians()` will deliver $\Omega(\mathbf{p^w})$.
3. $\mathcal{H}_m$ is the ownership matrix at the manufacturer level (ie: 1's if both products are owned by Diageo, Bacardi, etc.).
4. $\mathbf{s_t}$ are observed shares and we can plug into (B.1) to get $\mathbf{mc^m}$.
5. Because $\mathbf{mc^m}$ is backed out of (B.1) it is the combination of production costs and federal excise taxes. We never need to separate the two for any counterfactuals.

B.2. Micro Moments

B.2.1. Demographic Interactions

In `PyBLP` Conlon and Gortmaker (2023), all micro moments take the following form, where we match $\overline{v}_m$ with the model simulated analogue. We use the same number of Monte Carlo draws in each market t so that $w_{it} = \frac{1}{I}$ and the general formula simplifies:

$$v_{mt} = \frac{\sum_{i \in I_t} \sum_{j \in J_t \cup \{0\}} s_{ijt} w_{d_m ijt} v_{mijt}}{\sum_{i \in I_t} \sum_{j \in J_t \cup \{0\}} s_{ijt} w_{d_m ijt}}$$

Where $w_{d_m ijt}$ are the survey weights and v_{mijt} is the value. We match the following moments:

1. $w_{dijt} = 1\{j \neq 0\}$ and $v_{mijt} = 1\{\text{Income}_i \in \text{bin}_k\}$ for each market $t \in T$ and "inside" goods only. This allows us to match:

$$\mathbb{P}\left[\text{Income}_i \in \text{bin}_k \mid Purchase\right]$$

2. $w_{dijt} = 1\left\{j \neq 0, x_j = 750mL\right\}$ and $v_{mijt} = 1\{\text{Income}_i \in \text{bin}_k\}$ for each market $t \in T$ and "inside" goods only. This allows us to match:

$$\mathbb{P}\left[\text{Income}_i \in \text{bin}_k \mid 750mL\right]$$

3. $w_{dijt} = 1\left\{j \neq 0, x_j = 1750mL\right\}$ and $v_{mijt} = 1\{\text{Income}_i \in \text{bin}_k\}$ for each market $t \in T$ and "inside" goods only. This allows us to match:

$$\mathbb{P}\left[\text{Income}_i \in \text{bin}_k \mid 1750mL\right]$$

4. $w_{dijt} = 1\{j \neq 0, \text{Income}_i \in \text{bin}_k\}$ and $v_{mijt} = p^w_{jt}$ for each market $t \in T$ and "inside" goods only. This allows us to match:

$$\mathbb{E}\left[p^w_{jt} \mid \text{Income}_i \in \text{bin}_k \text{ and } Purchase\right]$$

We match a different set of values for each income bin. To avoid colinearity (probabilities sum to one) we exclude the middle income bin for the first three sets of moments. We match a different set of moments for each *year* from 2007-2013, rather than each *market* (a quarter). This is because the NielsenIQ Household Panelist data samples different households each year.

These moments are straightforward to calculate from the NielsenIQ Household Panelist data, and don't require any other data sources beyond the NielsenIQ data. The exception is that for each product, NielsenIQ reports the *retail price* and we must find the corresponding *wholesale price* because the model is defined in terms of *Wholesale Demand*.

B.2.2. Aggregate Elasticity

To capture aggregate elasticity we set $w_{dijt} = 1\{j \neq 0, t \in T\}$ and $v_{mijt} = \frac{\hat{s}^*_{ijt}(\mathbf{p}(1+\tau),\widehat{\theta})}{\hat{s}_{ijt}(\mathbf{p},\widehat{\theta})}$ and only use "inside goods" and average over all markets t:

$$v_m(\theta) = \frac{\sum_{t\in T}\sum_{i\in I_t}\sum_{j\in J_t} s_{ijt}(\theta)\cdot v_{mijt}}{\sum_{t\in T}\sum_{i\in I_t}\sum_{j\in J_t} s_{ijt}(\theta)} \approx \frac{1-\frac{1}{T}\sum_t \hat{s}^*_{0t}(\mathbf{p}(1+\tau),\widehat{\theta})}{1-\frac{1}{T}\sum_t s_{0t}(\theta)} = \frac{\sum_t \hat{Q}^*_t(\mathbf{p}(1+\tau),\widehat{\theta})}{\sum_t Q_t(\theta)} = 1+\varepsilon_{agg}.$$

This approximation is valid when $\hat{s}^*_{ijt}(\mathbf{p},\widehat{\theta}) \approx s_{ijt}(\mathbf{p},\theta)$. In practice, this requires an initial consistent estimate (just like two-step GMM or approximations to the optimal instruments), and that we evaluate the model at values of θ which generate choice probabilities close to those generated by our initial choice of $\widehat{\theta}$.

While this seems a bit non-standard, this enables us to incorporate the aggregate elasticity as a "micro moment" in the `PyBLP` software.

We construct a simple estimate of the aggregate elasticity of demand using the data from the NIAAA Surveillance Report (#119) available at `https://pubs.niaaa.nih.gov/publications/surveillance119/CONS20.htm` (Slater and Alpert, 2022). These data are constructed from administrative data provided by Alcohol and Tobacco Tax and Trade Bureau (TTB) which collects taxes from the manufacturer/distillers. Our main quantity measure is the annual gallons of spirits sold in Connecticut (this includes all manufacturer/distillers) and all products shipped to the state (all sizes and categories). These data also report gallons per legal drinking age adult (Per Capita Age 21+).

We also construct a volume-weighted price index using our main estimation sample (quarterly data). And then run the following two regressions:

$$\log Q_t = \beta_0 + \beta_1 \cdot POST_t + \beta_2 \cdot t + u_t$$
$$\log P_t = \gamma_0 + \gamma_1 \cdot POST_t + \gamma_2 \cdot quarter(t) + v_t$$

The idea is to construct a Wald-type estimator of the aggregate elasticity:

$$\frac{\mathbb{E}[\log Q|tax = 1.43] - \mathbb{E}[\log Q|tax = 1.18]}{\mathbb{E}[\log P|tax = 1.43] - \mathbb{E}[\log P|tax = 1.18]} = \frac{\beta_1}{\gamma_1} \approx \varepsilon_{agg}$$

For the annual quantity data we include a time-trend to capture that demand for spirits is rising over time. For the quarterly price data, we include a set of quarter dummies to capture seasonality in prices. In both regressions the main variable of interest is the coefficient on $POST$ an indicator for post July 2011 when the tax changes from \$1.18/L to \$1.43/L.[76] Our regression results are reported in Table B.1.

	Per Capita	Volume	Price	Price
Post	-0.021	-0.02047	0.0525	0.0279
	[0.015]	[0.011]	[0.006]	[0.009]
Trend	0.0202	0.0252		0.0021
	[0.001]	[0.0011]		[0.001]
Quarter FE			x	x
Adj R^2	0.982	0.993	0.848	0.907

Table B.1: Before/After Regression for Tax Change

When obtain nearly identical results when using the *Per Capita* or *Total Gallons* from the TTB data. Including a time trend in the price regression cuts the γ_1 coefficient nearly in half. Our preferred specification is to use the *Per Capita* quantity and data and exclude the time trend from our before/after regression. The resulting coefficient is $\varepsilon_{agg} = -0.41$ (or -0.39 if we use the total volume instead of per capita).

There is a large literature using the NIAAA quantity data to run panel regressions (using data from multiple states and tax changes). A meta-analysis of these results (Wagenaar et al., 2015) suggests that a typical aggregate price elasticity estimate for spirits is around $\varepsilon_{agg} = (-0.8, -0.29)$. If we instead assumed full-pass through we would replace $\gamma_1 = \frac{\tau_1 - \tau_0}{P_0} = 0.0143$ for an aggregate elasticity estimate of $\varepsilon_{agg} = -1.41$, which is much more elastic than most other studies. If we used $\gamma_1 = \frac{\tau_1 - \tau_0}{\tau_0} = 0.21$ which would yield an estimate of $\varepsilon_{agg} = -0.102$ (which is much less elastic than most other studies). Both of these make very strong (and incorrect) assumptions on the (percent on percent) pass-through rate $\frac{\partial \log p}{\partial \log \tau}$.

[76] For the annual quantity data we exclude 2011 from the sample because it is under two different tax regimes

C. Robustness Tests

C.1. Alternative to Figure 3

One concern about Figure 3 is that our choice of price bins may seem arbitrary. A better but more complicated way to address this concern is to rank all vodka products by their national price per liter and compare the CDF of purchases for Connecticut and Massachusetts. If cumulative sales are larger at each national price, then we can say that Connecticut consumes an inferior bundle of vodkas. (We could repeat the exercise for all products, but that might conflate preferences for different categories: Vodka vs. Tequila or Scotch Whisky for "quality"). We plot this in Figure C.1 and show that the bundle in CT nearly FOSD the bundle in MA (except for a few ties).

Figure C.1: CDF of Vodka Consumption by National Average Price Per Liter

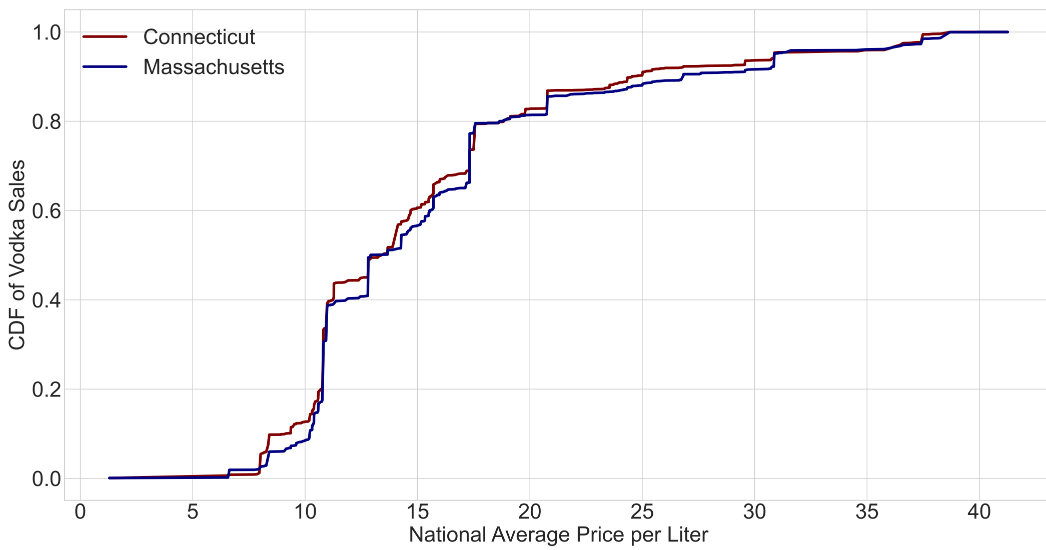

Note: The chart above shows the share of vodka consumption by national price per liter category. A product's national price category is determined using the average price per liter across all Nielsen markets outside of Connecticut designated market areas. For products only sold in Connecticut or Massachusetts the state price is used in place of the national price to calculate price per liter.

C.2. Sensitivity of Demand Estimates

C.2.1. Varying the ρ Parameter

We explore the sensitivity of our parameter estimates by fixing the nesting parameter ρ at different increments between $\rho = 0$ (plain logit) and $\rho = 1$ (all substitution within the nest). We do this by fixing ρ, and re-estimating the other parameters of the model. For this exercise we use the aggregate demand moments $\mathbb{E}[\Delta\xi_{jt} \mid z_{jt}^D] = 0$, and supply moments that come from matching the observed markups (16). For the demand moments, we compute the optimal instruments at our reported parameter estimates and use them for all values of ρ. We also include the micro-moments from (18). What we do not include is a moment that targets the aggregate elasticity. Instead, we report the aggregate elasticity at different values of ρ. The own pass-through rate is relatively unaffected by changes in ρ, but is over-shifted > 1. This is consistent with reduced-form estimates in Conlon and Rao (2020).

We present our results in Table C.1. Because we target the markups with the supply moments, we see there is little variation in the implied Lerner markups. However, larger values of ρ are associated with a higher average own elasticity and a lower aggregate (or market level) elasticity. This is rationalized so that as more consumers stay in the nest, the implied price sensitivity α_i must rise, while the overall taste for alcohol $\beta_{i,0}$ falls. The net result is that as ρ increases, markups stay fixed but consumers are much less likely to substitute to the outside good (which is in its own nest). Larger values of ρ require larger tax increases to hold ethanol consumption fixed. Larger values of ρ also imply somewhat smaller welfare gains, but similar patterns, as consumers become more price responsive and welfare is related to $\frac{1}{\alpha_i}$.

An alternative interpretation of Table C.1 is that it shows us what the aggregate elasticity moment does. It helps us select a value of ρ which governs both the extent to which consumers substitute from one vodka to another, but also the rate at which they substitute to the outside good in response to price increases. Absent these moments (and with some serious caveats about weighting matrices) we might be inclined to choose a small value of ρ suggesting consumers do not substitute within category. If we selected a specification based only on its aggregate elasticity, we might be inclined to choose $\rho \in [0.5, 0.55]$ to match the implied aggregate elasticity of $\varepsilon \approx -0.41$. These pull in different directions, and nearly all of our specifications including the full set of moments imply values of $\rho \in (0.4, 0.5)$ range. If we had second-choice data, we could measure the fraction of consumers whose first and second choice products were both vodka or both whiskey. This kind of variation would be extremely useful in estimating ρ, and would allow us to estimate category specific specific nesting parameters, or separate substitution within category from substitution to the outside good. See Conlon and Gortmaker (2023); Berry and Haile (2022) for a discussion of micro-data and second-choices.

C.2.2. Allowing for Wholesaling Costs

We might worry that the main results are driven by our assumption that in the absence of post and hold policies, that the wholesaler tier becomes perfectly competitive. A reasonable concern is that wholesaling is not costless, and unless wholesalers charge a markup above manufacturer prices, they may not be able to cover the costs of hiring drivers, and operating warehouses. To alleviate these concerns, we set $\mathbf{mc^w} = \mathbf{p^m} + 1$, so that the wholesaler incurs an additional cost of of \$1 per liter both when estimating the demand model, and when computing the counterfactual. We think this is reasonable, as it is in line with the wholesaler margins on the lowest margin items.[77] The exercise is slightly different from Table 7 where we hold the parameter estimates fixed, and allow for a \$1/L wholesale margin.

Qualitatively the patterns in Figure 8 in the main text, and Figure C.3 which allows for the \$1 per liter wholesaling cost, are nearly identical. The relative ranking of various tax instruments, and most importantly the fact that post and hold is clearly dominated by alternative taxes on a competitive market remains the same. Quantitatively, the somewhat higher cost means that the overall level of additional tax revenue that can be generated is reduced slightly, such that we can never increase revenue by more than 400%. The resulting equilibrium prices are highly similar, the main difference being that rather than capturing all of that as additional tax revenue, some must be used to cover the wholesaler costs.

[77] We obtain similar results if we consider larger wholesaling costs of $\mathbf{mc^w} = \mathbf{p^m} + 2$ or $\mathbf{mc^w} = \mathbf{p^m} + 3$.

Table C.1: Sensitivity to different values of nesting parameter ρ

ρ	Own Elas	Agg Elas	Lerner	Outside Div	Objective
0.05	-4.648	-0.672	0.238	0.608	3779.767
0.1	-4.680	-0.649	0.238	0.583	3799.554
0.15	-4.714	-0.626	0.238	0.557	3828.397
0.2	-4.750	-0.600	0.238	0.529	3865.996
0.25	-4.788	-0.574	0.238	0.501	3914.669
0.3	-4.829	-0.546	0.238	0.472	3976.333
0.35	-4.873	-0.516	0.238	0.442	4048.858
0.4	-4.918	-0.486	0.238	0.411	4126.725
0.45	-4.965	-0.453	0.238	0.380	4211.175
0.455	-4.965	-0.451	0.238	0.377	4212.997
0.46	-4.970	-0.448	0.238	0.374	4221.872
0.465	-4.975	-0.444	0.238	0.371	4231.000
0.47	-4.981	-0.441	0.238	0.367	4240.049
0.475	-4.986	-0.437	0.238	0.364	4249.514
0.48	-4.992	-0.434	0.238	0.361	4258.840
0.485	-4.998	-0.430	0.238	0.358	4268.100
0.49	-5.003	-0.427	0.238	0.354	4277.735
0.495	-5.009	-0.423	0.238	0.351	4287.355
0.5	-5.015	-0.420	0.238	0.348	4297.187
0.505	-5.021	-0.416	0.238	0.344	4306.946
0.51	-5.027	-0.413	0.238	0.341	4317.628
0.515	-5.033	-0.409	0.238	0.338	4327.625
0.52	-5.039	-0.406	0.238	0.334	4338.022
0.525	-5.045	-0.402	0.238	0.331	4348.346
0.53	-5.051	-0.399	0.238	0.327	4358.847
0.535	-5.058	-0.395	0.238	0.324	4369.809
0.54	-5.064	-0.391	0.238	0.321	4380.469
0.545	-5.071	-0.388	0.238	0.317	4391.425
0.55	-5.067	-0.385	0.238	0.315	4392.771
0.6	-5.122	-0.349	0.239	0.282	4499.914
0.65	-5.180	-0.311	0.239	0.248	4622.307
0.7	-5.242	-0.271	0.239	0.213	4762.314
0.75	-5.312	-0.229	0.240	0.178	4929.212
0.8	-5.389	-0.187	0.240	0.143	5132.264
0.85	-5.472	-0.143	0.240	0.108	5383.068
0.9	-5.561	-0.098	0.241	0.073	5592.372

Note: We profile demand estimates by varying the level of ρ. This uses the (aggregate) demand moments, the (aggregate) supply moments, and micro-moments from Nielsen Panelist data. It does NOT include the aggregate elasticity moments.
Markups, own elasticity, and outside good diversion are unweighted average over (j, t). Aggregate elasticitiy is market level reduction in purchase volume for a 1% sales tax averaged over markets. Pass-through is own $\frac{\partial p_j}{\partial c_j}$ (dollar for dollar) averaged over products in the final market.
Caution is required comparing GMM objectives across specifications since they have different weighting matrices.
Source: Authors' calculations

Table C.2: Distributional Impacts of Counterfactual Policies with $\mathbf{wc} = 1$

		% Change in CS					
No Change in Ethanol	% Total Revenue	% Overall	Below $25k	$25k-$45k	$45k-$70k	$70k-$100k	Above $100k
Ramsey (Ethanol)	-21.7	20.1	12.9	20.4	23.3	33.3	20.4
Minimum Price	-18.8	19.6	11.2	19.4	21.6	31.0	20.8
Ethanol	203.8	7.1	3.8	6.8	7.0	11.2	7.8
Volume	209.5	6.4	2.6	5.3	5.2	9.5	7.4
Sales+Volume	275.3	-1.6	-0.3	-0.6	-0.9	-0.4	-2.3
Ramsey (Revenue)	288.3	-3.5	-2.6	-2.7	-3.1	-4.4	-3.8
Sales	303.6	-7.5	-2.5	-4.8	-4.9	-7.4	-9.6
-10% Ethanol							
Ramsey (Ethanol)	-4.9	13.8	6.6	9.9	12.3	23.4	15.4
Minimum Price	11.5	13.1	4.1	8.6	9.4	19.4	15.9
Ethanol	239.6	0.4	-4.0	-3.6	-5.2	-0.4	2.8
Volume	246.0	-0.5	-5.5	-5.3	-7.1	-2.3	2.3
Sales+Volume	318.4	-10.1	-8.9	-12.3	-14.2	-13.7	-9.4
Ramsey (Revenue)	329.4	-11.8	-11.5	-14.6	-16.5	-17.7	-10.5
Sales	338.5	-15.7	-10.8	-16.3	-17.9	-20.0	-16.4
+10% Ethanol							
Ramsey (Ethanol)	-43.4	25.5	19.0	30.2	34.2	42.8	24.3
Minimum Price	-41.3	25.3	17.7	29.4	33.2	41.4	24.6
Ethanol	164.4	13.6	11.3	17.3	19.4	22.7	12.3
Volume	169.4	13.0	10.3	16.1	17.9	21.3	12.1
Sales+Volume	224.2	6.7	8.0	11.4	13.0	13.2	4.6
Ramsey (Revenue)	236.3	5.2	6.8	10.4	12.1	10.2	3.0
Sales	260.8	0.6	5.6	6.9	8.5	5.5	-2.9

Note: The table above reports estimates of the impacts of the counterfactual policy alternatives described in Table 5 on tax revenue collected, overall consumer surplus and the distribution of consumer surplus across the five income bins. All effects are reported as percentage changes relative to the PH baseline. The top panel describes the impact of alternative policies that limit ethanol consumption to the same aggregate level as under PH while panels B and C report the effects of alternative policies that reduce and increase ethanol consumption by 10%, respectively. Revenue is calculated as the additional tax revenue raised by the state compared to the existing excise tax collections.
Source: Authors' calculations

Figure C.2: Prices Under PH vs. Other Policy Alternatives with $\mathbf{wc} = 1$

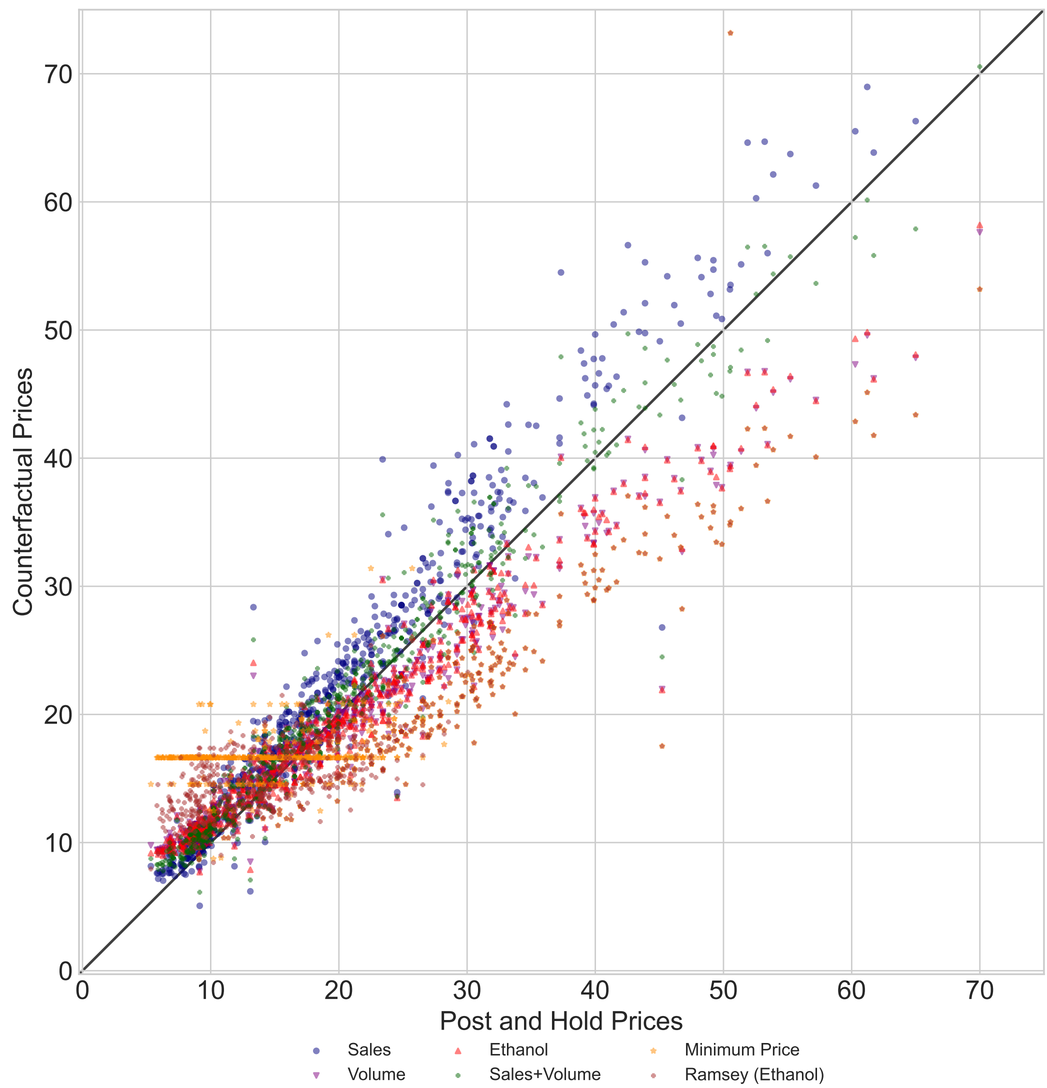

Note: The figure above plots product prices under PH against prices under our counterfactual policy alternatives. In each of our counterfactual scenarios we consider a tax rate that would keep the overall level of ethanol fixed at the status quo. Our taxes follow the definitions in Table 5, and are levied on a competitive market with a \$1/L additional wholesaling cost. The solid black 45 degree line illustrates prices unchanged from PH.

Figure C.3: Consumer Surplus vs. Tax Revenue and Ethanol Consumption Under Alternative Policies with $\mathbf{wc} = 1$

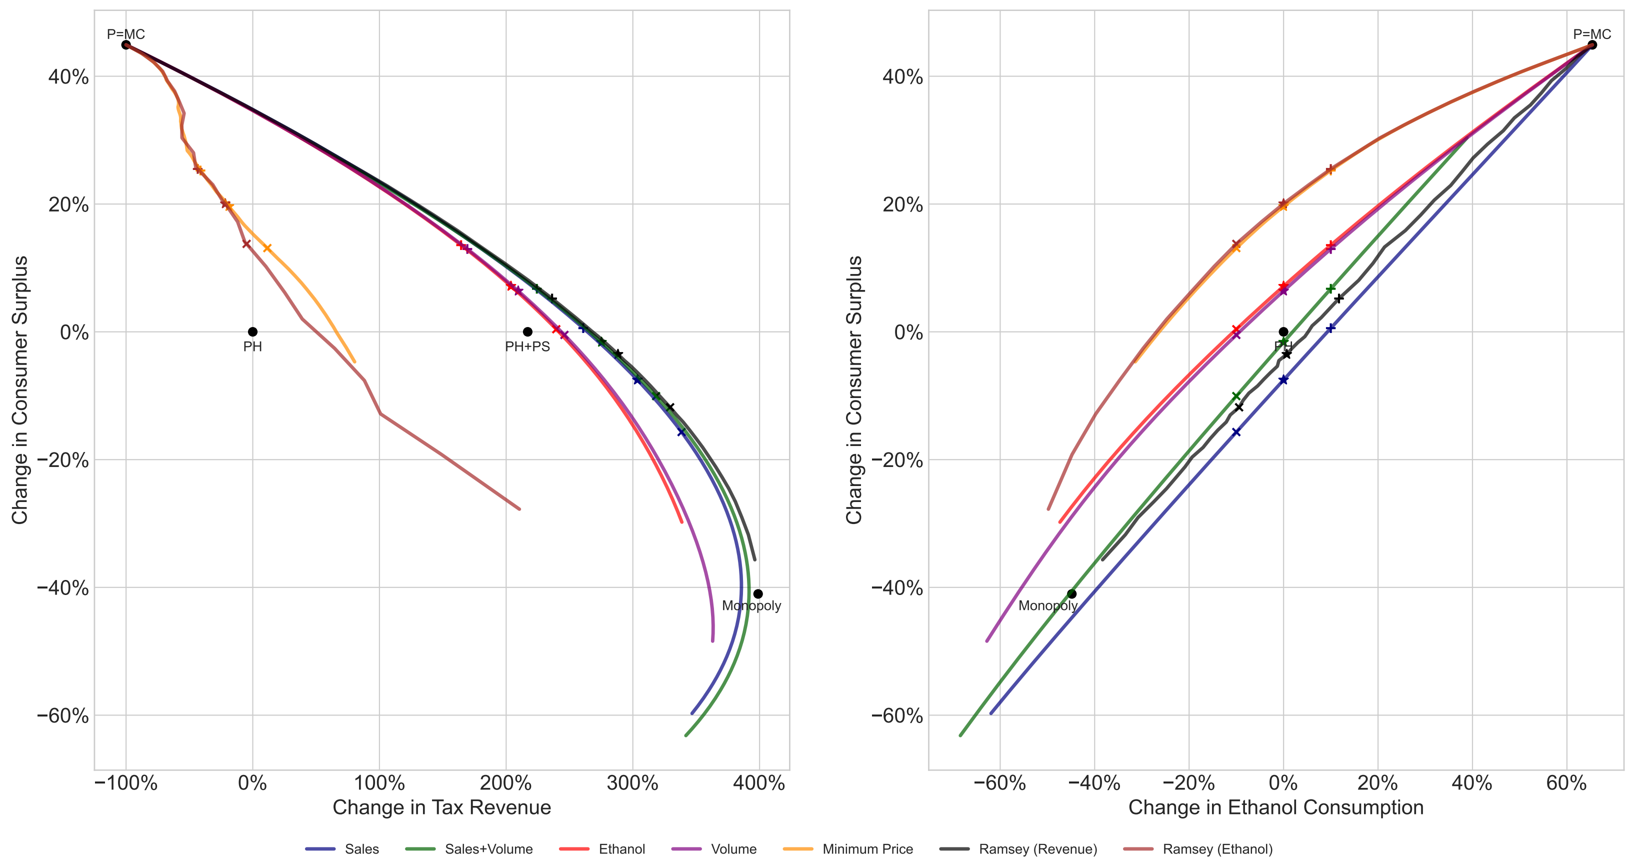

Note: The figure above plots the change tax revenue (left panel) and ethanol consumption (right panel) against the change in consumer surplus for each of the policy alternatives to PH detailed in Table 5 that we consider. The frontiers trace the trade-off between consumer surplus and tax revenue or ethanol consumption for each policy instrument. Stars indicate an aggregate ethanol consumption level equal to total ethanol under PH while x denotes 10% more and less ethanol consumption (in the left panel higher ethanol consumption corresponds to less tax revenue). We also mark competitive prices without taxes (denoted by P = MC), and PH pricing. In the left panel we indicate the revenue generate by existing excise taxes under PH pricing as well as the sum of tax revenue and wholesale profits generated by PH.

D. Panel Data Regressions

D.1. Cross-state Evidence on Consumption Effects of States Ending PH

Theory suggests that PH leads to higher markups, which is supported the price comparisons detailed in Section 4.1. As such it is natural to expect that these higher prices may reduce aggregate alcohol consumption at the state level, which may be a policy objective.

To assess the impact of PH laws on aggregate alcohol consumption, we assemble a panel of annual state data measuring wine, beer, and spirits consumption, demographic characteristics. These data are drawn from the National Institute on Alcohol Abuse and Alcoholism (NIAAA) *U.S. Apparent Consumption of Alcoholic Beverages*, which tracks annual consumption of alcoholic beverages for each state. We use the timing of when different states terminated PH laws (often as the result of lawsuit) to measure the association between regulation and alcohol consumption. Table D.1 reports PH termination dates. This table matches Cooper and Wright (2012), who also run a similar panel regression to the one we describe below (and obtain similar results):[78]

Table D.1: States with Post and Hold Laws

	Wine	Beer	Spirits
Connecticut	Y	Y	Y
Delaware	End 1999	End 1999	End 1999
Georgia	N	Y	Y
Idaho	Y	Y	N
Maine	Y	Y	N
Maryland	End 2004	End 2004	End 2004
Massachusetts	End 1998	End 1998	End 1998
Michigan	Y	Y	Y
Missouri	Y	N	Y
Nebraska	End 1984	N	End 1984
New Jersey	Y	Y	Y
New York	Y	Y	Y
Oklahoma	End 1990	End 1990	Y
Pennsylvania	N	End 1990	N
South Dakota	Y	N	Y
Tennessee	N	Y	N
Washington	End 2008	End 2008	N
West Virginia	N	N	Y

Note: The table above lists all states that have or have repealed PH regulations and details the types of alcoholic beverages covered by PH rules. Y denotes a state and beverage category with PH provisions. N denotes a state and beverage category was never subject to PH laws. The year of repeal is denoted for states that ended their PH regulations. No state adopted PH after the start of sample period, 1983. This table is a reproduction of Table 1 of Cooper and Wright (2012).

These state panel regressions are similar to those of Cooper and Wright (2012) and have the form:

$$Y_{it} = \alpha + \beta PH_{it} + X_{it}\gamma + \delta_t + \eta_i + \epsilon_{it} \tag{D.1}$$

The dependent variable is the log of apparent consumption per capita, where consumption is in

[78] In contrast, Saffer and Gehrsitz (2016) find a null effect of PH on prices, but rely on ACCRA data which tracks the price of only one brand each for: beer (Budweiser 6-pack), wine (Gallo Sauvignon Blanc) and distilled spirits (J&B Scotch).

ethanol-equivalent gallons and the relevant population is state residents age 14 and older. PH_{it} is a dummy variable equal to one if state i has a PH law in place at time t; X_{it} is a vector of control variables; and δ_t and η_i are time and state fixed effects, respectively. The coefficient of interest, β, describes the reduction in alcohol consumption associated with PH laws.

We report the results in Table D.2. The specification of column 1 includes only time and state fixed effects while column 2 adds state-specific linear time trends. Accounting for state differences in underlying consumption trends attenuates the wine coefficient, rendering it statistically insignificant, but increases the magnitude and precision of beer and spirits coefficients and makes them statistically significant.

The identifying variation comes from the handful of states ending their PH requirement. There are a number of reasons we should remain cautious about taking the regression estimates too seriously. The first is that we don't know why states terminate PH, though in several cases it was the result of losing a lawsuit rather than through the legislative process. The bigger issue is that when states eliminate PH, they tend to also change tax rates, and liberalize other laws regarding the distribution and sale of alcoholic beverages. We may wrongly attribute other factors (ending prohibitions on Sunday sales, etc.) to eliminating PH.

Though PH appears to reduce alcohol consumption, proponents of PH also mention the desire to protect small retailers from larger chains such as Wal-Mart and Costco. In Appendix D.2 we assess the impact of PH on retail liquor establishments. Again, we caution against taking these results too seriously, but they suggest that PH is also associated with: lower employment in the liquor retail sector, fewer retail stores per capita, but a larger share of very small (1-5 employee) establishments. This prevents us from calculating a "cost per job" measure because ending PH is associated with increased employment in the sector.

D.2. Cross-state Evidence on Employment Effects of States Ending PH

Advocates for PH argue that the regulation benefits small retailers by ensuring that they pay the same wholesale prices as large retailers such as Costco or BevMo. If PH does indeed protect small retailers, PH states like Connecticut should be home to more small-scale retail establishments. The impact of PH on employment and the total number of establishments, however, is less clear. While under PH small retailers enjoy uniform pricing, these uniform prices are the higher prices that result from non-competitive wholesaler pricing behavior. Having more small retailers in a retail sector that faces lower margins due to high wholesale prices could lead to either more or fewer establishments that overall employ more or fewer workers.

Table D.3 provides some empirical evidence regarding these questions. The regressions presented in Table D.3 are of the same form as equation D.1 and describe the impact of PH spirits regulations on three different outcomes: share of small retail establishments, log employment in the liquor retail sector, and log liquor stores per capita.[79]

The uppermost panel of Table D.3 examines the impact of PH regulations on the prevalence of small liquor retailers (that is, establishments with between one and four employees). Column one uses only only data from 2010 and includes demographic controls—state population and median income—and finds a marginally significant positive relationship between PH and share of small liquor retail establishments. Columns two through four use the full panel from 1986 through 2010. Adding state and year fixed effects does not yield a significant coefficient, as shown by column two.

[79]Panel data describing state liquor retail establishment counts and employment come from the Census County Business Patterns for 1986 through 2010.

Table D.2: Post and Hold Laws and State Alcohol Consumption

	(All)	(All)	(All)	(PH only)	(PH NE)
Wine					
PH	-0.0545***	-0.0215	-0.0197	-0.0277	-0.00360
	(0.0183)	(0.0192)	(0.0192)	(0.0182)	(0.0356)
R^2	0.965	0.984	0.984	0.985	0.988
Beer					
PH	-0.0155	-0.0218**	-0.0207**	-0.0192**	-0.0297**
	(0.0113)	(0.00968)	(0.00959)	(0.00859)	(0.0134)
R^2	0.891	0.968	0.968	0.954	0.980
Spirits					
PH	-0.00702	-0.0731***	-0.0725***	-0.0665***	-0.0851***
	(0.0175)	(0.0183)	(0.0181)	(0.0175)	(0.0279)
R^2	0.950	0.982	0.982	0.976	0.984
Year FE	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y
State Time Trends	N	Y	Y	Y	Y
Demog. Controls	N	N	Y	Y	Y
PH States	N	N	N	Y	Y
NE States	N	N	N	N	Y
Observations	1,428	1,428	1,428	532	168

Note: The table above presents coefficients from regression D.1. The outcome of interest is the log of apparent consumption per capita, where consumption is in ethanol equivalent gallons and the relevant population is state residents age 14 and older. Column 1 only includes state and time fixed effects. Column 2 adds state-specific time trends while column 3 also includes state demographic controls. Column 4 limits the sample to states that have had PH laws. Column 5 restricts the sample further to only northeastern states that once had PH laws. The alcohol consumption data are from the National Institute on Alcohol Abuse and Alcoholism, which is part of the National Institutes of Health; the demographic information comes from the Census Bureau's intercensal estimates. Standard errors in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Column three adds state-specific time trends, which control for changes in spirits consumption that vary by state. Adding these additional controls reveals that states with PH regulations do in fact have a larger share—4.8 percentage points larger—of small retail establishments. Dropping all states outside of the northeast does not substantively affect the coefficient but increases the precision of the estimate.

The middle panel examines the impact of PH regulations on employment in the alcohol retail sector. The dependent variable is the log of employment in the liquor retail sector per capita age 14 years and older. Looking at data from only 2010 does not suggest a statistically significant relationship between employment and PH laws. Adding year and state fixed effects as shown in column 2 reveals that states with PH laws actually have lower per-capita liquor retail employment. Including state time trends reduces the magnitude and precision of the coefficient from -1.762 (0.198) to -0.497 (0.239). Focusing on northeastern states (column 4) does not have an appreciable further impact on the estimates, though the estimate is less precise.

The bottom panel assesses how the number of establishments per capita is affected by PH regulations. As in the employment panel, examining the 2010 data alone does not suggest a statistically significant relationship between number of retailers and PH laws. Column two uses the full panel with state and time fixed effects, yielding a significant and negative coefficient. Controlling for state time trends reduces the coefficient to -0.608 (0.0914). As in the other panels, examining only northeastern states doesn't appreciably change the coefficient.

Table D.3: Post and Hold Laws and Alcohol Retailing

	2010 Only	All	All	Northeast
Share of 1-4 Employee Retailers	0.0728*	0.0339	0.0477*	0.0472**
	(0.0432)	(0.0209)	(0.0262)	(0.0227)
R-Squared	0.144	0.867	0.940	0.962
Log(Alcohol Employment/Pop 14+)	0.452	-1.762***	-0.497**	-0.422*
	(0.336)	(0.198)	(0.239)	(0.223)
R-Squared	0.064	0.467	0.740	0.821
Log(Liquor Stores Per Capita)	0.344*	-1.335***	-0.608***	-0.515***
	(0.204)	(0.0866)	(0.0914)	(0.103)
R-Squared	0.128	0.855	0.954	0.963
Obs	51	1,275	1,275	300
Demographic Controls	Y	Y	Y	Y
State FE	N	Y	Y	Y
Year FE	N	Y	Y	Y
State Specific Trends	N	N	Y	Y

Note: The table presents coefficients from regression D.1 where the outcome of interest is the share of retailers with 1-4 employees in the uppermost panel, the log of employment in the liquor retail sector per capita in the middle panel, and log of liquor stores per capita in the bottom panel. The reported coefficients correspond to a binary variable that is equal to one when spirits are subject to PH regulations. Column 1 uses only data from 2010 and includes demographic controls. Columns 2 through 4 use the full 1986 - 2010 panel. Column 2 adds state and year fixed effects. Column 3 adds state specific time trends and column 4 limits the sample to only northeastern states. Standard errors in parentheses.
*** p<0.01, ** p<0.05, * p<0.1