The Effect of Alcohol Sales Taxes on Off-Site Alcohol Purchases*

Ana Paula Milán Hinostroza†

May, 2026
Preliminary and Incomplete.
Please do not cite or circulate.

Abstract

While most research on alcohol taxation focuses on volume-based excise taxes, much less is known about the effects of sales taxes, which are applied as a percentage of retail prices and may differ in salience and incidence. This paper evaluates the impact of Maryland's 2011 three percent alcohol sales tax on household alcohol purchases using NielsenIQ Consumer Panel data and a synthetic difference-in-differences design. I find that the tax reduced monthly ethanol purchases per adult by approximately nine percent among households that purchase alcohol, corresponding to about 1.7 fewer standard drinks per adult per month. Declines are observed across beer, wine, and spirits, with no evidence of substitution toward any single beverage category. Reductions are present across the purchasing distribution, with the largest absolute declines among high-purchasing households, and broadly similar patterns are observed across demographic groups. Additional analyses provide little evidence that border-county households offset the tax through cross-border shopping. These findings provide new evidence on the behavioral effects of sales taxes and suggest that price-based alcohol policies can meaningfully reduce alcohol purchases.

*Acknowledgments. I would like to express my gratitude to Brady Horn, Xiaoxue Li, and Dave Dixon for their mentorship and guidance. I also thank the Western Economic Association International Graduate Student Workshop for valuable feedback and discussion. The opinions and conclusions presented in this paper are solely my own, and I take full responsibility for any errors.

The results are 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 researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing or preparing the results reported herein.

†The University of New Mexico. apaulamilan@unm.edu

1 Introduction

Excessive alcohol consumption is one of the leading preventable causes of death in the United States (U.S.) and is associated with over $249 billion in medical costs and lost productivity (Sacks et al., 2015). It accounts for approximately 385 deaths per day and more than 3.59 million years of potential life lost annually (National Center for Drug Abuse Statistics, 2024). A substantial body of evidence shows that alcohol excise taxes can effectively reduce consumption and mitigate the health and social harms of excessive drinking (Wagenaar et al., 2009), but little is known about the effects of alcohol sales taxes.

Alcohol can be taxed either by volume or by price. Volume-based excise taxes are levied per unit of alcohol (e.g., per gallon or liter), while price-based sales taxes are applied as a percentage of the beverage's retail price.[1] Although both taxation methods aim to regulate alcohol consumption, they operate through distinct mechanisms. Excise taxes are fixed per unit of alcohol and typically embedded in shelf prices, whereas sales taxes are applied at the point of sale and automatically adjust with inflation. These differences have implications for both the economic incidence of the tax and its behavioral effects.

There is broad evidence that alcohol excise taxes reduce alcohol consumption (see Wagenaar et al. (2009) for a review of the literature), but relatively little empirical research on the impact of alcohol sales taxes on consumption, as most states do not impose additional sales taxes on alcoholic beverages. Research suggests that excise taxes may be more effective than sales taxes, as the latter are less salient to consumers, being visible only at the point of sale (Chetty et al., 2009).

From standard economic theory, both excise and sales taxes increase the full price of alcohol and *should* reduce consumption. From a salience standpoint, sales taxes may be less salient because they are added at checkout rather than incorporated into posted prices; however, lower salience does not imply the absence of a behavioral response. If consumers are well informed or learn about the tax over time, sales taxes may still reduce alcohol consumption through the standard price mechanism. In particular, in settings with repeated purchasing decisions, consumers may internalize the higher final prices and adjust their behavior accordingly.

To examine this, I focus on Maryland's 2011 increase in the alcohol sales tax from 6% to 9%, the state's first such hike in four decades. Anecdotal evidence suggests the change

[1] Sales taxes are also known as ad valorem consumer taxes.

was adopted primarily to align Maryland's tax rate with those of neighboring states, which supports treating the increase as plausibly exogenous to underlying alcohol consumption trends. The policy, implemented on July 1, 2011, applied to both on-premise purchases (e.g., bars and restaurants) and off-premise purchases (e.g., liquor stores and supermarkets).[2]

Using the NielsenIQ Consumer Panel data, I construct a balanced household-level panel from July 2009 to June 2013. To estimate the tax's impact, I compare changes in ethanol ounces purchased per adult per household in Maryland relative to a weighted set of control households from other states, using a synthetic difference-in-differences approach.[3] I therefore estimate the effect of the tax on alcohol purchases rather than consumption.[4]

My preferred specification estimates that the tax increase led to a 9.2% reduction in ethanol ounces purchased per adult per month, equivalent to approximately 1.04 fewer ounces of ethanol per adult, or roughly 1.7 fewer standard drinks. There is no evidence of substitution across beverage categories, including beer, wine, and spirits. The reduction is observed across the purchasing distribution, with the largest absolute declines occurring among higher-purchasing households.

Across demographic groups, the results suggest broadly similar patterns, with most groups showing declines in alcohol purchases following the tax increase. Although some subgroup estimates differ in direction, inference is limited due to smaller sample sizes within these groups and should be interpreted cautiously. Finally, responses among households in Maryland's border counties are not significantly different from those in non-border counties.

I also look at the price paid per ounce of ethanol. Because this measure reflects both prices and purchasing behavior, it is important to distinguish between months with and without purchases. When averaging across all household-months, the price per ounce appears to decline after the tax increase, driven in part by reductions in purchase frequency and the inclusion of more zero-purchase months. However, when conditioning on months where households make alcohol purchases, the price per ounce increases, consistent with the mechanical effect of the tax. This suggests that the apparent decline in average prices reflects changes in purchasing behavior and selection into purchase months, rather than a true decrease in the price per ounce of ethanol faced by purchasing households.

[2]On-premise sales refer to alcohol purchased and consumed at the point of sale, such as in bars, restaurants, or clubs, while off-premise sales refer to alcohol purchased for consumption elsewhere, such as from liquor stores, supermarkets, or wholesalers.

[3]I focus on ethanol ounces rather than beverage volume because ethanol is the component of alcohol responsible for its harmful effects.

[4]The data do not include individual-level alcohol consumption, only purchases.

Overall, the findings suggest that, like excise taxes, alcohol sales taxes can reduce alcohol purchases. In settings with repeated purchasing decisions, consumers appear to respond to the overall price they face rather than to the specific tax instrument generating that price, which helps explain why sales taxes can produce effects comparable to those found in the excise tax literature. This reduction in purchases is a first-order impact that may lead to declines in alcohol-related externalities, including adverse health outcomes, crime, and traffic fatalities. Taken together, these findings highlight the potential for sales taxes to complement excise taxes in alcohol policy design. In addition, because sales taxes automatically adjust with inflation, they may retain their real value over time without requiring legislative updates.

2 Background

2.1 Alcohol taxation in the U.S.

Alcohol taxation is a common economic policy tool aimed at reducing excessive alcohol consumption and mitigating alcohol-related harms. In the U.S., alcohol taxes are imposed at both the federal and state levels. At the federal level, taxes are primarily excise taxes, levied based on alcohol volume rather than price. Excise tax rates vary by beverage type, with spirits taxed at the highest rate per gallon, followed by wine and beer. Federal excise rates have remained unchanged since 1991, though certain producers receive tax reductions based on annual production volume, leading to a decline in the real value of these taxes over time. Although excise taxes are usually levied on producers or retailers, they are generally passed on to consumers in the form of higher prices, with many estimates suggesting a pass-through rate greater than one (Kenkel, 1993, 2005; Nelson and Moran, 2019; Shrestha and Markowitz, 2016; Young and Bielińska-Kwapisz, 2002).

At the state level, alcohol taxation varies depending on the type of beverage, the point of sale (on-premise vs. off-premise), and the distribution system. In some states, known as "monopoly states," alcohol retail sales are directly controlled by the government, generating revenue through state-operated outlets rather than through traditional tax mechanisms. However, most states impose their own excise taxes, and some also apply additional sales taxes on alcohol. According to the Alcohol Policy Information System (APIS), alcohol sales taxes are defined as the difference between a state's alcohol retail sales tax and its general sales tax, reflecting the extra tax burden placed on alcohol compared to other goods (National

Institute on Alcohol Abuse and Alcoholism, 2024). Because sales taxes are calculated as a percentage of the product's price, they automatically increase with inflation, and result in higher tax burdens on more expensive alcoholic beverages.

Alcohol sales taxes are far less common than excise taxes, with only a few states having implemented them. In the past two decades, only Maryland, Washington, D.C., Kansas, Oklahoma, Texas, and South Carolina have applied additional sales taxes on alcohol. However, these policies differ substantially in scope and structure. Maryland is the only state to introduce a new, standalone alcohol sales tax that applies uniformly across all beverage types (beer, wine, and spirits) and across both on- and off-premise sales.[5] In contrast, other states' policies are either narrower in application or reflect different regulatory changes. For example, Texas applies sales taxes only to on-premise alcohol sales, South Carolina's policy applies only to spirits, and Oklahoma's changes are limited to beer (on-premise) and wine (on- and off-premise), with no comparable adjustment for spirits. Washington's policy change reflects a shift from being a control state rather than a comparable tax introduction, while D.C. and Kansas adjusted existing sales taxes with both increases and decreases over time, complicating a clear treatment definition. These differences limit comparability across states and make Maryland uniquely well-suited for isolating the effects of a comprehensive alcohol sales tax.

2.1.1 The 2011 Maryland Alcohol Sales Tax Increase

Maryland presents a unique case as one of the few states to ever implement an additional sales tax for alcohol. Prior to 2011, the state's alcohol tax structure had remained unchanged for decades. Maryland had not adjusted its excise tax on distilled spirits since 1955 or on beer and wine since 1972. As a result, the real value of these taxes declined substantially over time, making alcohol increasingly affordable. By 2011, Maryland's excise tax per drink was estimated to be less than two cents, among the lowest in the country. Compared to its five bordering jurisdictions, Maryland had the lowest distilled spirits excise tax and the second-lowest beer and wine excise tax (Esser et al., 2016).

On July 1, 2011, Maryland increased its state alcohol sales tax from 6% to 9%, a 50% increase (MD.§11–104(g)). Unlike excise taxes, which had remained unchanged for decades (MD.§5–105), this policy linked the tax burden directly to alcohol prices and applied to

[5] Tables A1 and A2 in the Appendix summarize state-level changes in alcohol excise and sales taxes during the main sample data period.

both on-premise and off-premise sales. This tax increase was the result of nearly a decade of advocacy rather than a response to short-term shifts in alcohol consumption. As early as 2002, policymakers introduced proposals to increase alcohol and tobacco taxes as revenue sources for state healthcare expansion plans, but these efforts initially failed (Ramirez and Jernigan, n.d.). Unlike tax changes aimed at reducing excessive consumption, Maryland's policy was primarily framed as a fiscal measure to align with neighboring states and support healthcare funding. This historical context suggests that the tax increase was not directly driven by fluctuations in the state's alcohol consumption levels, reinforcing its plausibly exogenous nature.

2.2 Alcohol taxes and consumption

There is a broad empirical literature examining the impact of alcohol taxes on consumption. Economic theory suggests that taxation increases the price of a good, which in turn reduces its consumption (Maldonado-Molina and Wagenaar, 2010; Xuan et al., 2015). Alcohol is generally considered a normal good (Hanson and Sullivan, 2016; Meng et al., 2014; Wagenaar et al., 2009), though its demand is inelastic, meaning consumption does not decrease in proportion to price increases (Ayyagari et al., 2013; Gehrsitz et al., 2021; Pryce et al., 2019). Price responsiveness varies across both beverage types (Clements et al., 2022; Gehrsitz et al., 2021; Son and Topyan, 2011) and consumer characteristics (Pryce et al., 2019; Wagenaar et al., 2009).

Among beverage types, beer tends to be less elastic than wine and spirits (Clements et al., 2022), suggesting that demand for spirits and wine is more sensitive to price changes (Gehrsitz et al., 2021; Son and Topyan, 2011). One possible explanation for these differences is substitution effects. For example, Gehrsitz et al. (2021) find that after the 2009 Illinois alcohol excise tax increase, consumers shifted to lower-priced alcohol, increasing their beer purchases and partially offsetting the overall decline in ethanol consumption. This suggests that when the price of one type of alcohol rises, consumers may switch to cheaper alternatives within the alcohol category to maintain their consumption levels.

Consumer responses to alcohol price changes may also vary across individuals and demographic characteristics (Meng et al., 2014). For example, heavy and binge drinkers tend to exhibit relatively inelastic demand, often substituting toward lower-cost options rather than reducing overall consumption (Pryce et al., 2019; Wagenaar et al., 2009). However, the empirical evidence on how price responses differ across other socioeconomic characteristics, such as gender and race or ethnicity, is more mixed (Kilian et al., 2023).

2.2.1 Contributions to the Literature

Most empirical studies examining the impact of alcohol taxes on consumption have focused on excise taxes. While both excise and sales taxes increase alcohol prices and can reduce consumption, they are not necessarily equivalent in terms of behavioral responses. In standard models with fully informed consumers, both taxes operate through the same price mechanism and should generate similar reductions in consumption. However, in practice, they differ along several dimensions that may affect behavior. Excise taxes are typically embedded in posted prices, while sales taxes are added at checkout and may be less salient (Chetty et al., 2009). In addition, excise taxes are levied in levels (per volume), whereas sales taxes are proportional to price, which can generate different incentives for substitution across products.

This paper contributes to the literature by offering the first economic empirical analysis of how a sales tax affects household-level alcohol purchases, using the NielsenIQ Consumer Panel and Maryland's 2011 alcohol sales tax increase as a case study. While previous research has evaluated the Maryland reform, earlier studies have largely focused on downstream health or behavioral outcomes, rather than purchases. For instance, Lavoie et al. (2016) document a decline in alcohol-positive drivers; Staras et al. (2016) find reductions in gonorrhea rates; and Smart et al. (2018) report fewer alcohol-related disciplinary actions among college students. These studies suggest a reduction in consumption and imply that the tax was at least partially passed through to prices, but they do not directly observe alcohol purchases.

Only one study, Esser et al. (2016), analyzes sales data following the tax. They find modest declines in per capita alcohol sales across all categories but estimate post-tax effects using a mixed-effects model without a control group. Their reliance on county-level data also precludes any assessment of household-level heterogeneity. By contrast, my analysis uses a synthetic difference-in-differences design with a flexible control group and household-level panel data, allowing for a causal interpretation of the tax's impact on purchasing behavior and rich heterogeneity analysis.

This study also builds on recent research using NielsenIQ Consumer Panel data. Saffer et al. (2024) study the 2009 Illinois alcohol excise tax increase using a difference-in-differences model, where the control group is constructed via a synthetic control method (SCM). Their estimation approach, following Donald and Lang (2007), aggregates monthly differences between Illinois (the treatment state) and the control aggregate from the SCM, reducing the effective sample size to 43. They find that both heavy and moderate drinkers, as well as low-income drinkers, reduce ethanol purchases in response to the tax. Their findings also indicate substitution effects, as heavy and moderate drinkers, along with middle- and high-income individuals, shift their purchases to avoid higher ethanol prices, while low-income drinkers end up paying more per unit of ethanol.

My research builds on this study but differs in two key ways. First, I examine an alcohol sales tax rather than an excise tax, which can generate different behavioral responses. Second, my identification strategy differs. I use a synthetic difference-in-differences model, which leverages a balanced panel of households from July 2009 to June 2013 and includes household and year-month fixed effects.

Unlike Saffer et al. (2024), who aggregate data to the state level, I exploit within-household variation in alcohol purchases, allowing for a more granular analysis. Similar to their findings, my results show that the tax reduces overall ethanol purchases. Their main objective is also to look at differential responses by heavy drinking and income levels, rather than doing this, I further look at demographic group differences (although I broadly do not find differential responses). I also compare Maryland border and non-border households and find no meaningful differences in their responses to the tax. By addressing these gaps, this study provides new empirical evidence on how sales taxes affect alcohol purchasing behavior.

3 Data

3.1 NielsenIQ Consumer Panel

To estimate the impact of Maryland's 2011 alcohol sales tax on alcohol purchases, I use household-level data from the NielsenIQ Consumer Panel. This dataset is particularly well suited for this analysis because it captures both overall changes in alcohol purchases and variation in responses across different types of consumers. The panel tracks a rotating sample of 40,000 to 60,000 households annually, with participating households recording their purchases using in-home scanners. Households are included in the sample only after consistently reporting purchases for at least 12 months. The dataset covers a wide range of consumer purchases across all retail channels and includes demographic information on household size, income, presence of children, and select characteristics of household heads, such as age and education.[6] [7]

My analysis focuses specifically on alcoholic beverage purchases. Because the data only capture off-premise alcohol sales, this study effectively examines the tax's impact on alcohol bought for home consumption rather than on-premise sales (e.g., restaurants and bars). Alcohol purchases are aggregated to the household-month level over 48 months, from July 2009 to June 2013. Since the tax was implemented in July 2011, this allows for approximately two years of pre- and post-tax observations. Data from earlier years are excluded to mitigate potential biases from the sharp decline in alcohol purchases during the 2007–2008 U.S. financial crisis, which may have had differential effects across states.[8] I will also show in Section 5 that results remain consistent when using less restrictive time cutoffs. To construct the analytic sample, I apply further restrictions at both the state and household levels.

At the state level, states that experienced any changes in alcohol sales or excise taxes during the study period are excluded from the sample, as these policy changes could affect alcohol purchases independently of Maryland's 2011 tax, making those states unsuitable comparison units. Specifically, Kansas and D.C. are excluded due to sales tax changes, while New York, Connecticut, North Carolina, Illinois, New Jersey, and Washington are excluded because they implemented significant alcohol excise tax increases.[9]

At the household level, I restrict the sample to households that participated for the full 48-month period surrounding the sales tax implementation to exploit household-level heterogeneity and control for time-invariant unobserved household characteristics. Households that changed their state of residence while in the sample are excluded to maintain consistent exposure to the policy, as their alcohol purchasing behavior could be influenced by differences in state-level policies or economic conditions. In addition, to define the relevant population for this analysis, I limit the sample to households that made at least one alcohol purchase in the 12 months preceding the tax increase, so that the analysis focuses on the tax's impact on active alcohol consumers rather than non-drinking households.[10] The final sample includes

[6]Retail channels include grocery stores, liquor stores, drug stores, mass merchandise retailers, superstores, club stores, and convenience stores. Purchases are recorded as long as the consumer scans them.

[7]Demographic characteristics are aggregated to the household level. For households with both a male and female head, I use the higher level of education, age, and other characteristics reported among the heads. For households with a single head, I use that individual's information. Household heads are 25 years and older.

[8]Maryland raised its general sales tax from 5% to 6% in 2007; while this change was not specific to alcohol, it could still have influenced alcohol purchases and introduced additional bias into the estimates.

[9]The specific tax changes are presented in Tables A1 and A2 in the appendix.

[10]Figure 6 includes models with the full sample which include non-drinking households.

households from 41 states, including Maryland, and consists of 10,653 households, with 205 in the treated group and 10,448 in the control group.[11]

3.1.1 Alcohol Measures

I evaluate both the quantity of alcohol purchased and the effective price paid for alcohol purchased for home consumption. To evaluate quantity, rather than using total ounces of alcohol purchased, I focus on ethanol ounces, which directly measure the alcohol content, the component with direct health implications. I calculate ethanol purchases by converting beverage ounces into ethanol equivalents, assuming standard alcohol contents of 5% for beer, 13% for wine, and 40% for spirits. By multiplying these percentages by the respective ounces purchased, I obtain a measure of monthly ethanol purchases per household.[12] [13] I then divide total monthly ethanol ounces by the number of adults in the household to create ethanol ounces per adult. In addition to total ethanol purchases, I construct analogous per-adult measures separately for beer, wine, and spirits.

To examine the effective price households paid for alcohol, I construct several price-related measures. Following Saffer et al. (2024), the first measure is the household's monthly alcohol spending, measured in 2013 dollars, divided by total ethanol ounces purchased in that month. I also construct beverage-specific unit values for beer, wine, and spirits by dividing category-specific spending in 2013 dollars by category-specific ethanol ounces purchased. These unit values are defined only in months with positive purchases of the relevant beverage. Because NielsenIQ expenditures do not include sales taxes, I adjust the overall price-per-ethanol measure and the beverage-specific unit values by increasing Maryland households' post-tax prices by 3%; thus, the price measures are tax inclusive. A limitation of this adjustment is that Maryland households may purchase alcohol outside the state. In that case, the measure would assign the Maryland tax to some purchases that may not have been taxed in Maryland. I return to this concern in Section 5 by examining whether border-county households respond differently in both quantities and effective prices. All price measures are expressed in 2013 dollars.

[11] NielsenIQ does not include data on Hawaii and Alaska.

[12] Formally, household-month ethanol ounces equal 0.05 times beer ounces plus 0.13 times wine ounces plus 0.40 times spirits ounces.

[13] Appendix Table A4 reports results using total volume (in ounces) rather than ethanol-equivalent ounces.

3.1.2 Household Classification

Since a key objective of this study is to examine behavioral responses across different consumer characteristics, I study heterogeneity by baseline purchasing intensity and household demographics. The primary heterogeneity analysis focuses on pre-tax alcohol purchasing intensity but I also examine differences by income, household composition, race, and education of the household head.

Purchasing patterns: I classify households by their baseline alcohol purchasing intensity using their average monthly ethanol ounces per adult in the pre-treatment period. Because isolated bulk purchases can mechanically inflate a household's pre-treatment average, I first identify anomalous purchasers, or stockpilers, as households with at least one pre-treatment month in which ethanol purchases exceed 200 ounces per adult.[14] These households are excluded from the purchaser-intensity classification but are otherwise retained in the sample. The purpose of this approach is to capture persistent differences in baseline purchasing intensity rather than temporary purchase spikes. For example, a household that makes an unusually large purchase before the tax and little afterward could otherwise be misclassified as a high-intensity purchaser, even though this pattern may reflect stockpiling rather than sustained high alcohol demand.

Among non-stockpiler households, I divide the distribution of pre-treatment average ethanol purchases per adult into terciles. Households in the bottom tercile are classified as low-intensity purchasers, those in the middle tercile as moderate-intensity purchasers, and those in the top tercile as high-intensity purchasers.[15] This classification is fixed at the household level and carried forward across all months of the panel. In the pre-treatment period, low-intensity purchasers buy about 0.6 ounces of ethanol per adult per month, moderate-intensity purchasers about three ounces, and high-intensity purchasers about 23 ounces.

[14] This is equivalent to more than 333 standard drinks per adult in a month, using 0.6 ethanol ounces per standard drink.

[15] Appendix Table A6 reports robustness to alternative purchaser-intensity cutoffs.

4 Methods

4.1 Identification strategy

My identification strategy exploits variation in household alcohol purchases before and after Maryland's 2011 alcohol sales tax increase. Specifically, I compare monthly ethanol ounces purchased by households in Maryland with those purchased by households in other states. Although the dataset contains many households, there is only one treated state, which makes conventional difference-in-differences (DiD) methods unsuitable as standard errors can be severely underestimated when there is a single treated unit or few aggregate treatment clusters. In such settings, a weighted combination of untreated units provides a more credible counterfactual than any single comparison unit (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Abadie, 2021).

To address this limitation, I use the synthetic difference-in-differences (SDiD) estimator, which combines the strengths of DiD and synthetic control method (SCM) approaches (Arkhangelsky et al., 2021). Whereas SCM is designed for cases with a single treated unit, constructing an optimally weighted combination of untreated units that closely reproduces the pre-treatment trajectory of the treated unit, SDiD extends this framework by incorporating unit and time fixed effects, as in the canonical DiD. This allows for persistent level differences between treated and control units while still matching on pre-treatment trends. By including these fixed effects, SDiD absorbs unobserved heterogeneity across households and over time, an important feature in this context, as households differ systematically in baseline alcohol consumption, purchasing frequency, and responsiveness to price changes.

Consistent with this intuition, I estimate an SDiD model that minimizes the following error (using terminology from Arkhangelsky et al. (2021)):

$$\arg\min_{\lambda,\gamma,\mu,\beta}\left\{\sum_{h=1}^{N}\sum_{t=1}^{T}\left(Y_{ht}-\mu-\gamma_h-\lambda_t-W_{ht}\tau\right)^2\hat{\omega}_h^{sdid}\hat{\lambda}_t^{sdid}\right\} \tag{1}$$

where Y_{ht} is the inverse hyperbolic sine (IHS) of monthly ethanol purchases per adult (in ounces) of household h at month-year t.[16] I estimate analogous specifications for alternative outcomes, including measures of ethanol prices. Because ounces of ethanol purchased per

[16] Appendix Table A3 presents results estimated using $\log(Y_{ht}+1)$ as an alternative transformation of the outcome variable. In addition, Appendix Table A4 presents results where the outcome is total volume (in ounces) rather than ethanol-equivalent ounces.

adult is non-negative, highly right-skewed, and contains many zeros, I use the IHS transformation, which approximates the natural log for large values but remains defined at zero.[17] Although nonlinear models such as Poisson regressions are often appropriate for skewed non-negative outcomes (Chen and Roth, 2024), the SDiD estimator is defined for linear models and does not extend naturally to nonlinear specifications. Following best practice in the SDiD literature, I therefore use a linear model with an IHS-transformed dependent variable to approximate percentage changes while preserving the identification properties of the estimator. Appendix Table A5 also presents the results estimated using a Poisson model.

W_{ht} is an indicator equal to one for households in Maryland during the post-tax period (July 2011 onward) and zero otherwise. γ_h and λ_t denote household and month-year fixed effects. τ is the average treatment effect on the treated and measures the change in household ethanol purchases after the tax relative to before, comparing Maryland to the synthetic control group. Household weights $\hat{\omega}_h^{SDiD}$ and time weights $\hat{\lambda}_t^{SDiD}$ are optimally chosen to minimize Maryland's pre-treatment prediction error.[18] Standard errors are clustered at the state level. Specifications also include household demographic covariates such as household size, income, race, age of the household head, and education of the household head. [19]

The SDiD estimator is doubly robust, meaning it remains consistent if either the parallel-trends assumption holds or if the synthetic control weights correctly approximate the treated unit's counterfactual path. This makes it particularly appealing in policy settings with a single treated state or a limited number of treated units. Simulation and empirical evidence show that SDiD generally outperforms both DiD and SCM (Arkhangelsky et al., 2021). For completeness, I also report DiD and SCM estimates for comparison in Section 5.

To assess the dynamics of the policy effect, I complement the main SDiD estimates with an event-study–style analysis based on the SDiD framework, following the implementation in Clarke et al. (2024). The model is estimated at the monthly level, consistent with the main specification, to fully capture the timing of the 2011 tax and short-run fluctuations in alcohol purchases. For presentation, however, I aggregate the resulting monthly coefficients to the quarterly level to facilitate interpretation and improve visual clarity, as quarterly effects provide a smoother depiction of post-treatment dynamics. This event-study analysis serves two purposes: first, to confirm that Maryland's pre-tax purchasing trends were closely aligned with its synthetic control; and second, to illustrate the temporal persistence of the

[17] Inverse hyperbolic sine transformation: $\text{asinh}(Y_{ht}) = \ln(Y_{ht} + \sqrt{Y_{ht}^2 + 1})$.

[18] The derivation of unit ($\hat{\omega}_h^{SDiD}$) and time ($\hat{\lambda}_t^{SDiD}$) weights is presented in Appendix B.1.

[19] To incorporate covariates, Arkhangelsky et al. (2021) recommend a residualization approach: first regress Y_{ht} on X_{ht} to obtain $\hat{\beta}$, then apply the SDiD algorithm to the residualized outcome $Y_{ht}^{res} = Y_{ht} - X_{ht}\hat{\beta}$.

treatment effect following the 2011 tax. Having outlined the SDiD estimator and its dynamic extension, I next turn to the inference procedures used to obtain valid standard errors for these estimates.

4.2 Inference

Following Arkhangelsky et al. (2021), I compute standard errors using a clustered bootstrap procedure, which provides valid inference when there are many observed units but a single treated cluster. The bootstrap proceeds as follows:

1. Draw B bootstrap resamples of households (resampled with replacement). Each resample must include both treated and control units; otherwise, it is discarded and redrawn.

2. For each resample b, re-estimate the SDiD model and obtain $\hat{\tau}_{sdid}^{(b)}$.

3. Define the bootstrap variance as:

$$\hat{V}_{\tau}^{(boot)} = \frac{1}{B}\sum_{b=1}^{B}\left(\hat{\tau}_{sdid}^{(b)} - \bar{\tau}^{boot}\right)^2 \tag{2}$$

   where $\bar{\tau}^{boot} = \frac{1}{B}\sum_{b=1}^{B}\hat{\tau}_{(b)}^{sdid}$ is the mean of the bootstrap estimates.

4. The standard error is then estimated as the square root of the variance:

$$\hat{SE}_{\tau}^{boot} = \sqrt{\hat{V}_{\tau}^{(boot)}}. \tag{3}$$

I perform $B = 1{,}000$ bootstrap replications for each model. Conceptually, this procedure repeatedly re-estimates the treatment effect on resampled data to approximate the empirical sampling distribution of $\hat{\tau}_{SDiD}^{(b)}$, thus producing consistent standard errors and confidence intervals under weak distributional assumptions. The same bootstrap procedure is used for the event-study analysis. While computationally intensive, this approach is well-suited to this setting because the data include thousands of households, providing substantial resampling variation despite having a single treated state.

4.3 Alternative model specifications

I test the sensitivity of the main estimates to several alternative specifications and diagnostic checks. These analyses evaluate whether the results are robust to alternative inference procedures, comparison-group construction, sample restrictions, possible spillovers, treatment timing, and macroeconomic conditions.

I begin by reporting estimates using alternative inference procedures to account for the single treated unit. In the main specification, I use clustered bootstrap standard errors, which are well-suited to this setting because the data include many households per state, providing substantial resampling variation. Although computationally intensive, this approach provides a flexible approximation to the sampling distribution in settings with a single treated cluster. As complementary checks, I also estimate standard errors using the jackknife and placebo inference procedures proposed by Arkhangelsky et al. (2021).[20]

I then report DiD and SCM estimates for comparison, although SDiD has been shown to outperform these approaches. In the DiD specification, Maryland households are matched to control households using propensity score matching based on pre-treatment alcohol purchases and household demographic characteristics. The treatment effect is then estimated using a two-way fixed-effects DiD model on the matched sample. This provides a complementary estimate based on observed covariate balance rather than SDiD unit and time weighting.[21]

Next, I examine the sensitivity of the estimates to the composition of the analytic sample. The main sample is restricted to households observed in a balanced sample over the two years before and two years after the tax change. As a robustness check, I re-estimate the model including non-drinking households to assess whether the results are driven by conditioning on households with positive alcohol purchasing behavior.

I also assess possible violations of the stable unit treatment value assumption (SUTVA). Because Maryland borders Delaware, Pennsylvania, Virginia, and West Virginia, households may have responded to the tax by shifting purchases across state lines, and households in neighboring states may also have been indirectly affected by cross-border shopping. To evaluate this possibility, I estimate one specification that excludes neighboring states from the donor pool and another that restricts the comparison group to neighboring states. These checks assess whether the estimated effect depends on potentially contaminated nearby con-

[20]Appendix B.2 describes the jackknife procedure, and Appendix B.3 describes the placebo inference procedure.

[21]Appendix C describes the SCM and propensity-score-matched DiD procedures.

trols or, conversely, on comparisons with more geographically proximate states.

I then examine the sensitivity of the estimates to the timing of treatment. The main specification defines the post-treatment period as beginning in July 2011, when the Maryland alcohol sales tax increase took effect. I also estimate a specification that treats June 2011 as the beginning of the post period. This anticipation check evaluates whether households changed purchasing behavior immediately before the tax took effect, for example through stockpiling or other anticipatory responses.

I also assess whether the estimates are influenced by macroeconomic conditions during the Great Recession recovery period. Although the main panel begins after the official end of the recession, the early pre-treatment months may still reflect differential recovery dynamics across states. I therefore estimate specifications that exclude July to December 2009, such that the pre-treatment period begins in January 2010.

Finally, I examine the sensitivity of the results to alternative sample windows. In addition to the main balanced panel, I consider longer unbalanced panels permitted by the data to assess whether the estimated effects depend on the time horizon. Because these specifications involve unbalanced panels, the SDiD framework is not applicable; I therefore estimate two-way fixed effects models, which should be interpreted as complementary, directional checks rather than causal estimates.

5 Results

Table 1 reports weighted pre-tax means for outcomes and household characteristics for Maryland and the comparison states, using the NielsenIQ Consumer Panel projection factor as a sample weight to improve national representativeness.[22] The weighted means indicate that Maryland and the comparison states are broadly similar across most pre-treatment characteristics. Maryland households have somewhat higher income and education levels, a larger share of Black households, and a smaller share of Hispanic and White households relative to the donor pool. These differences are not consequential for the analysis, as the SDiD procedure subsequently constructs a weighted synthetic comparison that matches Maryland households' pre-treatment alcohol purchasing trends.

[22] The NielsenIQ dataset provides sample weights designed to approximate the demographic composition of the U.S. Census. The weights are used only to present the summary statistics. The models do not use these weights, as they already incorporate weights from the SDiD procedure.

[Table 1 here]

Figure 1 shows the raw ethanol purchasing trends for Maryland (blue circles) and the control states (black squares).[23] The figure shows similar pre-tax trajectories in ethanol ounces purchased per adult across both groups, followed by a decline in Maryland relative to the comparison states after the tax increase. These raw trends provide descriptive evidence of comparable pre-treatment patterns and a decrease in ethanol purchases in Maryland relative to the control group.

[Figure 1 here]

To formally estimate the effect of the tax suggested by the trends in Figure 1, I estimate the SDiD model in 1 to construct a synthetic Maryland by assigning weights to households in the donor pool, $\hat{\omega}_h^{SDiD}$, and time weights, $\hat{\lambda}_t^{SDiD}$, over the pre-treatment period.[24] Figure 2 plots observed Maryland outcomes and the synthetic control, together with the estimated unit and time weights. The green bars represent the weights assigned to pre-treatment months. The dashed line (synthetic Maryland) closely tracks ethanol ounces purchased per adult for Maryland households during the pre-treatment period and diverges after the tax, suggesting a reduction in ethanol purchases relative to the synthetic control.

[Figure 2 here]

Table 2 reports the overall treatment effect estimates of the tax increase on the inverse hyperbolic sine of ethanol ounces purchased per adult per household. Column (1) presents the estimates without covariates, while column (2) presents the estimates with covariates. Across both specifications, the results remain consistent and precisely estimated. For ease of interpretation, I convert the IHS coefficient into an approximate percent effect as suggested by Bellemare and Wichman (2020).[25] Column (1) reports a coefficient of -0.098, corresponding to an approximate 9.4% reduction in monthly ethanol ounces purchased per adult in Maryland households relative to the pre-tax period. This reduction translates into 1.06 fewer ethanol ounces purchased per adult per month.

[23] Appendix Figure A1 presents analogous raw trends for alcohol expenditures as a share of total grocery spending. The figure shows a similar post-tax decline in Maryland relative to the comparison states.

[24] Appendix Figure A2 shows the state-level donor weights used to construct synthetic Maryland. Because the SDiD model is estimated at the household level, the underlying unit weights are household-level weights. I aggregate these to the state level by summing the positive household weights within each state.

[25] I calculated the percent change as follows: $100 \times \left[\exp\left(\hat{\tau} - \frac{1}{2}\widehat{SE}(\hat{\tau})^2\right) - 1\right]$.

[Table 2 here]

Similar to column (1), column (2) presents an estimate of -0.096, corresponding to an approximate 9.2% reduction in monthly ethanol ounces purchased per adult in Maryland households relative to the pre-tax period. Column (2) is my preferred specification, as the time-invariant covariates include important demographic characteristics associated with alcohol purchases, including household income, education, age, and household size. Including these covariates helps account for observable differences between Maryland and the donor pool while leaving the main results substantively unchanged. This reduction translates into 1.04 fewer ethanol ounces purchased per adult per month. All following models are estimated with covariates unless stated otherwise.

For this reduction to be interpreted causally, the estimate relies on the synthetic control providing a credible counterfactual path for Maryland in the absence of the tax increase. Figure 2 shows that Maryland and its synthetic control exhibit similar pre-treatment trends. To further assess this identifying assumption, I estimate the event-study version of equation 1. Figure 3 presents the dynamic event-study coefficients from the SDiD specification. Blue dots denote point estimates, and gray capped lines denote 95% confidence intervals based on standard errors from equation 3. Although the SDiD model is estimated at the monthly level, the event-study coefficients are aggregated to the quarterly level for ease of display. Unlike a conventional two-way fixed effects event study, the SDiD estimates are not identified relative to a single omitted reference period (such as event time -1), but relative to the weighted synthetic control path constructed from the full pre-treatment period. Accordingly, all pre-treatment coefficients are shown. Pre-treatment coefficients remain close to zero, and the post-treatment coefficients show a sustained decline in ethanol ounces purchased per adult, consistent with the raw trends. There is a decline in the event-time coefficient at month 0, which reflects the construction of the event-study coefficients: effects are aggregated to the quarterly level (excluding July), so the plotted quarterly estimates capture post-treatment adjustments, while month 0 is shown at the monthly level.

[Figure 3 here]

So far, the Maryland alcohol tax increase appears to have reduced alcohol purchases among consumers. I next examine whether this reduction was concentrated in particular beverage categories, whether responses differed across consumer groups, whether the effects varied for households living near state borders, and how the effective prices paid changed.

5.1 Beverage Types

Table 3 presents the estimated effect of the tax on ethanol ounces purchased by beverage type: beer, wine, and spirits per adult in a household. Thus, the outcomes are standardized in terms of pure ethanol content within each beverage category rather than total beverage volume purchased, to remain consistent with the main estimate. Column (1) reports the estimate for beer, column (2) for wine, and column (3) for spirits.

[Table 3 here]

Across all categories, the estimates indicate meaningful declines in beverage purchases. The estimated reductions are approximately 3.4% for beer, 5.3% for wine, and 4.5% for spirits, translating to monthly declines of roughly 0.13 ethanol ounces attributable to beer purchases, 0.20 ethanol ounces attributable to wine purchases, and 0.17 ethanol ounces attributable to spirits purchases per adult per household. Given that mean pre-tax ethanol ounces purchased by beverage type are fairly similar, the largest percent effect for wine also corresponds to the largest absolute reduction, followed by spirits and then beer.[26] The results do not provide evidence of substitution toward any single beverage type. Appendix Figure A3 presents the SDiD pre-treatment fit by alcohol type, and Appendix Figure A4 presents the event-study estimates.

5.2 Heterogeneous behavioral responses

So far, the results indicate an overall decrease in ethanol purchases, with no evidence of substitution across beverage types. However, these aggregate effects may mask different response patterns across households. Households may differ in their responsiveness to the tax depending on baseline alcohol purchasing behavior, income, household composition, age, education, and race or ethnicity. I therefore examine heterogeneity along several dimensions.

5.2.1 Purchasing levels

I first evaluate how monthly ethanol ounces purchased per adult change by household purchasing intensity. Households are classified into low-, moderate-, and high-intensity terciles

[26] Because each beverage category is estimated separately using the IHS transformation and outcome-specific SDiD weights, the category-specific reductions are not expected to sum exactly to the aggregate estimate.

based on their average pre-treatment ethanol purchases per adult, excluding anomalous purchasers/stockpilers so that the categories reflect persistent purchasing patterns rather than unusually large one-time purchases.[27] This analysis examines whether the overall reduction is concentrated among households with higher baseline alcohol purchasing or whether it is distributed more broadly throughout the purchasing distribution. Table 4 reports the estimates for (1) low-, (2) moderate-, and (3) high-intensity purchasing households.[28]

[Table 4 here]

The results are consistent across purchasing-intensity groups. Following the tax increase, Maryland households significantly reduced monthly ethanol ounces purchased per adult in each category. Low-purchasing households reduced purchases by 6.0%, moderate-purchasing households by 8.0%, and high-purchasing households by 6.8%, relative to the pre-tax period. These estimates correspond to approximately 0.04 fewer ounces of ethanol per adult per month for low-purchasing households, 0.23 fewer ounces for moderate-purchasing households, and 1.49 fewer ounces for high-purchasing households. Thus, the largest absolute reduction occurred among high-purchasing households. Appendix Figure A6 and Appendix Figure A7 present the SDiD fit and event-study estimates by purchasing intensity.

To assess whether higher-purchasing households substitute toward cheaper sources of ethanol, I further estimate the tax effect by beverage type and purchasing intensity. Appendix Table A7 reports these results. The estimated reductions remain broadly consistent across purchasing groups and beverage categories, providing no evidence that the main effects are driven by substitution toward cheaper alcohol types.

5.2.2 Demographic heterogeneity

I next examine heterogeneity across demographic and household characteristics. Specifically, I estimate separate effects by household income, household head education, household composition, and race or ethnicity.

[27] Appendix Figure A5 shows the distribution of alcohol purchasers and indicates that anomalous purchasers are located in the far right tail and represent a small share of the sample.

[28] Appendix Table A6 reports robustness estimates using alternative purchaser-intensity cutoffs. Rather than using the main classification, the table groups non-stockpiler households into the bottom quartile, middle 50 percent, and top quartile of the pre-treatment distribution of ethanol purchases per adult. The estimated effects remain consistent.

Figure 4 summarizes the demographic heterogeneity results. The figure plots the estimated SDiD treatment effect for each subgroup, with capped gray horizontal lines denoting 95% confidence intervals. Subgroups are defined using pre-treatment household characteristics: income, household head education, household composition, and race/ethnicity. For households with two heads, education is assigned using the higher reported category across heads.[29]

[Figure 4 here]

For most categories, the point estimates indicate declines in ethanol ounces purchased per adult following the tax increase. In contrast, households with income below $35,000 and those with less than a high school education exhibit positive point estimates. These estimates should be interpreted cautiously. When the sample is divided into narrower subgroups, the number of treated Maryland households becomes substantially smaller, limiting statistical precision. Accordingly, the figure is best viewed as descriptive evidence on the direction and relative magnitude of subgroup responses rather than definitive causal estimates. Average treatment effects on the treated by demographic category are reported in Appendix Tables A8 to A11.

5.2.3 Border counties

Another potential response to alcohol taxes is cross-border shopping, particularly for households located near lower-tax jurisdictions. In this setting, however, one objective of the policy was to align Maryland's alcohol tax rates more closely with those in neighboring states, limiting the scope for cross-border price differentials and making this margin of adjustment less likely. Nonetheless, Maryland borders several states, creating opportunities for consumers to avoid the tax by purchasing alcohol elsewhere. To assess whether proximity to other states affects purchasing patterns, I compare ethanol ounces purchased by households residing in Maryland border counties with those in non-border counties.

Because this analysis focuses exclusively on Maryland households, the SDiD framework is not applicable. Instead, I estimate a traditional TWFE model with household and time

[29] Household-composition heterogeneity is defined using mutually exclusive bins based on household size and the presence of children. Because two-person households with children constitute a very small share of the treated sample (0.6%), estimating a separate effect for this subgroup would not yield valid estimates for inference. I therefore group all two-person households together, while continuing to distinguish between three-or-more-person households with and without children.

fixed effects, clustering standard errors at the county level. By interacting the event-time indicators with an indicator for border-county residence, this specification identifies whether the post-tax changes in ethanol purchases differ for border households relative to non-border households. Figure 5 presents the event study estimates from this model.

[Figure 5 here]

Although this approach does not directly observe tax avoidance behavior, it assesses whether households in border counties changed their overall ethanol purchases differently from those in non-border counties after the tax increase. Because NielsenIQ records purchases scanned by households, including out-of-state purchases when reported, substantial cross-border shopping would attenuate the decline in measured purchases for border households relative to non-border households. If so, border households would be expected to exhibit smaller post-tax reductions in ethanol purchases. However, the estimates show no consistent pattern. For the most part, border and non-border households do not differ significantly in either the pre-treatment or post-treatment periods, with only a small number of quarters showing statistically significant differences. Overall, the results suggest that border proximity did not meaningfully offset the tax response.[30]

5.3 Prices paid for ethanol

I next examine prices paid for ethanol. This analysis is more challenging because the NielsenIQ Consumer Panel does not report posted shelf prices for all household-month observations. Instead, I observe unit values implied by household purchases: real alcohol spending divided by ounces of ethanol purchased. These unit values reflect both the prices consumers face and endogenous purchasing choices, including the products, package sizes, and quantities selected. I therefore report three complementary measures to distinguish changes in effective prices from changes in purchase incidence.

[Table 5 here]

Panel A of Table 5 estimates effects on the inverse hyperbolic sine of the tax-inclusive price per ounce of ethanol.[31] For the SDID specification, non-purchase months are coded as zero so

[30] Appendix Table A13 reports the estimated treatment effects.

[31] The tax-inclusive measure is constructed by multiplying the observed unit value by 1.03 for Maryland household-month observations after the tax increase.

that the outcome is observed for the balanced household panel. The estimates are negative overall and for beer, wine, and spirits. The overall estimate implies an approximate 2.9% decline in dollars spent per ounce of ethanol, or about $0.02 relative to the Maryland pre-treatment mean of $0.69 per ounce. This corresponds to roughly one cent per standard drink, since a standard drink contains 0.6 ounces of ethanol. Although statistically significant, the implied economic magnitude is small. Moreover, this measure should be interpreted cautiously because it combines changes in transaction prices with changes in whether and what households purchase.

Panel B examines purchase incidence, defined as whether the household made any alcohol purchase in a given month. The estimates indicate that the tax reduced monthly purchase incidence overall and across beverage categories. This matters for interpreting Panel A because the SDiD price measure includes zero-valued non-purchase months; reductions in purchase frequency can therefore mechanically lower the unconditional effective price measure even if prices conditional on purchase do not fall.

Panel C estimates two-way fixed effects models for the tax-inclusive price per ounce of ethanol conditional on a positive purchase. Because this outcome is observed only in purchase months, the panel is no longer balanced and is not well suited to the SDID framework. These estimates should also be interpreted cautiously because identification comes from one treated state and conventional clustered standard errors may understate standard errors. Conditional on purchase, the estimated effects are positive overall and by beverage type. This pattern suggests that the negative estimates in Panel A are driven primarily by changes in purchase incidence and purchasing behavior, rather than by meaningful declines in the underlying price paid per ounce of ethanol among households.[32]

5.4 Robustness checks

I look at the robustness of the main SDiD estimate using alternative estimators, inference procedures, sample definitions, donor-pool restrictions, and treatment timing assumptions. These checks evaluate whether the estimated effect of the Maryland sales tax is sensitive to how the comparison group is constructed, the method used for inference, or the composition

[32]Shares are not reported for Panels B and C because these panels are intended to decompose the unconditional price measure rather than estimate compositional outcomes. Panel B is already an extensive-margin measure, indicating whether any purchase occurred in the household-month. Panel C restricts the sample to positive-purchase months and estimates the price paid per ounce of ethanol conditional on purchase. In that conditional sample, non-purchase months are excluded, so purchase shares over the full household-month panel are not defined in the same way and would answer a different question.

of the household sample.

Figure 6 summarizes these robustness checks using a coefficient plot. Each point represents an estimate of the average treatment effect of the Maryland sales tax on the inverse hyperbolic sine of monthly ethanol ounces purchased per adult under a different specification. Horizontal capped lines denote 95 percent confidence intervals. The baseline estimate, shown with blue circles, reproduces the preferred SDiD specification for reference.

[Figure 6 here]

First, I examine whether inference is sensitive to alternative standard error procedures. The baseline specification uses bootstrap standard errors clustered at the state level. I compare this to jackknife standard errors, which recompute uncertainty by sequentially omitting observations, and to placebo-based inference, which uses the distribution of placebo treatment assignments to assess uncertainty. Confidence intervals are somewhat wider under these alternative procedures, but the estimates remain negative and statistically significant at the 5% level.

I next assess robustness to alternative empirical strategies commonly used in settings with a single treated unit: the synthetic control method and difference-in-differences using a propensity-score matched control group. The propensity-score matched difference-in-differences specification first matches Maryland households to comparison households using pre-treatment alcohol purchases and household covariates, and then estimates a two-way fixed effects difference-in-differences model on the matched sample. This provides a check based on observable pre-treatment similarity rather than SDiD weighting. The results remain consistent across these alternative approaches, with the propensity score matching giving an even larger estimate.

I then consider whether the results depend on the household sample or donor pool. The main analysis restricts the sample to households that purchased alcohol at least once in the pre-treatment period. I therefore re-estimate the model using a broader sample that includes non-purchasing households. This assesses whether the estimated decline is driven by the restriction to pre-treatment alcohol purchasers. The effect is attenuated but remains negative and statistically significant.

I also report two donor-pool restrictions. One excludes Maryland's neighboring states from the donor pool, addressing the possibility that cross-border shopping or other spillovers may affect nearby control states which would be a violation of SUTVA. The other uses only neighboring states—Delaware, Pennsylvania, Virginia, and West Virginia—as the donor pool, providing a geographically closer comparison group. The estimates remain negative under these alternative donor-pool definitions, although the magnitudes vary.

I next examine sensitivity to treatment timing and macroeconomic conditions. I redefine the treatment date as June 2011, one month before the tax took effect, to assess whether the estimates are sensitive to anticipatory behavior. Then, I also drop July through December 2009, the first six months of the sample, which coincide with the immediate post-recession recovery period. This checks whether the main estimate is influenced by differential consumption dynamics during the early recovery from the Great Recession. The estimated effects remain negative and statistically significant in both cases.

Finally, I examine alternative sample windows. In the main specification, I construct a balanced household panel to exploit within-household purchasing behavior over time, which requires dropping households not continuously observed throughout the sample period. To assess sensitivity to this restriction, I consider additional sample windows permitted by the data. First, I use the full NielsenIQ sample available prior to COVID-19, covering 2004 to 2019. Second, I construct a 10-year panel from 2006 to 2016. Third, using the same four-year window as the main specification (July 2009 to June 2013), I retain all available households rather than restricting the sample to households observed for the full period.

Because these specifications contain unbalanced household panels, the SDiD framework is not applicable, as it requires balanced panels. I therefore estimate two-way fixed effects models, alternating between state fixed effects and household fixed effects, reported in Table 6. Columns (1) and (2) use the full pre-COVID sample (2004–2019), columns (3) and (4) use the 10-year panel (2006–2016), and columns (5) and (6) use the main four-year window (July 2009–June 2013) while retaining all available households. Within each pair, the first column includes state fixed effects and the second includes household fixed effects. These estimates should be interpreted cautiously. Although they exploit variation across many treated households, they do not address the single treated-state setting in the same way as the main SDiD design and are therefore best viewed as directional robustness checks rather than causal estimates.

[Table 6 here]

Across the different sample windows, treatment effects range from declines of 4.8% to

6.8%. Across all specifications, the estimated effects remain negative, suggesting a reduction in monthly ethanol ounces purchased per adult, although magnitudes vary. Together, these results suggest that the estimated effect of the tax is stable across inference procedures, empirical strategies, sample definitions, donor-pool restrictions, treatment timing assumptions, macroeconomic conditions, and alternative sample windows.

6 Discussion

A large literature documents that increases in the price of alcohol lead to reductions in alcohol consumption (Wagenaar et al., 2009; Maldonado-Molina and Wagenaar, 2010; Saffer et al., 2024; Xuan et al., 2015). However, the magnitude and consistency of these responses depend on the relatively inelastic nature of alcohol demand and on heterogeneity across consumers (Ayyagari et al., 2013; Gehrsitz et al., 2021; Pryce et al., 2019). While most empirical work has focused on excise taxes, comparatively little is known about the behavioral effects of alcohol sales taxes, which differ in both structure and salience. This paper contributes to that literature by examining the impact of Maryland's 2011 alcohol sales tax on household alcohol purchases and related behavioral responses.

The primary takeaway is that the tax led to a meaningful reduction in alcohol purchases. Using a synthetic difference-in-differences framework and data from the NielsenIQ Consumer Panel, I estimate a 9.2% decline in monthly ethanol ounces purchased per adult among alcohol-purchasing households. This corresponds to approximately 1.04 fewer ounces of ethanol per adult per month, or roughly 1.7 fewer standard drinks. This magnitude is economically meaningful and consistent with the broader literature showing that higher alcohol prices reduce consumption. Importantly, the results suggest that a sales tax, despite concerns about lower salience relative to excise taxes (Chetty et al., 2009), can generate behavioral responses comparable to those typically found in excise tax studies.

While these estimates are not directly comparable, I relate them to the excise tax literature using results from Saffer et al. (2024). They study the 2009 Illinois alcohol excise tax increase, which raised spirits taxes from \$4.50 to \$8.55 per gallon, wine taxes from \$0.73 to \$1.39 per gallon, and beer taxes from \$0.19 to \$0.23 per gallon. They find an approximately 6.4% reduction in ethanol purchases, corresponding to roughly 0.7 fewer ounces of ethanol based on their pre-treatment mean. In comparison, I estimate a reduction of 1.04 ounces of ethanol. Despite differences in tax structure, baseline prices, and empirical setting, the estimates are similar in magnitude, with the Maryland estimate somewhat larger.

One possible explanation for these findings is that, although sales taxes are often considered less salient in a single purchasing decision, they are applied to the final retail price and therefore scale with consumer choice. Because the tax is proportional to price, it raises the overall cost of alcohol consumption while leaving relative prices across products largely unchanged.[33] Consistent with this, I find no evidence of substitution across beverage types: purchases of beer, wine, and spirits all decline following the tax increase. By contrast, prior work on excise taxes has documented substitution toward lower-priced products (Gehrsitz et al., 2021). While I do not directly examine substitution within beverage categories, the absence of substitution across beverage types suggests that consumers respond to the overall price they face rather than reallocating purchases across alcohol categories.

Importantly, this analysis uses panel data that follow households over time, allowing consumers to repeatedly face and adjust to the higher post-tax prices. In this setting, the salience of the tax may be less relevant than the persistent increase in the final price, as households learn and update their purchasing behavior across shopping trips. This dynamic adjustment helps explain why the estimated effects are comparable to those found in the excise tax literature, despite differences in how the taxes are presented at the point of purchase.

Consistent with this interpretation, the reductions in alcohol purchases are broadly distributed across households rather than concentrated among specific groups. Households across the purchasing distribution reduce their alcohol purchases following the tax increase, indicating that the response is not limited to a particular group of consumers, although the largest absolute reductions occur among higher-purchasing households. Across demographic groups, the results show broadly similar patterns, with most groups exhibiting declines in alcohol purchases. Although some subgroup estimates differ in direction, these estimates are less precise due to smaller sample sizes and should be interpreted cautiously. In particular, households with lower income and lower levels of education show point estimates suggesting increases in ethanol purchases. These patterns do not necessarily imply higher alcohol spending, but may instead reflect differences in purchasing behavior or product choice. Overall, the evidence does not point to large or systematic differences in responsiveness across income, education, household composition, or race and ethnicity, suggesting that the tax operates broadly across the population rather than targeting a specific demographic group.

[33] By relative prices, I refer to the ratio of prices across products (e.g., beer relative to wine or spirits). Under sales taxes, prices increase proportionally, so these ratios remain largely unchanged.

As an additional behavioral margin, I examine households living in Maryland border counties, where cross-border shopping is more accessible. I find no consistent evidence of differential purchasing patterns in border counties relative to non-border counties. This suggests that, despite having the option to shop in nearby states, border households reduced their ethanol purchases similarly to households in non-border areas.

Finally, to provide additional context for these results, I also examine prices paid per ounce of ethanol. This analysis highlights the importance of distinguishing between observed unit values and underlying transaction prices. The price measure used in the data is constructed from household purchases and therefore reflects both the prices consumers face and their purchasing decisions. When averaging across all household-months, including months in which no alcohol is purchased, the price per ounce of ethanol appears to decline after the tax increase. However, this pattern is driven in part by a reduction in purchase incidence, as more zero-purchase months enter the calculation. To isolate changes in prices when alcohol is actually purchased, I examine the price per ounce only in months with positive purchases. In these months, the price per ounce increases following the tax, consistent with the mechanical effect of the tax on retail prices. This comparison indicates that the apparent decline in average prices is not driven by lower transaction prices, but rather by changes in purchasing behavior.

Together, these findings indicate that alcohol sales taxes can meaningfully reduce alcohol purchases and may serve as a useful complement to traditional excise taxes. By increasing the final retail price across a wide range of products, sales taxes can generate broad-based reductions in demand without inducing substantial substitution across beverage types or geographic avoidance behavior. These results have important implications for alcohol tax policy, suggesting that policymakers may consider both excise and sales taxes as part of a comprehensive strategy to reduce excessive alcohol consumption and its associated social harms.

7 Conclusion

Overall, the 2011 alcohol sales tax increase in Maryland appears to have been an effective tool for reducing household alcohol purchases. Across specifications, the tax is associated with a decline in monthly ethanol ounces purchased per adult per household of approximately 9.2%, equivalent to roughly 1.7 fewer standard alcoholic drinks per month. Declines are observed across beer, wine, and spirits, with no evidence of substitution toward any single beverage category. Reductions are also present throughout the purchasing distribution. While households with low, moderate, and high baseline purchasing levels all reduce alcohol purchases following the tax increase, the largest absolute declines happen among high-purchasing households.

Subgroup estimates suggest broadly similar patterns across demographic groups defined by income, education, household composition, and race or ethnicity, with most estimates indicating declines in ethanol ounces purchased per adult per household, although these estimates are less precise. Finally, there appears to be little evidence that households in border counties offset the tax through cross-border shopping. Taken together, the findings indicate that a broad-based sales tax can meaningfully reduce alcohol purchases. Alcohol sales taxes may be a useful complement to traditional excise taxes in efforts to reduce excessive alcohol consumption and its associated social harms.

References

Abadie, Alberto, "Using synthetic controls: Feasibility, data requirements, and methodological aspects," *Journal of economic literature*, 2021, *59* (2), 391–425.

_ **, Alexis Diamond, and Jens Hainmueller**, "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American statistical Association*, 2010, *105* (490), 493–505.

_ **and Javier Gardeazabal**, "The economic costs of conflict: A case study of the Basque Country," *American economic review*, 2003, *93* (1), 113–132.

Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager, "Synthetic difference-in-differences," *American Economic Review*, 2021, *111* (12), 4088–4118.

Ayyagari, Padmaja, Partha Deb, Jason Fletcher, William Gallo, and Jody L Sindelar, "Understanding heterogeneity in price elasticities in the demand for alcohol for older individuals," *Health economics*, 2013, *22* (1), 89–105.

Bellemare, Marc F and Casey J Wichman, "Elasticities and the inverse hyperbolic sine transformation," *Oxford Bulletin of Economics and Statistics*, 2020, *82* (1), 50–61.

Chen, Jiafeng and Jonathan Roth, "Logs with zeros? Some problems and solutions," *The Quarterly Journal of Economics*, 2024, *139* (2), 891–936.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and taxation: Theory and evidence," *American economic review*, 2009, *99* (4), 1145–1177.

Clarke, Damian, Daniel Pailanir, Susan Athey, and Guido Imbens, "On synthetic difference-in-differences and related estimation methods in Stata," *The Stata Journal*, 2024, *24* (4), 557–598.

Clements, Kenneth W, Marc Jim M Mariano, George Verikios, and Berwyn Wong, "How elastic is alcohol consumption?," *Economic Analysis and Policy*, 2022, *76*, 568–581.

Donald, Stephen G and Kevin Lang, "Inference with difference-in-differences and other panel data," *The Review of Economics and Statistics*, 2007, *89* (2), 221–233.

Esser, Marissa B, Hugh Waters, Mieka Smart, and David H Jernigan, "Impact of Maryland's 2011 alcohol sales tax increase on alcoholic beverage sales," *The American Journal of Drug and Alcohol Abuse*, 2016, *42* (4), 404–411.

Gehrsitz, Markus, Henry Saffer, and Michael Grossman, "The effect of changes in alcohol tax differentials on alcohol consumption," *Journal of public economics*, 2021, *204*, 104520.

Hanson, Andrew and Ryan Sullivan, "Incidence and salience of alcohol taxes: Do consumers overreact?," *Public Finance Review*, 2016, *44* (3), 344–369.

Imbens, Guido W and Donald B Rubin, *Causal inference in statistics, social, and biomedical sciences*, Cambridge university press, 2015.

Kenkel, Donald S, "Drinking, driving, and deterrence: The effectiveness and social costs of alternative policies," *The Journal of Law and Economics*, 1993, *36* (2), 877–913.

_ , "Are alcohol tax hikes fully passed through to prices? Evidence from Alaska," *American Economic Review*, 2005, *95* (2), 273–277.

Kilian, Carolin, Julia M Lemp, Laura Llamosas-Falcon, Tessa Carr, Yu Ye, William C Kerr, Nina Mulia, Klajdi Puka, Aurelie M Lasserre, Sophie Bright et al., "Reducing alcohol use through alcohol control policies in the general population and population subgroups: a systematic review and meta-analysis," *EClinicalMedicine*, 2023, *59.*

Lavoie, Marie-Claude, Gordon S Smith, Patricia Langenberg, Andres Villaveces, Patricia Dischinger, Linda Simoni-Wastila, and Kathleen Hoke, "303 Does increasing alcohol taxes always reduce alcohol-related crashes: Maryland sales tax increase?," 2016.

Maldonado-Molina, Mildred M and Alexander C Wagenaar, "Effects of alcohol taxes on alcohol-related mortality in Florida: time-series analyses from 1969 to 2004," *Alcoholism: clinical and experimental research*, 2010, *34* (11), 1915–1921.

Meng, Yang, Alan Brennan, Robin Purshouse, Daniel Hill-McManus, Colin Angus, John Holmes, and Petra Sylvia Meier, "Estimation of own and cross price elasticities of alcohol demand in the UK—A pseudo-panel approach using the Living Costs and Food Survey 2001–2009," *Journal of health economics*, 2014, *34*, 96–103.

National Center for Drug Abuse Statistics, "Alcohol Abuse Statistics," https://drugabusestatistics.org/alcohol-abuse-statistics/ 2024. Accessed: 2025-03-04.

National Institute on Alcohol Abuse and Alcoholism, "Alcohol Beverages Taxes," Alcohol Policy Information System 2024. https://alcoholpolicy.niaaa.nih.gov/apis-policy-topics.

Nelson, Jon P and John R Moran, "Effects of alcohol taxation on prices: a systematic review and meta-analysis of pass-through rates," *The BE Journal of Economic Analysis & Policy*, 2019, *20* (1), 20190134.

Pryce, Robert, Bruce Hollingsworth, and Ian Walker, "Alcohol quantity and quality price elasticities: quantile regression estimates," *The European Journal of Health Economics*, 2019, *20*, 439–454.

Ramirez, Rebecca L and David H Jernigan, "A Case Study of the 2011 Maryland Campaign to Raise Alcohol Taxes," *Center for Public Health Advocacy*, n.d.

Rubin, Donald B, "Using propensity scores to help design observational studies: application to the tobacco litigation," *Health Services and Outcomes Research Methodology*, 2001, *2*, 169–188.

Sacks, Jeffrey J, Katherine R Gonzales, Ellen E Bouchery, Laura E Tomedi, and Robert D Brewer, "2010 national and state costs of excessive alcohol consumption," *American journal of preventive medicine*, 2015, *49* (5), e73–e79.

Saffer, Henry, Markus Gehrsitz, and Michael Grossman, "Do the effects of alcohol excise taxes differ by drinking level and by income level?," *National Tax Journal*, 2024, *77* (4).

Shrestha, Vinish and Sara Markowitz, "The pass-through of beer taxes to prices: evidence from state and federal tax changes," *Economic Inquiry*, 2016, *54* (4), 1946–1962.

Smart, Mieka J, Safiya S Yearwood, Seungyoung Hwang, Roland J Thorpe Jr, and C Debra Furr-Holden, "Impact of Alcohol Tax Increase on Maryland College Students' Alcohol-Related Outcomes," *Substance use & misuse*, 2018, *53* (6), 1015–1020.

Son, Chong Hwan and Kudret Topyan, "The effect of alcoholic beverage excise tax on alcohol-attributable injury mortalities," *The European Journal of Health Economics*, 2011, *12*, 103–113.

Staras, Stephanie AS, Melvin D Livingston, and Alexander C Wagenaar, "Maryland alcohol sales tax and sexually transmitted infections: A natural experiment," *American journal of Preventive Medicine*, 2016, *50* (3), e73–e80.

Wagenaar, Alexander C, Matthew J Salois, and Kelli A Komro, "Effects of beverage alcohol price and tax levels on drinking: a meta-analysis of 1003 estimates from 112 studies," *Addiction*, 2009, *104* (2), 179–190.

Xuan, Ziming, Frank J Chaloupka, Jason G Blanchette, Thien H Nguyen, Timothy C Heeren, Toben F Nelson, and Timothy S Naimi, "The relationship between alcohol taxes and binge drinking: evaluating new tax measures incorporating multiple tax and beverage types," *Addiction*, 2015, *110* (3), 441–450.

Young, Douglas J and Agnieszka Bielińska-Kwapisz, "Alcohol taxes and beverage prices," *National tax journal*, 2002, *55* (1), 57–73.

Figures

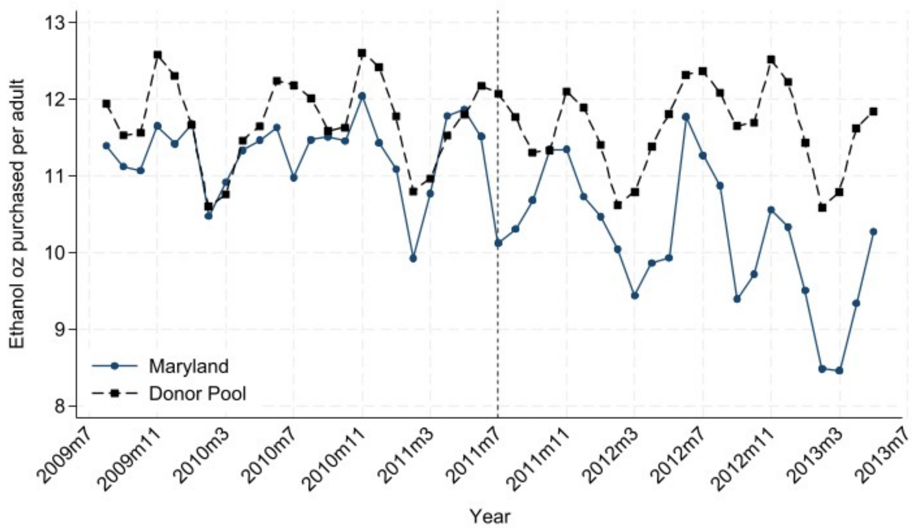

Figure 1. Trends in ethanol ounces purchased per adult in Maryland and donor pool

Note. The Y-axis shows the three-month moving average of ethanol ounces purchased per adult per household, and the X-axis represents the month–year. The figure presents two series: average alcohol consumption in Maryland (blue circles, solid line) and in the comparison states (black squares, dashed line). The vertical dashed line indicates the month of Maryland's sales tax increase (July 2011). Comparison states exclude households in Connecticut, Illinois, New Jersey, New York, North Carolina, Washington, the District of Columbia, and Kansas, as these states experienced alcohol tax changes during the sample period. Data come from the NielsenIQ Consumer Panel.

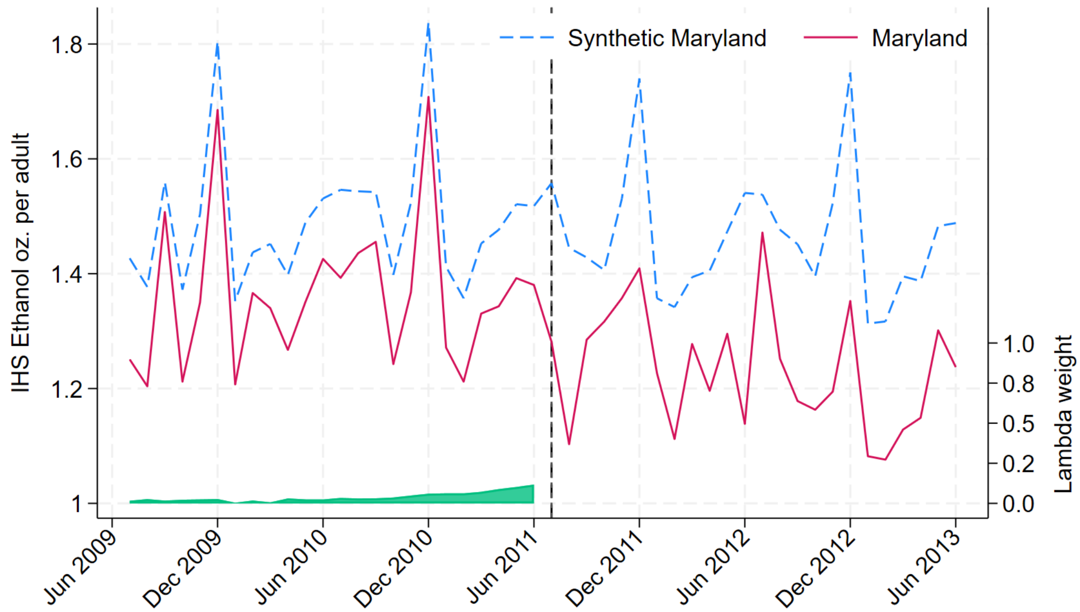

Figure 2. Synthetic Maryland fit

Note. Trends in ounces of ethanol purchased per adult for Maryland households using Synthetic Difference-in-Differences (SDiD) weights. The time weights (λ_t green bar) are estimated by minimizing the error term in Equation 1. The solid line shows the trend for Maryland, and the dashed line shows the corresponding trend for the synthetic control. The vertical dashed line denotes July 2011, when Maryland implemented the alcohol sales tax increase. Comparison states exclude Connecticut, Illinois, New Jersey, New York, North Carolina, Washington, the District of Columbia, and Kansas, which experienced alcohol tax changes during the sample period. Data are from the NielsenIQ Consumer Panel.

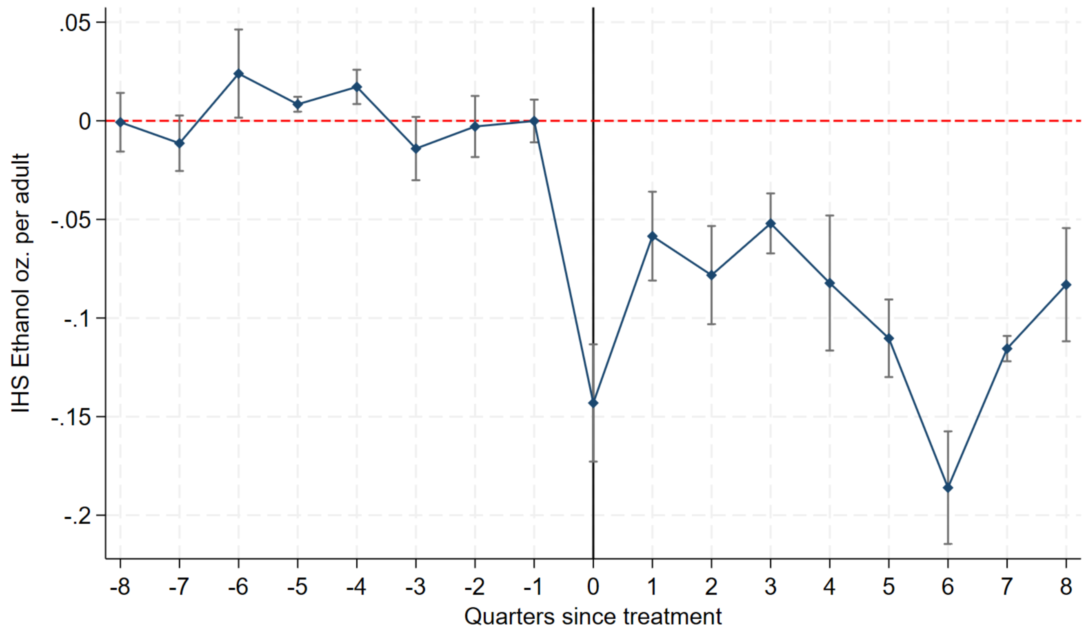

Figure 3. Dynamic Event study estimates for ethanol ounces purchased

Note. The figure shows dynamic treatment effects estimated using synthetic difference-in-differences (SDiD). The vertical line at event time zero marks the month of the tax change (July 2011). The dependent variable is the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month. Monthly event-time coefficients were averaged to the quarter level. Blue diamonds indicate point estimates, and gray capped lines represent 95% confidence intervals computed from 1,000 bootstrap replications. Standard errors were clustered at the state level. The model includes income, household size, race, marital status, presence of children, education, and age controls. Data on alcohol purchases come from the NielsenIQ Consumer panel for the period from June 2009 to July 2013.

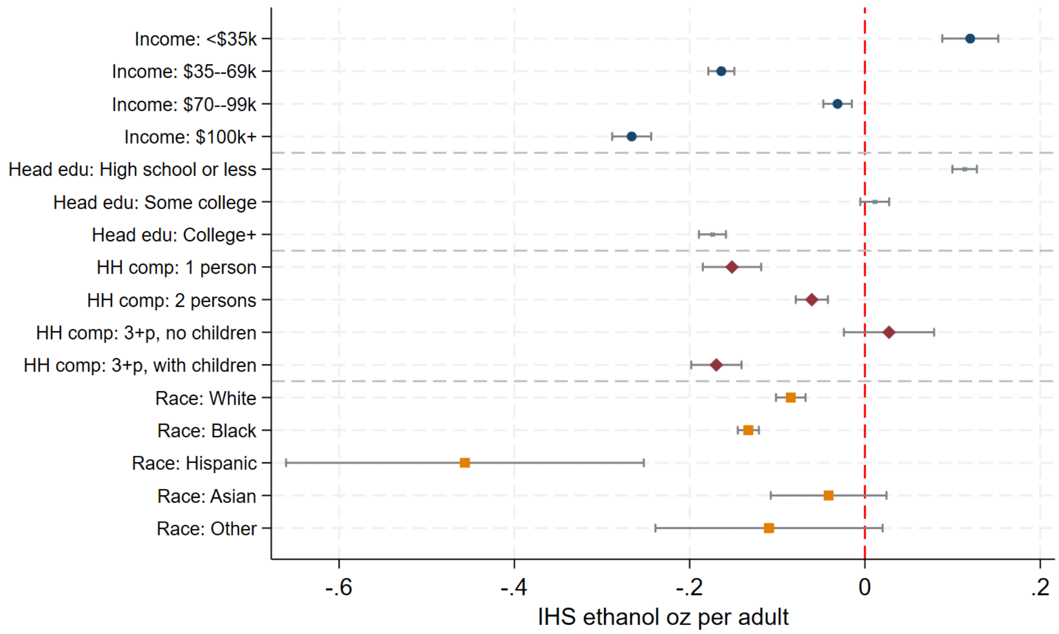

Figure 4. Effect of Maryland's alcohol sales tax on ethanol ounces purchased across demographic groups

Note. The figure plots subgroup-specific average treatment effect estimates across the following demographic categories: income (blue circles), head of household education (teal x), household composition (red diamonds), and race or ethnicity (yellow squares). The outcome variable is the inverse hyperbolic sine of monthly ethanol ounces purchased per adult per household. Average treatment effects on the treated are estimated using synthetic difference-in-differences (SDiD). Gray capped lines represent 95% confidence intervals computed from 1,000 bootstrap replications. Standard errors are clustered at the state level. All specifications include household income, household size, marital status, presence of children, race or ethnicity, education, and age controls, except that the corresponding category is omitted when estimating subgroup-specific effects. For example, income controls are omitted when estimating effects by income category. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period June 2009 to July 2013.

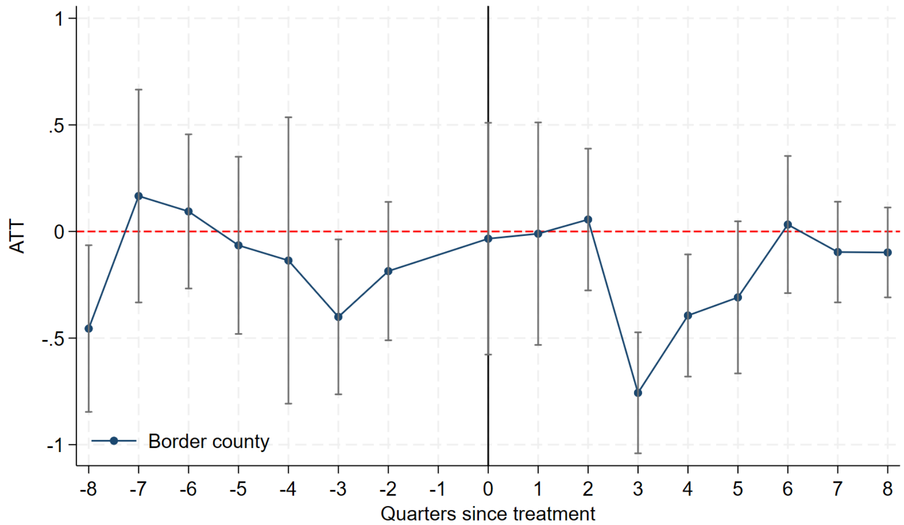

Figure 5. Event study estimates for ethanol ounces purchased for Maryland border counties

Note. The figure shows treatment effects estimated using a TWFE model restricted to Maryland households. The dependent variable is the inverse hyperbolic sine of monthly ethanol ounces purchased per adult per household. Monthly event-time coefficients are averaged to the quarter level for display. The vertical line at event time zero marks the month of the tax change. Blue circles indicate point estimates, and gray capped lines represent 95% confidence intervals. Standard errors are clustered at the county level. The model includes income, household size, marital status, presence of children, race, education, and age controls. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period June 2009 to July 2013.

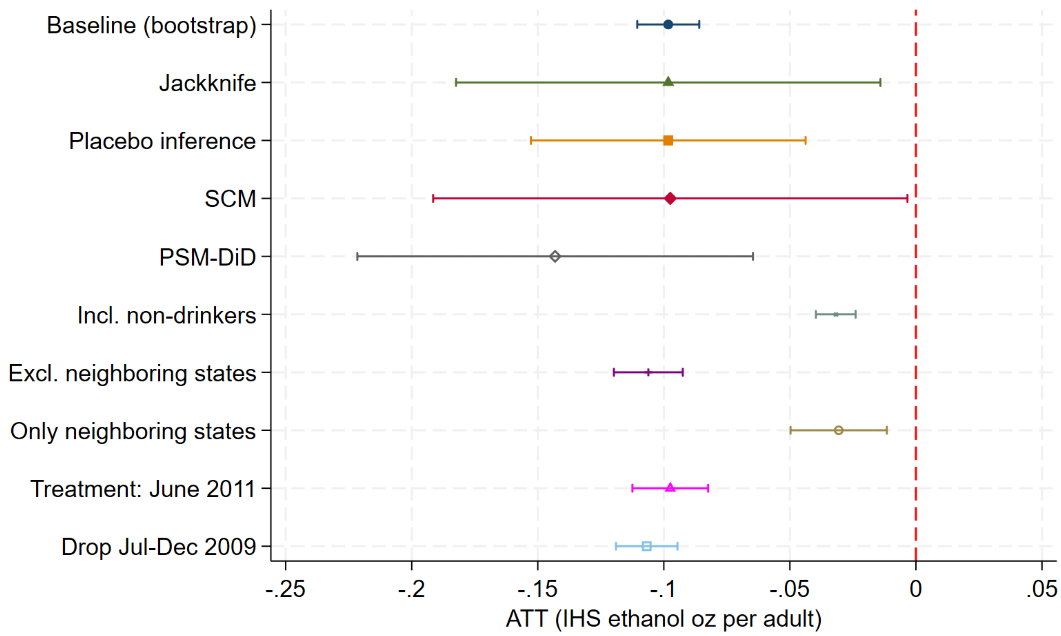

Figure 6. Robustness of estimates to alternative specifications

Note. The figure plots average treatment effect estimates across alternative inference procedures, sample restrictions, and model specifications. The outcome variable is the inverse hyperbolic sine of monthly ethanol ounces purchased per adult per household. The blue circle denotes the main synthetic difference-in-differences (SDiD) estimate using clustered bootstrap inference. The green triangle denotes the SDiD estimate using jackknife inference, and the yellow square denotes the SDiD estimate using placebo inference. The red diamond denotes the synthetic control method (SCM) estimate, and the hollow gray diamond denotes the estimate from an alternative propensity score matching difference-in-differences specification. The teal cross denotes the SDiD estimate from a sample that includes non-drinkers. The purple plus sign denotes the SDiD estimate excluding neighboring states, while the hollow green circle denotes the SDiD estimate using only nearby donor states (Alabama, Virginia, New Jersey, and Delaware). The hollow pink triangle denotes the SDiD estimate using an alternative treatment date of June 2011 rather than July 2011. The empty light blue square denotes the SDiD estimate dropping July through December 2009 during the Great Recession. Gray capped lines represent 95% confidence intervals. All models include The model includes income, household size, marital status, presence of children, race, education, and age controls. Standard errors are clustered at the state level for all specifications except the jackknife, which uses household-level clustering.

Tables

Table 1. Weighted pre-treatment means for Maryland and the donor pool

	Donor Pool Mean	Maryland Mean
Ethanol oz per adult	12.35	11.96
Beer ethanol oz per adult	4.48	3.12
Wine ethanol oz per adult	2.90	3.98
Spirits ethanol oz per adult	4.97	4.85
Low-intensity purchasers	0.34	0.34
Moderate-intensity purchasers	0.34	0.26
High-intensity purchasers	0.32	0.39
Price per oz ethanol, overall	0.64	0.66
Price per oz ethanol, beer	0.36	0.36
Price per oz ethanol, wine	0.39	0.42
Price per oz ethanol, spirits	0.18	0.23
Any alcohol purchase this month	0.44	0.42
Any beer purchase this month	0.24	0.20
Any wine purchase this month	0.21	0.22
Any spirits purchase this month	0.16	0.19
Household income $<$35k	0.28	0.20
Income: \$35,000-\$49,999	0.15	0.10
Income: \$50,000–\$69,999	0.17	0.14
Income: \$70,000-\$99,999	0.17	0.26
Income: \$100,000+	0.24	0.30
Head age: 25–44 Years	0.19	0.17
Head age: 45–64 Years	0.56	0.59
Head age: 65+ Years	0.25	0.25
Head education: $<$ High school graduate	0.02	0.00
Head education: High school graduate	0.28	0.22

Continued on next page

Table 1. Weighted pre-treatment means for Maryland and the donor pool (continued)

	Donor Pool Mean	Maryland Mean
Head education: Some college	0.31	0.28
Head education: College graduate+	0.38	0.50
1 person	0.28	0.31
2 persons	0.39	0.40
3+ persons, no children	0.11	0.08
3+ persons, with children	0.22	0.22
Married	0.46	0.45
With child	0.24	0.26
White	0.81	0.58
Asian	0.03	0.05
Black	0.10	0.32
Hispanic	0.11	0.03
Other race	0.07	0.06
Observations	250,000	4,920
Unique households	10,426	205

Notes. This table presents pre-treatment means for households in Maryland and comparison states. Each observation is at the household-month level. Comparison states exclude households in Connecticut, Illinois, New Jersey, New York, North Carolina, Washington, the District of Columbia, and Kansas, as these states experienced alcohol tax changes during the sample period. Values for demographic and household characteristics, including drinking levels, are expressed as proportions. For age and education, the reported values correspond to the highest level among male and female household heads when both are present. Data come from the NielsenIQ Consumer Panel, and means are weighted using the panel's projection factor.

Table 2. Impact of the Maryland 2011 sales tax on ethanol ounces purchased.

	4-year balanced panel	
	(1)	(2)
Treat x Post	-0.098***	-0.096***
	(0.006)	(0.006)
Percent effect	-9.4	-9.2
MD Pre-treatment mean	11.29	11.29
Observations	510,288	510,288
Household FE	YES	YES
Year-month FE	YES	YES
Household controls	NO	YES

Note. The table reports estimates from two model specifications using the same outcome variable: the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household. Column (1) presents estimates from the baseline synthetic difference-in-differences (SDiD) model. Column (2) presents estimates from the SDiD specification that additionally includes household demographic covariates: income, education, age, marital status, presence of children, and race. Both specifications include household and year-month fixed effects. Percent effect converts the IHS coefficient into an approximate percentage change using $\left[\exp\left(\hat{\beta} - \frac{1}{2}\widehat{\mathrm{SE}}(\hat{\beta})^2\right) - 1\right] \times 100$, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011. Standard errors are clustered at the state level. For SDiD estimates, standard errors are obtained via bootstrapping with 1,000 replications. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** p<0.01, ** p<0.05, * p<0.10.

Table 3. Impact of the Maryland 2011 sales tax on ounces of beer, wine and spirits purchased.

	Beer	Wine	Spirits
	(1)	(2)	(3)
Treat x Post	-0.034***	-0.055***	-0.046***
	(0.006)	(0.003)	(0.008)
Percent effect	-3.4	-5.3	-4.5
MD Pre-treatment mean	3.76	3.75	3.77

Note. Note. The table reports synthetic difference-in-differences (SDiD) estimates for three outcome variables: the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household for (1) beer, (2) wine, and (3) spirits. All specifications include household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 for each beverage category. Standard errors, reported in parentheses, are obtained via bootstrapping with 1,000 replications and clustered at the state level. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
** p<0.01, ** p<0.05, * p<0.10.

Table 4. Impact of the Maryland 2011 sales tax on ethanol purchases by pre-treatment purchasing intensity.

	Low	Moderate	High
	(1)	(2)	(3)
Treat x Post	-0.062***	-0.084***	-0.070***
	(0.005)	(0.018)	(0.014)
Percent effect	-6.0	-8.0	-6.8
MD Pre-treatment mean	0.58	2.89	21.86
Share of sample (%)	33.3	32.3	34.4

Note. The table reports synthetic difference-in-differences (SDiD) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household across three purchaser-intensity groups: (1) low-intensity purchasers, defined as households in the bottom tercile of the pre-treatment purchasing distribution; (2) moderate-intensity purchasers, defined as households in the middle tercile; and (3) high-intensity purchasers, defined as households in the top tercile. Purchaser groups are based on average monthly ethanol purchases prior to the tax increase and exclude anomalous purchasers, defined as households purchasing more than 200 ounces of ethanol in any single month. All specifications include household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each purchaser-intensity group. Share of sample reports the proportion of households in each purchaser-intensity group. Standard errors, reported in parentheses, are obtained via bootstrapping with 1,000 replications and clustered at the state level. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 5. Effects on purchase incidence

	(1) Overall	(2) Beer	(3) Wine	(4) Spirits
Panel A: Price per oz ethanol				
Treat x Post	-0.029***	-0.010***	-0.028***	-0.012***
	(0.003)	(0.002)	(0.003)	(0.002)
Percent effect	-2.9	-1.0	-2.8	-1.2
Pre-treatment mean	0.69	0.37	0.46	0.22
Panel B: Any purchase in month				
Treat x Post	-0.032***	-0.010***	-0.024***	-0.019***
	(0.002)	(0.002)	(0.001)	(0.002)
Pre-treatment mean	0.403	0.210	0.222	0.159
Panel C: Price per oz ethanol, conditional on purchase				
Treat x Post	0.045***	0.030***	0.045***	0.036***
	(0.004)	(0.003)	(0.006)	(0.004)
Pre-treatment mean	1.70	1.77	2.06	1.41

Note. The table presents estimates for three outcome variables: the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household (Panel A), an indicator for any alcohol purchase in a month (Panel B), and the tax-inclusive price per ounce of ethanol in months with positive purchases (Panel C). Panels A and B are estimated using synthetic difference-in-differences (SDiD), while Panel C is estimated using a two-way fixed effects model. Column (1) presents the overall estimate, column (2) the estimate for beer, column (3) for wine, and column (4) for spirits. All specifications include household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Percent effect reports the approximate percentage change implied by the IHS coefficient in Panel A, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean of each outcome variable in Maryland prior to July 2011. Standard errors, reported in parentheses, are clustered at the state level; for SDiD specifications, they are obtained via bootstrapping with 1,000 replications. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** p<0.01, ** p<0.05, * p<0.10.

Table 6. Impact of the Maryland 2011 sales tax on ethanol ounces purchased per adult in different time windows

	Unbalanced panel					
	Full		10-year panel		4-year panel	
	(1) HH FE	(2) State FE	(3) HH FE	(4) State FE	(5) HH FE	(6) State FE
Treat x Post	-0.068***	-0.067***	-0.056***	-0.050***	-0.060***	-0.049***
	(0.011)	(0.011)	(0.008)	(0.008)	(0.005)	(0.005)
Percent effect	-6.5	-6.5	-5.5	-4.8	-5.9	-4.8
MD Pre-treatment mean	4.70	4.70	4.58	4.58	4.64	4.64
Observations	10,844,478	10,844,478	7,175,864	7,175,864	2,935,548	2,935,548
Household FE	YES	NO	YES	NO	YES	NO
State FE	NO	YES	NO	YES	NO	YES
Year-month FE	YES	YES	YES	YES	YES	YES
Household controls	YES	YES	YES	YES	YES	YES

Note. The table reports two-way fixed effects estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household. Columns (1) and (2) present estimates using the full pre-COVID sample from 2004 to 2019. Columns (3) and (4) present estimates using a 10-year unbalanced panel from 2006 to 2016. Columns (5) and (6) present estimates using the main sample window from July 2009 to June 2013 while retaining all available households. Observations are at the household-month level. All specifications include household demographic covariates: income, education, age, marital status, presence of children, and race, as well as year-month fixed effects. Columns (1), (3), and (5) include state fixed effects, while columns (2), (4), and (6) include household fixed effects. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each sample window. Standard errors, reported in parentheses, are clustered at the state level. Data on alcohol purchases come from the NielsenIQ Consumer Panel.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$.

A Appendix

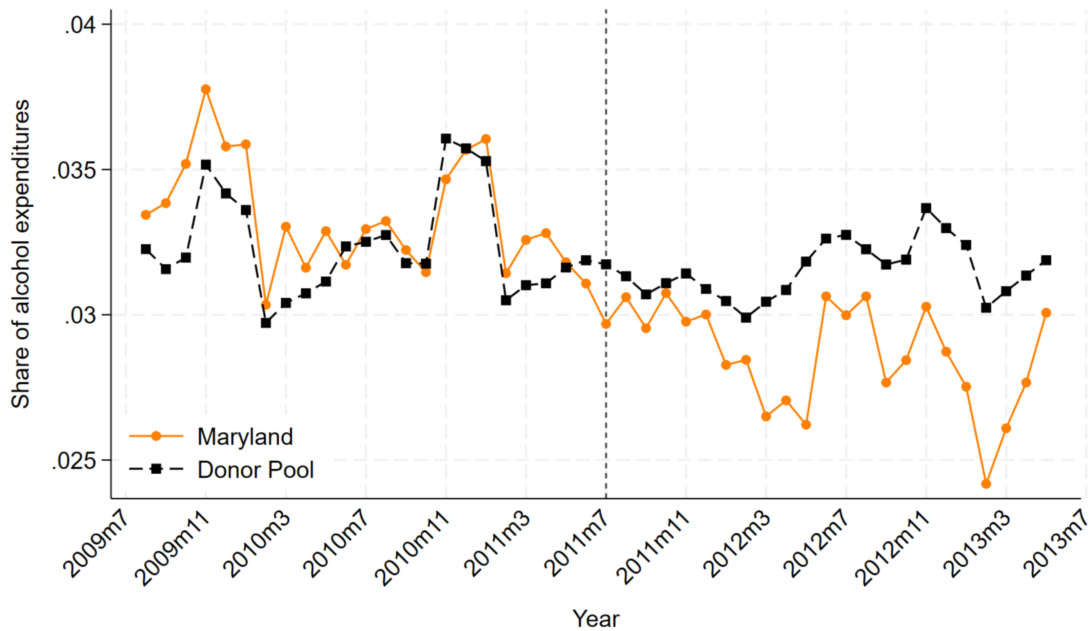

Figure A1. Trends in alcohol expenditure share in Maryland and donor pool.

Note. The y-axis shows the three-month moving average of alcohol expenditure as a share of total household expenditures, and the x-axis shows the month–year. The figure plots this share for Maryland (yellow circles, solid line) and for the comparison states (black squares, dashed line). The vertical dashed line denotes July 2011, when Maryland implemented the alcohol sales tax increase. Comparison states exclude Connecticut, Illinois, New Jersey, New York, North Carolina, Washington, the District of Columbia, and Kansas, which experienced alcohol tax changes during the sample period. Data are from the NielsenIQ Consumer Panel.

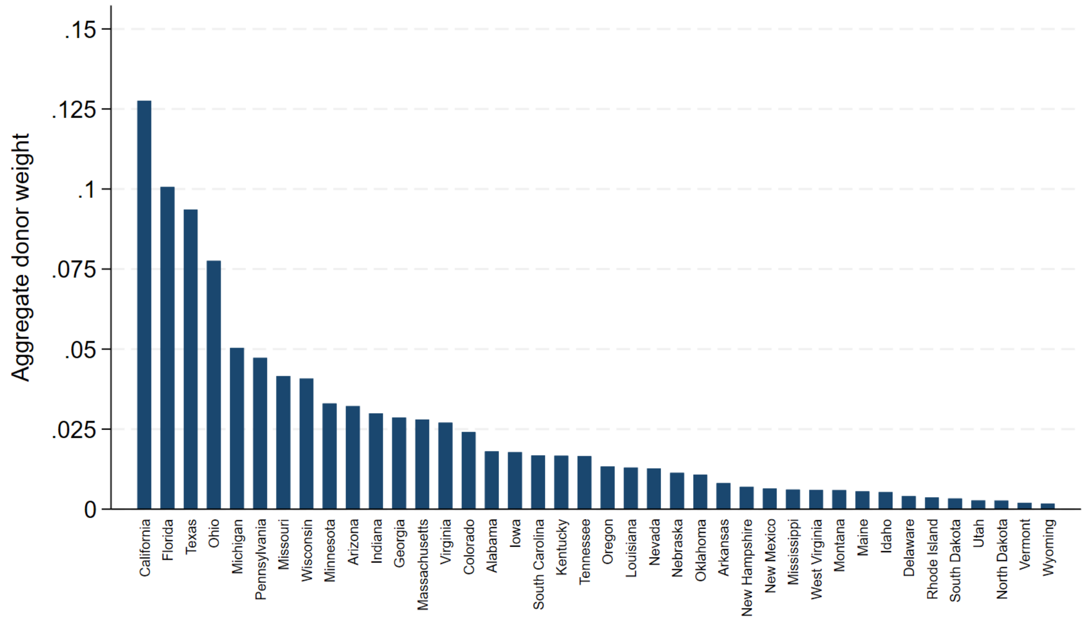

Figure A2. Aggregate SDiD donor weights by state.

Note. Household-level SDiD donor weights are aggregated to the state level by summing positive household weights within each state. The y-axis shows each state's aggregate donor weight, and the x-axis lists donor states with positive weights. The comparison group excludes Connecticut, Illinois, New Jersey, New York, North Carolina, Washington, the District of Columbia, and Kansas because these states experienced alcohol tax changes during the sample period. Hawaii and Alaska are not observed in the analytic sample. Data are from the NielsenIQ Consumer Panel.

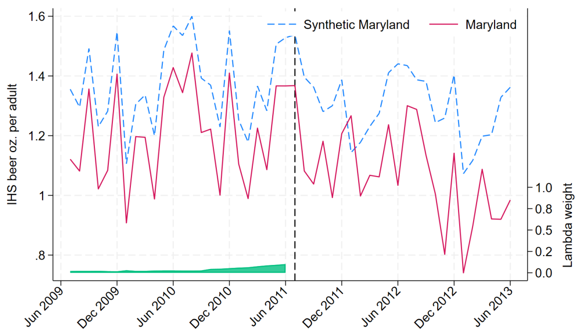

(a) Beer

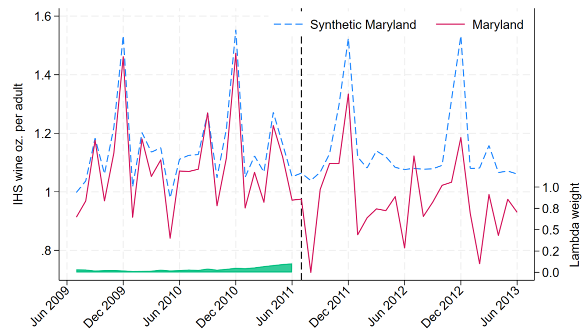

(b) Wine

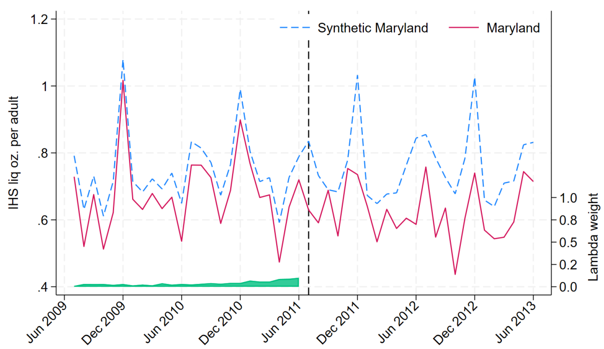

(c) Spirits

Figure A3. Trends in alcohol purchases by beverage for Maryland and synthetic control

Note. Trends in ounces of ethanol purchased per adult for Maryland households using Synthetic Difference-in-Differences (SDiD) weights, disaggregated by alcohol type: beer, wine, and spirits. Panel (a) reports trends for beer, panel (b) for wine, and panel (c) for spirits. The time weights (λ_t) are estimated by minimizing the error term in Equation 1. The solid line shows the trend for Maryland, and the dashed line shows the corresponding trend for the synthetic control.

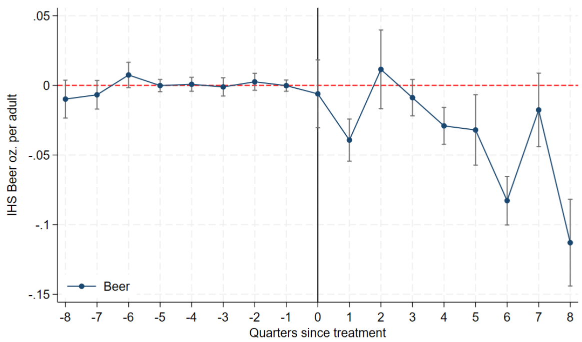

(a) Beer

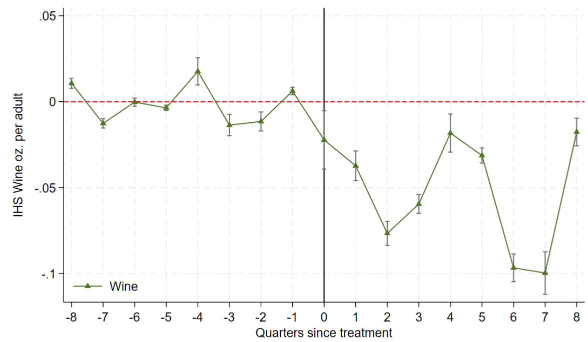

(b) Wine

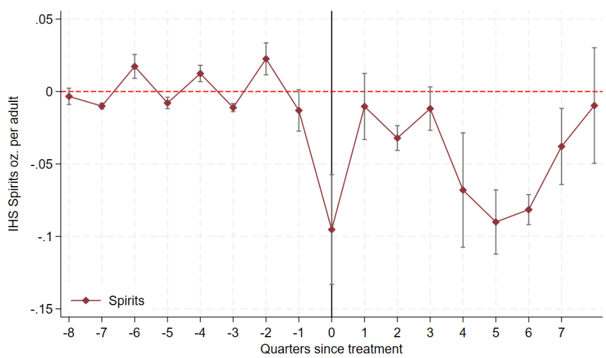

(c) Spirits

Figure A4. Event study estimates by alcohol type: beer, wine, and spirits

Note. The figure presents event study estimates of the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month by alcohol type, estimated using the Synthetic Difference-in-Differences (SDID) method. Monthly event-time coefficients were averaged to the quarter level for presentation. The vertical line at event time zero marks the month of the tax change. Panel (a) displays results for beer (blue circles), panel (b) for wine (green triangles), and panel (c) for spirits (maroon diamonds). Gray capped lines represent 95% confidence intervals estimated from 1,000 bootstrap replications. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period June 2009–July 2013.

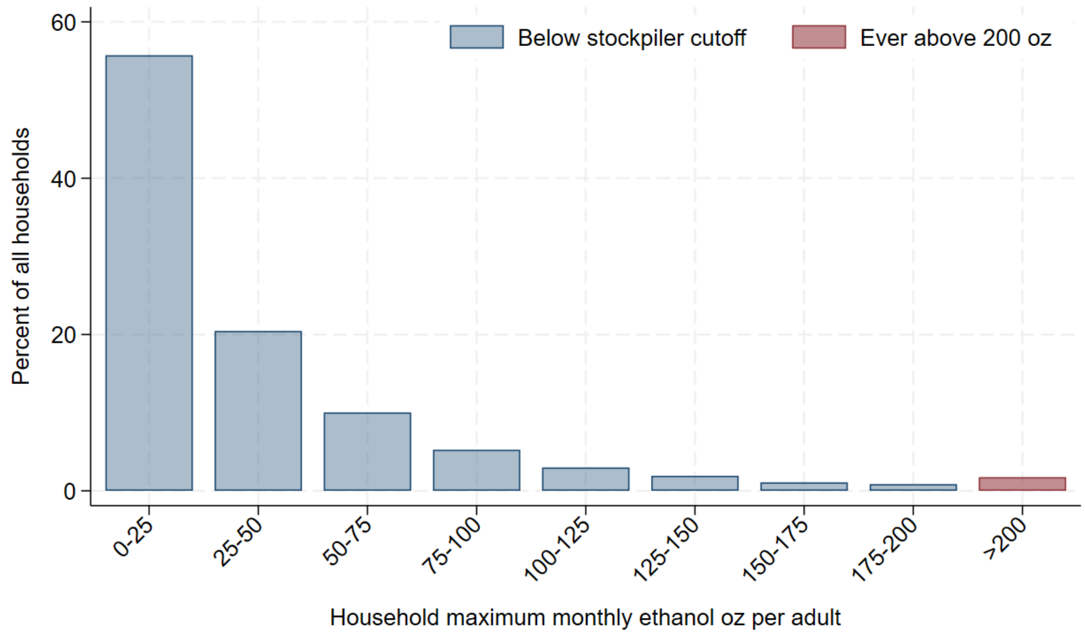

Figure A5. Distribution of maximum monthly ethanol purchases by household.

Note: This figure shows the distribution of each household's maximum monthly ethanol purchases per adult during the pre-treatment period. The unit of observation is the household, and bins report the share of all households whose largest pre-treatment monthly purchase falls in each range. Households with any month above 200 ethanol ounces per adult are highlighted in red and classified as anomalous purchaser or stockpilers. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the pre-treatment period.

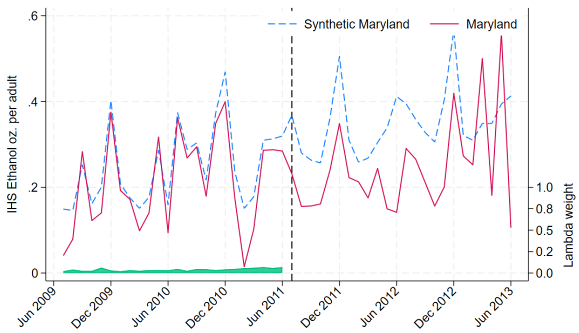

(a) Low-intensity

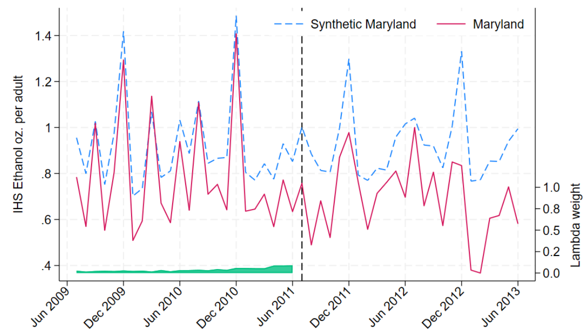

(b) Moderate-intensity

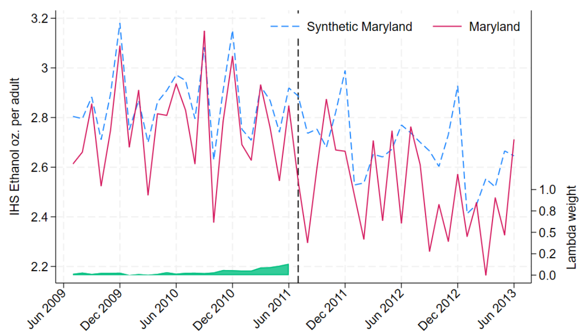

(c) High-intensity

Figure A6. Trends in ethanol purchases by baseline purchase intensity for Maryland and synthetic control

Note. Trends in ounces of ethanol purchased per adult for Maryland households using Synthetic Difference-in-Differences (SDiD) weights, disaggregated by heavy, moderate, light, and low-income drinkers.Panel (a) displays low-intensity purchasers (blue circles), panel (b) moderate-intensity purchasers (green triangles), and panel (c) high-intensity purchasers (maroon diamonds). Purchaser groups are defined using pre-treatment terciles of average monthly purchasing levels and exclude anomalous purchasers, defined as households purchasing more than 200 ounces of ethanol in any single month. The solid line shows the trend for Maryland, and the dashed line shows the corresponding trend for the synthetic control. The time weights (λ_t) are estimated by minimizing the error term in Equation 1. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period June 2009–July 2013.

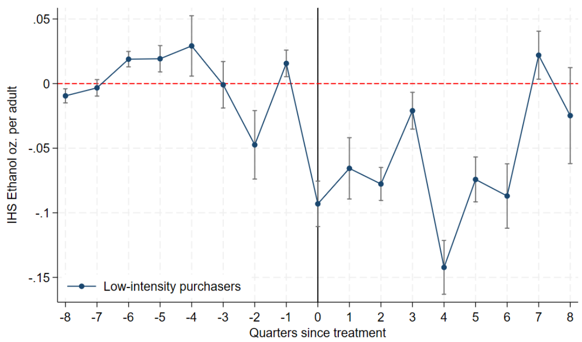

(a) Low-intensity

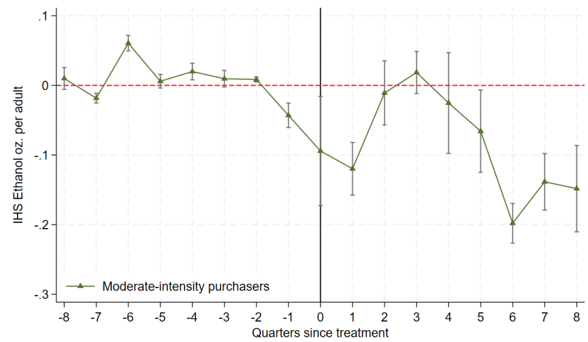

(b) Moderate-intensity

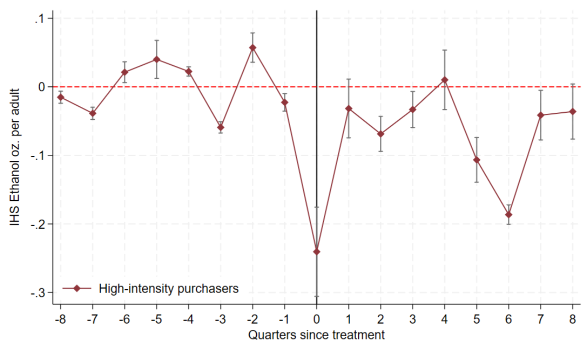

(c) High-intensity

Figure A7. Event study estimates by baseline purchase intensity

Note. The figure presents event-study estimates of the inverse hyperbolic sine transformation of monthly ethanol ounces purchased per adult per household, estimated using the Synthetic Difference-in-Differences (SDID) method. Monthly event-time coefficients were averaged to the quarter level for visual clarity. The vertical line at event time zero denotes the month of the tax change. Panel (a) displays estimates for low-intensity purchasers (blue circles), panel (b) for moderate-intensity purchasers (green triangles), and panel (c) for high-intensity purchasers (maroon diamonds). Purchaser groups are defined using pre-treatment terciles of average monthly purchasing levels and exclude anomalous purchasers, defined as households purchasing more than 200 ounces of ethanol in any single month. Gray capped lines indicate 95% confidence intervals based on 1,000 bootstrap replications. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period June 2009–July 2013.

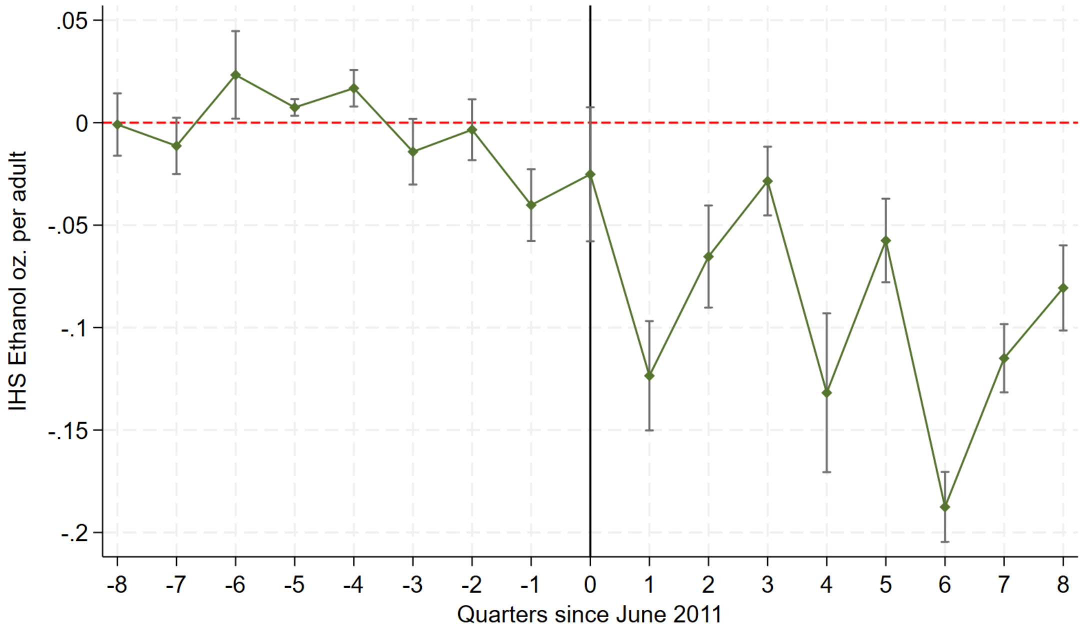

Figure A8. Anticipation event study.

Note. The figure shows treatment effects estimated using a synthetic difference-in-differences (SDiD) model, where June 2011 is used as the treatment date to assess potential anticipatory responses. The dependent variable is the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household. Monthly event-time coefficients are averaged to the quarter level for display. The vertical line at event time zero marks June 2011. Green circles indicate point estimates, and gray capped lines represent 95% confidence intervals. The model includes household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Standard errors are clustered at the state level and obtained via bootstrapping with 1,000 replications.

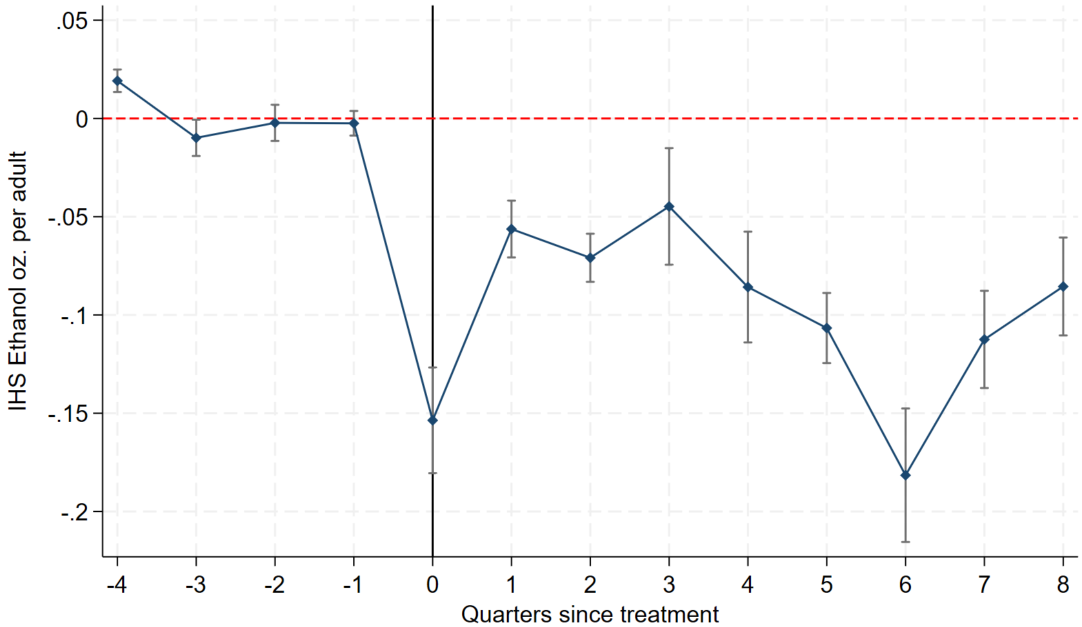

Figure A9. Great Recession robustness event study.

Note. The figure shows treatment effects estimated using a synthetic difference-in-differences (SDiD) model after excluding July through December 2009, which correspond to the early post–Great Recession recovery period. The dependent variable is the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household. Monthly event-time coefficients are averaged to the quarter level for display. The vertical line at event time zero marks the month of the tax change. Blue circles indicate point estimates, and gray capped lines represent 95% confidence intervals. The model includes household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Standard errors are clustered at the state level and obtained via bootstrapping with 1,000 replications.

Table A1. States with Alcohol Sales Tax Changes from June 2009 to July 2013

State	FIPS	Date	On-Premise	Off-Premise
District of Columbia	11	9/30/2009	↓ (4.25% to 4%)	↓ (3.25% to 3%)
Kansas	20	6/30/2010	↓ (4.70% to 3.70%)	↓ (2.70% to 1.70%)
Kansas	20	6/30/2013	↑ (3.70% to 3.85%)	↑ (1.70% to 1.85%)
Maryland	24	6/30/2011	↑ (3%)	↑ (3%)

Note. This table lists states that implemented changes in alcohol sales tax rates between 2009 and 2013, based on data from the Alcohol Policy Information System (APIS) (National Institute on Alcohol Abuse and Alcoholism, 2024). The table details the direction and magnitude of tax changes for on-premise (e.g., bars, restaurants) and off-premise (e.g., retail stores) alcohol purchases.

Table A2. States with Alcohol Excise Tax Increases from June 2009 to July 2013

State	FIPS	Date	Beer tax	Wine tax	Spirit tax
Connecticut	9	5/4/2011	$\uparrow$ \$0.04	$\uparrow$ \$0.12	$\uparrow$ \$1.00
Illinois	17	9/1/2009	$\uparrow$ \$0.04	$\uparrow$ \$0.66	$\uparrow$ \$4.05
New Jersey	9	8/1/2011	$\uparrow$ \$0.04	$\uparrow$ \$0.18	$\uparrow$ \$1.10
New York	36	5/1/2009	$\uparrow$ \$0.03	$\uparrow$ \$0.11	-
North Carolina	37	9/1/2009	$\uparrow$ \$0.09	$\uparrow$ \$0.21	-
Washington (FIPS 53) became a monopoly state on 12/8/2011.					

Note. This table lists states that implemented alcohol tax increases between 2009 and 2013, based on data from the Alcohol Policy Information System (APIS) (National Institute on Alcohol Abuse and Alcoholism, 2024). The table details the tax increase amounts for beer, wine, and spirits, where applicable.

Table A3. Impact of the Maryland 2011 sales tax on the natural log of ethanol ounces purchased.

	4-year balanced panel	
	(1)	(2)
	HH FE	State FE
Treat x Post	-0.080***	-0.079***
	(0.005)	(0.005)
MD Pre-treatment mean	11.29	11.29
Household FE	YES	YES
Year-month FE	YES	YES
Household controls	NO	YES

Note. The table reports estimates from two model specifications using the same outcome variable: the natural log of monthly ethanol ounces purchased per adult per household plus one. Column (1) presents estimates from the baseline synthetic difference-in-differences (SDiD) model. Column (2) presents estimates from the SDiD specification that additionally includes household demographic covariates: income, education, age, marital status, presence of children, and race. Both specifications include household and year-month fixed effects. Coefficients can be interpreted as approximate percentage changes. MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011. Standard errors are clustered at the state level; for SDiD estimates, they are obtained via bootstrapping with 1,000 replications. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013. *** p<0.01, ** p<0.05, * p<0.10.

Table A4. Impact of the Maryland 2011 sales tax on alcohol ounces purchased.

	4-year balanced panel	
	(1) HH FE	(2) State FE
Treat x Post	-0.150***	-0.153***
	(0.011)	(0.009)
Percent effect	-13.9	-14.2
MD Pre-treatment mean	113.47	113.47
Observations	510,288	510,288
Household FE	YES	YES
Year-month FE	YES	YES
Household controls	NO	YES

Note. The table reports estimates from two model specifications using the same outcome variable: the inverse hyperbolic sine of monthly alcohol ounces purchased per adult per household. Column (1) presents estimates from the baseline synthetic difference-in-differences (SDiD) model. Column (2) presents estimates from the SDiD specification that additionally includes household demographic covariates: income, education, age, marital status, presence of children, and race. Both specifications include household and year-month fixed effects. Coefficients can be interpreted as approximate percentage changes. MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011. Standard errors are clustered at the state level; for SDiD estimates, they are obtained via bootstrapping with 1,000 replications. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** p<0.01, ** p<0.05, * p<0.10.

Table A5. Impact of the Maryland 2011 sales tax on ethanol ounces purchased, Poisson specification.

	4-year balanced panel	
	(1) Baseline	(2) Covariates
Treat x Post	-0.091***	-0.101***
	(0.007)	(0.007)
MD Pre-treatment mean	11.29	11.29
Household FE	YES	YES
Year-month FE	YES	YES
Household controls	NO	YES

Note. The table reports estimates from two model specifications using the same outcome variable: monthly ethanol ounces purchased per adult per household. Column (1) presents estimates from a Poisson specification. Column (2) presents estimates from the Poisson specification that additionally includes household demographic covariates: income, education, age, marital status, presence of children, and race. Both specifications include household and year-month fixed effects. Coefficients can be interpreted as approximate percentage changes. MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011. Standard errors are clustered at the state level. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table A6. Impact of the Maryland 2011 sales tax on ethanol purchases by pre-treatment purchasing intensity (alternative thresholds).

	Low	Moderate	High
	(1)	(2)	(3)
Treat x Post	-0.046***	-0.079***	-0.047***
	(0.008)	(0.010)	(0.017)
Percent effect	-4.5	-7.6	-4.6
MD Pre-treatment mean	0.32	3.31	26.61
Share of sample (%)	21.5	52.8	25.6

Note. The table reports synthetic difference-in-differences (SDiD) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household across three purchaser-intensity groups defined using alternative thresholds: (1) low-intensity purchasers, defined as households below the 25th percentile of the pre-treatment purchasing distribution; (2) moderate-intensity purchasers, defined as households between the 25th and 75th percentiles; and (3) high-intensity purchasers, defined as households at or above the 75th percentile. Purchaser groups are based on average monthly ethanol purchases prior to the tax increase and exclude anomalous purchasers, defined as households purchasing more than 200 ounces of ethanol in any single month. All specifications include household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each purchaser-intensity group. Share of sample reports the proportion of households in each purchaser-intensity group.Standard errors, reported in parentheses, are clustered at the state level and obtained via bootstrapping with 1,000 replications. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period from July 2009 to June 2013.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table A7. Beverage substitution by purchaser-intensity group.

	Beer			Wine			Spirits		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	(Low)	(Mod)	(High)	(Low)	(Mod)	(High)	(Low)	(Mod)	(High)
Treat x Post	-0.005*	-0.022*	-0.034**	-0.025***	-0.058***	-0.049***	-0.048***	-0.034***	-0.054***
	(0.003)	(0.013)	(0.015)	(0.003)	(0.008)	(0.005)	(0.004)	(0.007)	(0.014)
Percent effect	-0.5	-2.2	-3.4	-2.5	-5.7	-4.8	-4.6	-3.4	-5.2
MD Pre-treatment mean	0.13	0.70	5.57	0.25	1.19	9.00	0.20	1.00	7.29
Share of sample (%)	33.3	32.3	34.4	33.3	32.3	34.4	33.3	32.3	34.4

Note. The table reports synthetic difference-in-differences (SDiD) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household by beverage type and purchaser-intensity group. Estimates are shown separately for beer, wine, and spirits. Within each beverage category, results are reported for three purchaser-intensity groups: low-intensity purchasers, defined as households in the bottom tercile of the pre-treatment purchasing distribution; moderate-intensity purchasers, defined as households in the middle tercile; and high-intensity purchasers, defined as households in the top tercile. Purchaser groups are based on average monthly ethanol purchases prior to the tax increase and exclude anomalous purchasers, defined as households purchasing more than 200 ounces of ethanol in any single month. All specifications include household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each beverage type and purchaser-intensity group. Share of sample reports the proportion of households in each purchaser-intensity group. Standard errors, reported in parentheses, are clustered at the state level and obtained via bootstrapping with 1,000 replications. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$

Table A8. Ethanol purchases across household income groups.

	(1) <35k	(2) $35-69k	(3) $70-99k	(4) $100k+
Treat x Post	0.120***	-0.164***	-0.031***	-0.266***
	(0.016)	(0.008)	(0.008)	(0.011)
Percent effect	12.8	-15.1	-3.1	-23.4
MD Pre-treatment mean	4.04	15.71	10.73	11.36
Share of sample (%)	14.1	26.8	32.7	26.3

Note. The table reports synthetic difference-in-differences (SDiD) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household across four income groups: (1) less than $35,000, (2) $35,000 to $69,999, (3) $70,000 to $99,999, and (4) $100,000 or more. All specifications include household demographic covariates: education, age, marital status, presence of children, and race. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each income group. Share of sample reports the proportion of households in each income group. Standard errors, reported in parentheses, are obtained via bootstrapping with 1,000 replications and clustered at the state level. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
** p<0.01, ** p<0.05, * p<0.10.

Table A9. Ethanol purchases across education groups.

	(1) High school or less	(2) Some college	(3) College+
Treat x Post	0.114***	0.011	-0.174***
	(0.007)	(0.008)	(0.008)
Percent effect	12.0	1.1	-16.0
MD Pre-treatment mean	12.60	10.90	11.18
Share of sample (%)	12.2	22.9	64.9

Note. The table reports synthetic difference-in-differences (SDiD) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household across three education groups: (1) high school or less, (2) some college, and (3) college or more. Education is defined based on the highest reported level of schooling of the household head. All specifications include household demographic covariates: income, age, marital status, presence of children, and race. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each education group. Share of sample reports the proportion of households in each education group. Standard errors, reported in parentheses, are obtained via bootstrapping with 1,000 replications and clustered at the state level. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** p<0.01, ** p<0.05, * p<0.10.

Table A10. Household composition heterogeneity.

	(1) 1 person	(2) 2 persons	(3) 3+p no child	(4) 3+p with child
Treat x Post	-0.152***	-0.061***	0.027	-0.170***
	(0.017)	(0.009)	(0.026)	(0.015)
Percent effect	-14.1	-5.9	2.7	-15.6
MD Pre-treatment mean	12.27	13.60	3.39	7.49
Share of sample (%)	27.3	47.3	9.8	15.6

Note. The table reports synthetic difference-in-differences (SDiD) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household across four household composition groups: (1) one-person households, (2) two-person households, (3) households with three or more members without children, and (4) households with three or more members with children. Household composition is defined based on the number of household members and the presence of children. All specifications include household demographic covariates: income, education, age, marital status, and race. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each household composition group. Share of sample reports the proportion of households in each group. Standard errors, reported in parentheses, are obtained via bootstrapping with 1,000 replications and clustered at the state level. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** p<0.01, ** p<0.05, * p<0.10.

Table A11. Race heterogeneity

	White	Black	Hispanic	Asian	Other
	(1)	(2)	(3)	(4)	(5)
Treat x Post	-0.085***	-0.133***	-0.456***	-0.041	-0.109*
	(0.009)	(0.006)	(0.104)	(0.034)	(0.066)
Percent effect	-8.1	-12.5	-37.0	-4.1	-10.6
MD Pre-treatment mean	13.32	6.19	21.29	2.98	15.54
Share of sample (%)	68.8	23.9	2.0	3.9	3.4

Note. The table reports synthetic difference-in-differences (SDiD) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household across five race and ethnicity groups: (1) White, (2) Black, (3) Hispanic, (4) Asian, and (5) Other. Race and ethnicity are defined based on the reported characteristics of the household head. All specifications include household demographic covariates: income, education, age, marital status, and presence of children. Percent effect reports the approximate percentage change implied by the IHS coefficient, following Bellemare and Wichman (2020). MD pre-treatment mean refers to the mean monthly ethanol ounces purchased per adult per household in Maryland prior to July 2011 within each race and ethnicity group. Share of sample reports the proportion of households in each group; shares do not sum to 100 percent because ethnicity is defined separately from race. Standard errors, reported in parentheses, are obtained via bootstrapping with 1,000 replications and clustered at the state level. Data on alcohol purchases are from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** p<0.01, ** p<0.05, * p<0.10.

Table A12. Maryland border counties.

County Name	**FIPS Code**
Allegany County	001
Calvert County	005
Caroline County	011
Carroll County	013
Cecil County	015
Charles County	017
Frederick County	021
Garrett County	023
Harford County	025
Kent County	029
Montgomery County	031
Prince George's County	033
St. Mary's County	037
Washington County	043
Worcester County	047

Note. Border counties correspond to Maryland counties that share a border with neighboring states or the District of Columbia. County FIPS codes correspond to the three-digit county identifiers.

Table A13. Border-county analysis within Maryland.

	(1) Ethanol (oz) per adult
MD Border x Post	0.040
	(0.086)
Pre-treatment mean: border counties	11.25
Pre-treatment mean: interior counties	11.38

Note. The table reports two-way fixed effects (TWFE) estimates for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household using a sample restricted to Maryland households. The coefficient on MD Border $\times$ Post captures the differential change in ethanol purchases for households located in border counties relative to those in non-border counties after the tax increase. The model includes household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Pre-treatment means are reported separately for border and interior counties. Standard errors, reported in parentheses, are clustered at the county level. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$.

B Appendix: Synthetic difference-in-difference

This appendix provides additional details on the SDiD specification, including the construction of the unit and time weights and the alternative inference procedures based on jackknife and placebo methods.

B.1 Unit and time weights

Equation 1 describes the error that SDiD minimizes. Part of this objective are the unit and time weights, $\hat{\omega}_h^{SDiD}$ and $\hat{\lambda}_t^{SDiD}$. The weights are chosen to minimize the pre-treatment prediction error. This appendix provides additional detail on the construction of these weights, following Arkhangelsky et al. (2021).

Unit weights. Unit weights are obtained by solving:

$$(\hat{\omega}_0, \hat{\omega}^{sdid}) = \arg \min_{\omega_0 \in \mathbb{R},\, \omega \in \Omega} \sum_{t=1}^{T_{pre}} \left(\omega_0 + \sum_{h=1}^{N_{co}} \omega_h Y_{ht} - \bar{Y}_t^{tr} \right)^2 + \zeta_{T_{pre}}^2 \|\omega\|_2^2, \tag{B1}$$

where

$$\bar{Y}_t^{tr} = \frac{1}{N_{tr}} \sum_{h=N_{co}+1}^{N} Y_{ht}. \tag{B2}$$

Here, $\|\omega\|_2$ is the Euclidean norm and ζ is a regularization described in Arkhangelsky et al. (2021). This procedure yields a vector of non-negative unit weights, $\hat{\omega}^{sdid}$, and an intercept ω_0. The weights ω_h ensure that the synthetic control created from untreated households matches Maryland's pre-treatment trajectory, while ω_0 allows for level differences between Maryland and its synthetic counterpart. Thus, SDiD matches trends, not necessarily looking for a perfect pre-treatment fit, unlike the synthetic control method, which requires exact pre-treatment fit.

Time weights. Time weights are obtained through an analogous optimization:

$$(\hat{\lambda}_0, \hat{\lambda}^{sdid}) = \arg \min_{\lambda_0 \in \mathbb{R},\, \lambda \in \Lambda} \sum_{h=1}^{N_{co}} \left(\lambda_0 + \sum_{t \leq T_{pre}} \lambda_t Y_{ht} - \bar{Y}_{post}^{tr} \right)^2 + \zeta_{N_{co}}^2 \|\lambda\|_2^2, \tag{B3}$$

where

$$\bar{Y}^{tr}_{post} = \frac{1}{T_{post}} \sum_{t>T_{pre}} \bar{Y}^{tr}_t. \tag{B4}$$

Intuitively, the time weights $\hat{\lambda}^{sdid}$ assign more weight to pre-treatment periods whose outcomes most closely resemble the post-treatment period.** This allows SDiD to create a synthetic "time path" that places greater emphasis on pre-tax months that are most predictive of post-tax behavior. In other words, while unit weights make the control group resemble Maryland, time weights make the pre-treatment periods used for comparison resemble the dynamics of the treated period. The resulting estimator therefore balances treated vs. control and early vs. late periods in a way that minimizes pre-treatment prediction error.

B.2 Jackknife inference

As an alternative to the bootstrap, I implement the jackknife variance estimator proposed by Arkhangelsky et al. (2021). This approach is computationally less intensive and provides conservative standard errors when the number of observed units is large.

Let $\hat{\tau}^{sdid}$ denote the treatment effect estimated from the full sample. The jackknife procedure proceeds as follows:

1. Sequentially remove one household i from the sample and re-estimate the SDID model, keeping the optimal unit and time weights $\hat{\omega}^{sdid}$ and $\hat{\lambda}^{sdid}$ fixed from the original estimation.

2. Denote the resulting estimate as $\hat{\tau}^{sdid}_{(-i)}$.

3. Compute the jackknife variance as

$$\hat{V}^{(jack)}_{\tau} = \frac{N-1}{N} \sum_{i=1}^{N} \left(\hat{\tau}^{sdid}_{(-i)} - \bar{\tau}^{sdid}\right)^2,$$

   where $\bar{\tau}^{sdid}$ is the mean of the leave-one-out estimates.

4. The standard error is then estimated as the square root of the variance:

$$\hat{SE}^{jack}_{\tau} = \sqrt{\hat{V}^{(jack)}_{\tau}}.$$

The jackknife is advantageous because it avoids recomputing the optimal weights in each iteration, substantially reducing computation time relative to the bootstrap. It tends to yield slightly larger (more conservative) confidence intervals, particularly when the number of treated units is small.

B.3 Placebo inference

I also implement the placebo, or permutation-based, inference procedure proposed by Arkhangelsky et al. (2021). This approach is particularly useful when inference is conducted with a small number of treated aggregate units, as in my setting where Maryland is the only treated state. Unlike the bootstrap and jackknife procedures, placebo inference constructs the sampling distribution of the estimator by randomly assigning placebo treatment only among untreated units.

Let Y^{co} denote the outcome matrix for control households and let N_{tr} denote the number of treated households in Maryland. The placebo procedure proceeds as follows:

1. Keep only households in the donor pool, excluding Maryland households from the placebo assignment procedure.
2. For each replication $b = 1, \ldots, B$ (in this case $B = 1,000$), randomly select N_{tr} households from the N_{co} control households without replacement to receive the placebo treatment.
3. Construct a placebo treatment matrix, $W_{co}^{(b)}$, that assigns the same treatment timing as the actual Maryland tax. In my application, placebo-treated households are assigned treatment beginning in July 2011.
4. Re-estimate the SDiD model using only control households, treating the randomly selected placebo households as if they were treated. Denote the resulting placebo estimate as $\hat{\tau}_{(b)}^{sdid}$.
5. Repeat this procedure $B = 1,000$ times to obtain a distribution of placebo estimates.
6. Compute the placebo variance as

$$\hat{V}_{\tau}^{(placebo)} = \frac{1}{B} \sum_{b=1}^{B} \left(\hat{\tau}_{(b)}^{sdid} - \bar{\tau}^{placebo} \right)^2 ,$$

where

$$\bar{\tau}^{placebo} = \frac{1}{B}\sum_{b=1}^{B}\hat{\tau}_{(b)}^{sdid}.$$

7. The standard error is then estimated as the square root of the variance:

$$\hat{SE}_{\tau}^{placebo} = \sqrt{\hat{V}_{\tau}^{(placebo)}}.$$

Intuitively, this procedure asks how large the estimated SDiD effect would be if the Maryland treatment were reassigned to untreated households that did not actually experience the tax change. If the Maryland estimate is large relative to the distribution of placebo estimates, this provides evidence that the estimated effect is unlikely to be driven by random variation in untreated households' alcohol purchases.

An important limitation is that placebo inference relies on variation from placebo assignments within the donor pool. As discussed by Arkhangelsky et al. (2021) and Clarke et al. (2024), this procedure requires the number of control units to exceed the number of treated units and relies on a homoskedasticity assumption across units. For this reason, I report placebo inference as an additional robustness check rather than as the sole basis for inference.

C Appendix: Alternative methods

My primary specification estimates the effect of the alcohol sales tax using an estimator designed to minimize the error structure described in equation 1. To assess the robustness of these results, I also re-estimate the model using alternative procedures for constructing the control group. Each approach generates a different set of comparison weights and carries its own strengths and limitations. The following sections describe these weighting strategies and the assumptions underlying each.

C.1 Synthetic Control Method

First, I use the synthetic control method (SCM) suggested by Abadie and Gardeazabal (2003), where I estimate a DiD model that minimizes the following error (using terminology from Arkhangelsky et al. (2021)):

$$\arg\min_{\lambda,\mu,\beta}\Big\{\sum_{s=1}^{S}\sum_{t=1}^{T}(Y_{ht}-\mu-\lambda_t-W_{ht}\tau)^2\,\hat{\omega}_h^{sc}\Big\} \tag{C1}$$

where $\hat{\omega}_h^{sc}$ is a vector of household control weights statistically-derived according to Abadie et al. (2010). Equation C1 has similar parameters to equation 1 but it omits the unit (household) weights.This technique constructs a synthetic control group that approximates the outcome in the treatment group as closely as possible by selecting a weighted combination of untreated units based on pre-treatment characteristics.

The SCM is used for comparative case study cases where there is a single treated unit. In such cases, a combination of untreated units might be a better suited comparison than a single treated unit (Abadie and Gardeazabal, 2003; Abadie, 2021; Abadie et al., 2010). Although Maryland is the only treated state, the unit of observation is at the household-month level, therefore there are several treated households. Since the SCM assumes only one treated unit, it omits the inclusion of household level intercepts which help account for time-invariant differences in alcohol consumption across households.

C.2 Propensity-score-matched difference-in-differences

I also estimate a propensity-score-matched difference-in-differences specification. This approach constructs a comparison group by matching Maryland households to households in untreated states with similar observed pre-treatment characteristics, and then estimates a conventional two-way fixed-effects DiD model on the matched sample.

The matching procedure proceeds in two steps. First, I estimate propensity scores using the logistic regression model:

$$P(Treatment_h = 1 \mid X_h) = \frac{e^{X_h'\beta}}{1+e^{X_h'\beta}}, \tag{C2}$$

where T_h is an indicator for Maryland households and X_h includes the household's mean pre-treatment IHS ethanol purchases per adult and household demographic covariates, including income, race/ethnicity, household size, and education. Matching is based only on pre-treatment characteristics so that post-tax outcomes do not influence the construction of the comparison group.

For each Maryland household, I select one nearest neighbor from the donor pool based on the closest propensity score (1:1 matching without replacement), imposing a common-support restriction. This process yields a matched sample that is balanced on observed characteristics. Using the matched sample, I then estimate the DiD specification that minimizes the following error:

$$\arg\min_{\lambda,\gamma,\mu,\tau}\left\{\sum_{h=1}^{N}\sum_{t=1}^{T}\left(Y_{ht}-\mu-\gamma_h-\lambda_t-W_{ht}\tau\right)^2\right\}. \tag{C3}$$

Equation C3 includes the same treatment indicator, household fixed effects, and month-year fixed effects as Equation 1, but omits the SDiD unit and time weights. The credibility of this approach therefore depends on the quality of the matched comparison group. I assess match quality using covariate balance diagnostics. Following Imbens and Rubin (2015), standardized differences below 0.25 are interpreted as evidence of adequate balance. Figure C1 presents visual diagnostics, including standardized percentage bias and the variance ratio of covariates orthogonal to the propensity-score index. A variance ratio close to one indicates stronger balance (Rubin, 2001).

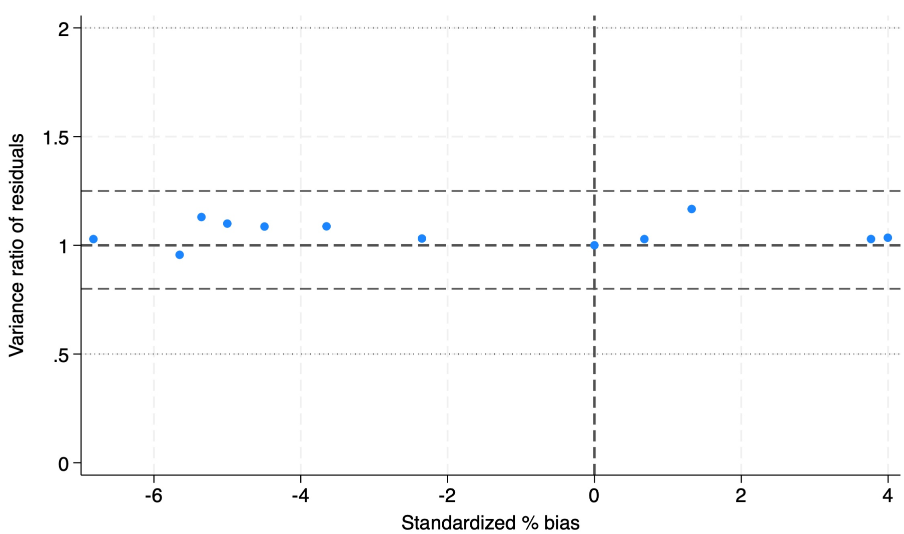

Figure C1. Standardized Bias and Variance ratio of the residuals

Note. The figure plots standardized percentage bias against the residual variance ratio for covariates orthogonal to the linear propensity-score index. Standardized differences are expressed as percentage bias. Following Rubin (2001), the variance ratio compares the variance of the residualized covariates between treated and matched control households.