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

Ana Paula Milán Hinostroza†

October, 2026

Preliminary and Incomplete. Please do not cite or circulate.

Abstract

The majority of the literature on alcohol taxation focuses on volume-based excise taxes, but not much is known about how consumers respond to alcohol-specific sales taxes. Although economic theory suggests that both types of taxes should reduce alcohol purchases, behavioral responses might differ due to the distinct tax structures. Specifically, sales taxes are calculated as a percentage of the retail price, are added only at the register, and increase automatically as prices rise. To understand the impact of alcohol-specific sales taxes on off-site alcohol purchases, I study Maryland's 2011 increase in its alcohol sales tax from 6 to 9 percent. Using NielsenIQ Consumer Panel data and a synthetic difference-in-differences design, I find that, after the tax increase, monthly ethanol purchases per adult decline by approximately 5.5 percent in the full sample and 9.2 percent among households that purchased alcohol in the year prior to the tax increase. These estimates correspond to approximately 0.6 and 1.7 fewer standard drinks per adult per month. Purchases decline across beer, wine, and spirits, with no evidence of substitution towards a particular alcohol category. The reductions are consistent throughout the pre-tax purchasing distribution and are broadly similar across demographic groups, with the largest declines among the pre-tax heaviest purchasers. These findings provide new evidence on the behavioral effects of alcohol sales taxes and suggest that price-based alcohol policies can produce reductions in alcohol purchases and can be used as a complement for traditional excise taxes.

*Acknowledgments. I would like to express my gratitude to Brady Horn, Xiaoxue Li, and Dave Dixon for their mentorship and guidance. I also thank participants at the NBER Health and Healthcare Variations Conference, the Western Economic Association International Graduate Student Workshop, and ASHEcon 2026 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 (US) 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 associated with excessive drinking (see Wagenaar et al. (2009) for a review of the literature), 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 included in shelf prices, whereas sales taxes are applied at the point of sale and automatically adjust with inflation.

These differences in the tax mechanism have implications for both the economic incidence of the tax and its behavioral effects, as consumers may respond differently depending on what taxes they see when shopping. In particular, Chetty et al. (2009) find that consumers respond less to taxes excluded from posted prices and that increases in alcohol excise taxes reduce consumption more than comparable increases in sales taxes. Their findings suggest that the visibility of the full price of the product can influence the effectiveness of alcohol taxation. However, they do not specifically study alcohol sales taxes.

From standard economic theory, both excise and sales taxes increase the full price of alcohol and *should* reduce consumption. Lower salience *may* attenuate responses to sales tax increases, but it does not imply that purchases will remain unchanged (Chetty et al., 2009). In particular, an alcohol-specific sales tax increase raises the full price of alcohol relative to other goods. Understanding whether and to what extent alcohol sales taxes reduce alcohol purchases, therefore, remains an unanswered empirical question.

To examine this, I focus on Maryland's 2011 increase in its alcohol sales tax from 6% to 9%. This was the state's first such hike in four decades. Anecdotal evidence suggests the change was adopted primarily to align Maryland's tax rate with those of neighboring

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

states, which supports treating the increase as plausibly exogenous to underlying alcohol consumption trends.[2] 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).[3]

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 take advantage of the panel structure and exploit within household variation to compare changes in ethanol ounces purchased per adult per household in Maryland relative to a weighted set of control households from other states. I do so using a synthetic difference-in-differences approach.[4] Importantly, I estimate the effect of the tax on alcohol purchases rather than actual consumption.[5]

I estimate that the tax increase led to a sustained 5.5% reduction in ethanol ounces purchased per adult per month in the full sample and a 9.2% reduction among households that purchased alcohol in the 12 months preceding the tax increase, which is my preferred sample for the remaining analyses. These reductions correspond to approximately 0.33 and 1.04 fewer ounces of ethanol per adult per month, respectively, or roughly 0.6 and 1.7 fewer standard drinks. This result is robust to different model specifications and a battery of other robustness checks.

Purchases decline across beer, wine, and spirits, while purchases of non-alcoholic beer and wine increase slightly, with no economically meaningful changes in tobacco-related products. 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.

Finally, I also look at alcohol spending as a secondary outcome, and I find that adults in households in the full sample and households that purchased alcohol in the 12 months preceding the tax reduced their spending after the tax increase. In addition, the share spent on alcohol also declined after the tax increase. Note that these results capture changes in

[2]This interpretation is further supported by the fact that ethanol purchases in Maryland, although volatile, did not show a sustained increasing trend in the years preceding the tax increase (see Appendix Figure A1).

[3]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.

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

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

pretax spending. Taking into account the higher sales tax would partly offset the reduction in pretax expenditure, so these estimates do not directly represent out-of-pocket reductions.

Overall, the findings suggest that, like excise taxes, alcohol sales taxes can reduce alcohol purchases. 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. 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.

These findings contribute to the alcohol-tax literature in two ways. First, to my knowledge, this paper provides the first household-level economic analysis of how an alcohol-specific sales tax affects alcohol purchases. Second, by using household-level panel data, I am able to exploit within household variation to explore whether responses differ across beverage types, the pre-tax purchasing distribution, and household demographic characteristics.

The paper proceeds as follows: Section 2 describes the institutional background of alcohol taxation in the US and Maryland's 2011 tax increase. Section 3 situates the paper relative to the alcohol taxation and tax salience literature. Sections 4 through 8 describe the data, empirical strategy, results, discussion and conclusion.

2 Institutional 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 US, 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. That is, they generate 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.[6]

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

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

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

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 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.[7]

Montgomery County, unlike other Maryland counties, differs from the rest of the state as it operates under a county-controlled alcohol distribution system, although the statewide sales tax increase applied there as well.

3 Related Literature

3.1 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). In addition, 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.,

[7]Figure A1 plots average monthly ethanol purchases for Maryland before the tax increase and shows no pronounced changes in purchasing patterns leading to the tax change.

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 alcohol to maintain their consumption levels.

Price changes may also affect purchases outside alcoholic beverages. Evidence on the relationship between alcohol and tobacco is mixed. Goel et al. (2016) find that beer and cigarettes behave as complements, while Pryce et al. (2024) find no consistent pattern in cross-price effects across alcohol and tobacco products. Nonalcoholic beer may be another potential alternative; in a hypothetical purchasing experiment, Acuff et al. (2026) find stronger evidence of substitution away from alcoholic beer when nonalcoholic beer is available at a lower price.

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). 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). However, nearly all the evidence above comes from the excise tax literature, and may differ from alcohol sales tax responses which have not been studied in such detail.

3.1.1 Contributions to the Literature

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 buyers, both taxes operate through the same price mechanism and should generate similar reductions in purchases. However, in practice, they differ along several dimensions that may affect behavior. Excise taxes are typically embedded in posted prices (the price that the consumer sees on the shelf when shopping), 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.

Chetty et al. (2009) study how salience differs between excise and sales taxes. They run a supermarket experiment and complement it with aggregate beer consumption data. They find smaller responses to sales taxes compared to similar increases in excise taxes. However, my study is different from this. First, I look at a specific sales tax, and, more importantly, I study a different question. I am specifically interested in how household alcohol purchases respond to a sales tax increase specifically targeting alcoholic beverages. That is, although salience is an important consideration, I am not comparing which tax mechanism reduces purchases more.

This paper contributes to the literature by providing, to my knowledge, the first household-level analysis of how an alcohol sales tax affects alcohol purchases. Previous research on Maryland's 2011 reform has primarily examined downstream health and behavioral outcomes. For example, 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. Although these findings are consistent with a reduction in alcohol use, these studies do not directly observe changes in alcohol purchases.

The closest study of the Maryland reform is Esser et al. (2016), who analyze county-level alcohol sales and find modest declines in per capita sales across beverage categories. However, they estimate post-tax changes using a mixed-effects model without a control group, and their aggregate data do not permit an examination of household-level responses. 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 analyses.

This study also relates closely to 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 may generate different behavioral responses. Second, unlike Saffer et al. (2024), who aggregate the data into state-level treatment-control differences, I apply synthetic difference-in-differences to the household-month panel. This approach combines unit and time weighting with household and year-month fixed effects, allowing me to account for differences between Maryland and the comparison group while retaining household-level variation in alcohol purchases.

4 Data

4.1 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.[8,9]

My analysis focuses on alcoholic beverage purchases. Because the data only capture off-premise alcohol sales, I effectively examine the tax's impact on alcohol bought for home consumption rather than on-premise sales (e.g., restaurants and bars). I aggregate alcohol purchases 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

[8]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.

[9]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.

post-tax observations. I exclude data from earlier years 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.[10] I will also show in Section 6 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.[11]

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.

I consider both the full sample of households and a drinking sample restricted to households that made at least one alcohol purchase in the 12 months preceding the tax increase. This restriction focuses the analysis on households with demonstrated pre-tax alcohol purchasing behavior, while the full sample also includes households with no observed alcohol purchases in the 12 months preceding the tax increase. Both samples include households from 41 states, including Maryland. The full sample consists of 17,255 households, with 385 in Maryland and 16,870 in the control states, while the drinking sample consists of 10,631 households, with 205 in Maryland and 10,426 in the control states.[12]

4.1.1 Alcohol Measures

I look at the quantity of alcohol purchased and how much consumers spent on alcohol. To evaluate quantity, rather than using total ounces of alcohol purchased, I focus on ethanol

[10] 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.

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

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

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.[13,14] I then divide total monthly ethanol ounces by the number of adults in the household to create an ethanol ounces per adult measure. In addition to total ethanol purchases, I construct analogous measures separately for beer, wine, and spirits.

To examine changes in alcohol spending, I construct two expenditure measures. The first is monthly alcohol expenditure per adult (adjusted to 2011 dollars), and the second is the share spent on alcohol, with the denominator being the total recorded monthly expenditure (including both alcohol and non-alcohol purchases). The latter measure is at the household level. The NielsenIQ price data exclude sales taxes; therefore, both expenditure measures represent pre-tax spending.

4.1.2 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.[15] 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

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

[14]Appendix Table A3 Panel B reports results using total volume (in ounces) rather than ethanol-equivalent ounces.

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

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.[16] 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.[17]

4.1.3 Other Measures

In addition to alcohol purchases, I construct measures of purchases of non-alcoholic beer and wine and tobacco products. Similar to the alcohol beverage measures, all continuous quantity outcomes are aggregated to the household-month level and divided by the number of adults in the household. For non-alcoholic beer and wine, the outcome is ounces purchased per adult in each household-month.[18] For cigarettes and cigars, the outcomes are counts of the respective products purchased per adult in each household-month. For chewing and smoking tobacco, the outcomes are ounces purchased per adult in each household-month. I also construct an extensive-margin outcome, any tobacco purchase, which is an indicator equal to one if the household purchased any tobacco product during that month.

5 Methods

5.1 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

[16] Appendix Table A4 reports robustness to alternative purchaser-intensity cutoffs.

[17] This is based on the drinking sample as this classification conditions households on having at least one pre-treatment month where they purchased alcohol.

[18] Volumes of non-alcoholic beer and wine are standardized and converted to fluid ounces.

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 a 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, this is relevant in this context, as households differ systematically in baseline alcohol consumption, purchasing frequency, and responsiveness to price changes.

Consistent with this intuition, I estimate a 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.[19] I estimate analogous specifications for alternative outcomes, including measures of ethanol prices. Because ounces of ethanol purchased per 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.[20] Although nonlinear models such as Poisson 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 es-

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

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

timator. Panel C in Appendix Table A3 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.[21] Standard errors are clustered at the state level. I also include household demographic covariates such as household size, income, race, age of the household head, and education of the household head. [22]

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 6.

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 improve visual clarity, as quarterly effects provide a smoother depiction of post-treatment dynamics. This event-study analysis serves two purposes: (1) to confirm that Maryland's pre-tax purchasing trends were closely aligned with its synthetic control; and (2) to illustrate the persistence of the treatment effect during the post treatment period. Having outlined the SDiD estimator and its dynamic extension, I next turn to the inference procedures I used to obtain valid standard errors for these estimates.

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

[22] 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 $Y_{ht}^{res} = Y_{ht} - X_{ht}\hat{\beta}$.

5.2 Inference

Following Arkhangelsky et al. (2021), I compute standard errors using a state-clustered bootstrap. Bootstrap inference approximates the sampling distribution of an estimator by repeatedly drawing samples with replacement from the observed data and re-estimating the model in each sample.

In this setting, alcohol sales taxes vary at the state level, and households within the same state may be exposed to common economic and policy shocks. Resampling households independently could therefore understate uncertainty by treating households within a state as independent. I accordingly adapt the bootstrap procedure in Arkhangelsky et al. (2021) by resampling states, rather than households, with replacement. The procedure is as follows:

1. Draw B bootstrap resamples of states (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 $\hat{\tau}^{(b)}_{sdid}$.
3. Define the bootstrap variance as:

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

   where $\bar{\tau}^{boot} = \frac{1}{B}\sum_{b=1}^{B}\hat{\tau}^{sdid}_{(b)}$ 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. This procedure approximates the sampling distribution of the SDID estimator by repeatedly resampling complete state clusters and re-estimating both the SDID weights and the treatment effect. However, because Maryland is the only treated state, resampling it over and over does not introduce new variation in shocks specific to Maryland. Thus, the state-clustered bootstrap accounts for dependence among households within states.[23]

[23] I do not use the placebo inference for the main estimator because the analytical units are households, whereas treatment is assigned at the state level and states contain different numbers of households. I instead use the state-clustered bootstrap as the primary inference procedure and report household-level placebo inference as a robustness check.

5.3 Alternative model specifications

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

First, I look at three alternative sample windows that do not require households to be continuously observed: the full pre-COVID-19 NielsenIQ sample from 2004 to 2019, a 10-year sample from 2006 to 2016, and the main four-year window from July 2009 to June 2013 with all available households retained. Because these specifications produce unbalanced household panels, the SDiD framework is not applicable. I therefore estimate two-way fixed-effects models with either state or household fixed effects. I view these models as complementary, directional checks under alternative sample definitions as they do not reproduce the identifying variation or weighting used by my preferred SDiD design.

I then report estimates using alternative inference procedures. The main specification uses bootstrap standard errors obtained by resampling households while preserving their complete panel observations. This procedure is computationally intensive but captures sampling variation in the composition of the household panel but does not capture variation in state-level shocks specific to Maryland. I therefore also report standard errors using the jackknife and placebo inference procedures proposed by Arkhangelsky et al. (2021).[24]

I also estimate the effect using SCM and DiD estimates for comparison, although SDiD has been shown to outperform these approaches. The synthetic control method is well suited to settings with a single treated state because it constructs a weighted combination of untreated states that approximates the treated state's pre-treatment alcohol purchasing path. 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 model on the matched sample. This provides a complementary estimate based on observed covariate balance rather than SDiD unit and time weighting.[25]

I 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

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

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

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 estimate a donut specification that excludes June and July 2011 to determine whether purchases immediately surrounding the tax increase drive the estimated effect.

Then, I look at 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 also test for 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 controls or, conversely, on comparisons with more geographically distant states.

6 Results

Table 1 reports weighted pre-tax means for outcomes and household characteristics in Maryland and the donor pool for both the full sample and the drinking sample.[26] I use the NielsenIQ Consumer Panel projection factor as a sample weight to improve national representativeness.[27] The weighted means suggest that Maryland and the donor pool are broadly similar across most pre-treatment characteristics in both samples. Maryland households purchase less beer and more wine, face similar average prices, and have higher income and education levels, a larger share of Black households, and smaller shares of White and Hispanic households than households in the donor pool.

[26] As described in Section 4, the drinking sample is defined based on alcohol-purchasing behavior in the 12 months preceding the tax increase.

[27] 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 and the donor pool for the full sample in Panel 1a and the drinking sample in Panel 1b. In both samples, Maryland and the donor pool have similar pre-treatment trends in ethanol purchases. Following the tax increase, purchases in Maryland decline relative to the donor pool. This difference is modest and concentrated in the later post-treatment periods in the full sample, but is more pronounced in the drinking sample. These trends provide descriptive evidence of a reduction in ethanol purchases following the tax increase, particularly among households that purchased alcohol before the policy.

[Figure 1 here]

To formally estimate the effect of the tax suggested by the trends in Figure 1, I estimate the SDiD model in equation 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.[28] Figure 2 plots observed outcomes for Maryland and the synthetic group, together with the estimated time weights. Panel 2a presents the full sample, while Panel 2b presents the drinking sample. The green bars represent the weights assigned to pre-treatment months. Following the tax increase, purchases in Maryland appear to fall slightly relative to the synthetic comparison in the full sample, with the difference somewhat more visible in the drinking sample. These patterns are consistent with a reduction in ethanol purchases following the tax increase.

[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 for the full sample of households (columns 1 and 2) and the drinking sample (columns 3 and 4). Columns (2) and (4) present the estimates controlling for household-level covariates discussed in Section C. All specifications include household and year-month fixed effects.

[28] 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.

Across all 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).[29] For example, column (1) reports a coefficient of -0.032, corresponding to an approximate 3.1% reduction in monthly ethanol ounces purchased per adult in Maryland households relative to the pre-tax period. Similarly, column (2) reports at 5.5% decrease after controlling for household-level characteristics. This percent decrease translates into 0.33 fewer monthly ounces of ethanol per adult per household.

[Table 2 here]

The magnitude of the decrease is larger among drinking households, with columns (3) and (4) presenting a reduction of 9.4% and 9.2% in monthly ethanol ounces purchased per adult in Maryland households relative to the pre-tax period respectively. Column (4) is my preferred specification, as it includes important time-invariant demographic characteristics associated with alcohol purchases, including household income, education, age group, and household size, and uses the drinking sample, which represents the population directly exposed to the tax. Including these covariates helps account for observable differences between Maryland and the donor pool while leaving the main results substantively unchanged.

The estimate in column (4) translates into 1.04 fewer ethanol ounces purchased per adult per month. Given that the estimates are not sensitive to the composition of the analytic sample, specifically whether non-drinking households are included or excluded, all subsequent analyses use the drinking sample and include household-level covariates unless noted otherwise.

For the estimated 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 have similar pre-treatment trends. To further assess this identifying assumption, I estimate the event-study version of equation 1 using my preferred specification.[30]

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,

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

[30] Appendix Figure A3 presents the analogous event study for the full sample that includes households that did not purchase alcohol during the 12 months preceding the tax increase.

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. Consistent with the raw trends, pre-treatment coefficients remain close to zero, and post-treatment coefficients show a sustained decline in ethanol ounces purchased per adult, especially at the start of treatment.[31]

[Figure 3 here]

6.1 Beverage Types

I next examine whether this reduction was concentrated in particular beverage categories. 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]

There are significant and consistent reductions in ethanol ounces purchased per adult across all beverage type. 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.[32] The results do not provide evidence of substitution toward any single beverage type. Appendix Figure A5 shows the SDiD pre-treatment fit by alcohol type, and Appendix Figure A6 plots the event-study estimates.

[31] Appendix Figure A4 alternatively presents the event study estimates when then June 2011 is consider the start of treatment rather than July 2011.

[32] 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.

6.2 Potential cross-product responses

To explore whether the reduction in alcohol purchases coincided with changes in purchases of other products, I apply the main SDiD specification to non-alcoholic beer and wine and to tobacco products. Non-alcoholic beer and wine may be considered close alternatives to alcoholic beverages, while tobacco and nicotine products provide evidence on whether the tax generated broader changes in purchases of other potentially related goods. Table 4 presents estimates of the impact of the tax on non-alcoholic beer and wine (Panel A) and tobacco products (Panel B). Column (1) presents the treatment effect estimate, column (2) the standard error, column (3) the approximate percent effect calculated following Bellemare and Wichman (2020), and column (4) Maryland's mean ounces or units purchased before the tax increase.

[Table 4 here]

Panel A presents the results for non-alcoholic beer and wine, measured in ounces purchased per adult. Purchases of non-alcoholic beer and wine significantly increased by approximately 2% which corresponds to about 0.014 additional ounces per adult relative to the pre-treatment mean. Panel B presents estimates for any tobacco purchase, counts of cigarettes and cigars, and ounces of chewing and smoking tobacco purchased per adult. Across tobacco products, the estimates are consistently negative but are statistically significant only for smoking tobacco. Smoking tobacco purchases declined by approximately 0.6%, corresponding to a reduction of about 0.0013 ounces per adult relative to the pre-treatment mean. Appendix Figures A7 and A8 show the SDiD pre-treatment fit event-study estimates for these models.

6.3 Heterogeneous behavioral responses

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 these dimensions.

6.3.1 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 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.[33] 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 5 reports the estimates for (1) low-, (2) moderate-, and (3) high-intensity purchasing households.[34]

[Table 5 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 A10 and Appendix Figure A11 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 A5 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.

[33] Appendix Figure A9 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.

[34] Appendix Table A4 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.

6.3.2 Demographic heterogeneity

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

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.[35]

[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 show positive point estimates. However, these estimates should be interpreted cautiously. When the sample is divided into narrower subgroups, the number of treated Maryland households becomes substantially smaller. Accordingly, the figure is best viewed as descriptive evidence on the direction and relative magnitude of subgroup responses rather than a causal estimate. Average treatment effects on the treated by demographic category are reported in Appendix Tables A6 to A9. Appendix Figures A12 through A19 report the SDiD fit and event study estimates for each respective demographic group.

6.4 Robustness checks

I also assess the robustness of the main SDiD estimate using alternative sample periods, inference procedures, empirical strategies, donor-pool definitions, treatment timing assumptions, macroeconomic adjustments, and sample and treatment definitions.

[35] 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.

First, I look at whether the results are sensitive to the balanced-panel restriction and the selected sample period. Table 6 presents two-way fixed-effects estimates for the three alternative sample windows described in Section 5.3. Columns (1) and (2) use the full pre-COVID-19 sample from 2004 to 2019, columns (3) and (4) use the 10-year sample from 2006 to 2016, and columns (5) and (6) use the four-year window that I use in the main specification but retaining all available households rather than imposing a panel restriction. Within each pair, the first column includes state fixed effects and the second includes household fixed effects.

Across the alternative sample windows, the estimates range from declines of 4.8% to 6.8%. All estimates are negative, suggesting that the main finding is not driven by the balanced-panel restriction or the selected four-year sample period. These estimates should nevertheless be interpreted as directional robustness checks because they rely on different identifying assumptions from my preferred SDiD design.

[Table 6 here]

Furthermore, Figure 5 summarizes additional robustness checks. Each point represents an estimate of the effect of the Maryland alcohol sales tax increase on the inverse hyperbolic sine of monthly ethanol ounces purchased per adult, and the horizontal capped lines represent 95% confidence intervals. The baseline estimate (with bootstrapped standard errors), is shown with a blue circle and reproduces my preferred SDiD specification for reference.

[Figure 5 here]

Looking at Figure 5, I first 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.[36]

I next compare the preferred SDiD estimate with estimates obtained using the synthetic control method and propensity-score-matched difference-in-differences. The estimated effects

[36] I further discuss these alternative SDiD standard error procedures in Appendix B, and Appendix figures B1 and B2 present the event studies for the Jacknife and placebo procedures.

remain negative across both alternative empirical strategies. The synthetic control estimate is similar to the preferred SDiD estimate, while the propensity-score-matched difference-in-differences specification produces a larger estimated reduction in alcohol purchases.[37]

Finally, the effects remain negative when I exclude Maryland's neighboring states from the donor pool and when I restrict the donor pool to these states only. Although the magnitudes vary across the two specifications, the results do not appear to depend on the inclusion of neighboring states in the donor pool or on comparisons with distant states.[38]

6.4.1 Border counties

The state-level checks presented in Figure 5 do not isolate the neighboring households most likely to have purchased alcohol in Maryland. Maryland borders four states and the District of Columbia (see Appendix Figure A21). Before the tax increase, households in neighboring counties outside Maryland may have purchased alcohol in the state. If the tax changed their total alcohol purchases, including them as untreated households in the donor pool could affect the estimated counterfactual for Maryland.

[Figure 6 here]

Figure 6 examines this possibility using two alternative treatment definitions. First, I add households in neighboring-state counties that border Maryland to the treated group, allowing the estimated effect to include households potentially exposed across the state line (Panel 6a). Second, I retain Maryland as the treated group but exclude those neighboring-state border households from the donor pool (Panel 6b). Both specifications show post-tax declines in ethanol purchases per adult comparable to the main estimate.[39]

Since Montgomery county in Maryland has a different alcohol distribution system, I also further look alcohol purchasing trends particularly in this county to see if they differ from the rest of the state. This is presented in Appendix D.

[37] Appendix C presents a detailed explanation on how the synthetic control method and the propensity score matched control group specifications were estimated.

[38] I also look at heterogeneity by season, and the estimates remained significant (see Figure A20).

[39] Appendix Table A11 reports the estimated treatment effects.

6.5 Alcohol spending

Although the main analysis looks at changes in the quantity of alcohol purchased, I also briefly study changes in alcohol spending.[40] Table 7 shows the SDiD estimates for alcohol expenditures and alcohol expenditures as a share of total expenditures for the full sample (Panel A) and the drinking sample (Panel B). Across both panels, the estimates for the two outcomes are negative and statistically significant.

[Table 7 here]

For the full sample, the estimates suggest a reduction of $0.28 less in monthly alcohol expenditure per adult for the full sample, and $1.52 less per adult per household in the drinking sample. Recall that the expenditure measures do not include the 3% point increase that Maryland faced from July 2011 and on. Because NielsenIQ expenditures exclude sales taxes, these results capture reductions in the pre-tax value of alcohol purchases and do not directly measure changes in total payments including tax. Alcohol's share of total recorded expenditures decreased by 0.127 percentage points for the full sample and 0.285 percentage points for the drinking sample, equivalent to 7.1% and 8.6% of their respective pre-tax means. Appendix Figure A23 and Appendix Figure A24 present the SDiD fit and event-study estimates for overall alcohol spending and alcohol expenditure as a share of total expenses for the full sample and the drinking sample.

7 Discussion

This paper provides new household-level evidence on how alcohol purchases respond to an alcohol specific sales tax increase. Although sales taxes are added at checkout and may be less salient than excise taxes which are embedded in posted prices (Chetty et al., 2009), the evidence I present suggests that Maryland's 2011 alcohol sales tax increase led to a meaningful reduction in household alcohol purchases. This finding contributes to the broader literature on alcohol taxation by showing that purchasing behavior could also respond to price-based sales taxes, and not solely to excise taxes.

Using a synthetic difference-in-differences framework and data from the NielsenIQ Consumer Panel, I estimate a 5.5% decline in monthly ethanol ounces purchased per adult in the

[40] Appendix Figure A22 plots the raw trends in the overall alcohol expenditures as well as share of expenditures on alcohol for both the full and drinking sample.

full sample and a 9.2% decline among households that purchased alcohol in the 12 months preceding the tax increase. These estimates correspond to approximately 0.33 and 1.04 fewer ounces of ethanol per adult per month, or approximately 0.6 and 1.7 fewer standard drinks. The larger estimate among households with pre-tax alcohol purchases indicates that the full-sample effect is attenuated by the inclusion of households with little or no baseline alcohol purchasing.

The smaller estimate for the full sample should not be interpreted as evidence that households with no pre-tax alcohol purchases have a less-elastic demand. Rather, the two specifications estimate average effects for different populations. The full-sample estimate includes households with no recorded alcohol purchases in the 12 months before the tax increase and therefore little observed purchasing behavior over which to adjust. The larger estimate among households with pre-tax alcohol purchases is consistent with the response being concentrated among households with more consistent participation in the alcohol market prior to the tax increase.

Looking at the event studies, I find a sustained decrease in alcohol purchases across the post periods. Because the analysis follows the same households over time, consumers repeatedly face the post-tax purchasing environment and have opportunities to adjust their behavior across shopping trips. The persistence of the estimated reduction suggests that the effect of the alcohol-specific sales tax on purchases was not limited to the time of its implementation. Importantly, the results suggest that a sales tax, despite concerns about lower salience relative to excise taxes (Chetty et al., 2009), can still generate behavioral responses.

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. Because their drinking-level samples are defined using positive alcohol purchases before the tax increase, the estimate for my alcohol-purchasing sample provides the more comparable benchmark. For this sample, I estimate a 9.2% reduction, corresponding to 1.04 fewer ounces of ethanol per adult per month. However, my full-sample estimate of a 5.5% reduction is closer in magnitude to their estimate. Despite differences in tax structure, baseline prices, sample construction, and empirical setting, the estimates suggest reductions in alcohol purchases sometimes even larger than excise taxes.

The beverage-specific results provide additional evidence on how households adjusted their purchases. Because a sales tax is proportional to price, it raises the cost of alcoholic beverages while leaving relative prices across beverage types largely unchanged.[41] Consistent with this structure, I find reductions in purchases of beer, wine, and spirits, with no evidence that households shifted purchases across beverage types. This differs from some evidence from excise tax settings, where consumers may substitute toward relatively less expensive products (Gehrsitz et al., 2021).

I also find relatively small increases in purchases of non-alcoholic beer and wine, and no economically meaningful changes for tobacco-related products. These results provide little evidence of movement toward non-alcoholic alternatives, although the magnitude of the response is small. Broadly, there is little evidence that the reduction in alcohol purchases was accompanied by changes in purchases of other related products.

In addition, 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.

The main results remain robust to alternative sample periods, inference procedures, empirical strategies, and donor pools. Specifically for the border analysis, the reduction in ethanol ounces purchased per adult per household is similar when neighboring-state border households are excluded from the donor pool and when they are added to the treated group,

[41] By relative prices, I refer to the ratio of prices across products (e.g., beer relative to wine or spirits). Under a proportional sales tax, these ratios remain unchanged if pre-tax prices do not differentially adjust across products.

with the latter specification estimating an effect for a broader population.

Finally, to provide additional context for the purchase results, I examine alcohol spending. In both the full sample and the drinking sample, the estimates suggest reductions in monthly alcohol expenditure per adult and in alcohol's share of total expenditure.[42] These findings suggest that households reduced their spending on alcohol relative to what they would have spent in the absence of the tax increase and allocated a smaller share of their spending to alcohol. Because NielsenIQ expenditures exclude sales taxes, these results capture changes in pretax spending. Accounting for the higher sales tax would offset part of the reduction in pretax expenditure, so the estimates should not be interpreted as equivalent reductions in total out-of-pocket spending.

The reductions in spending are consistent with the reductions in ethanol purchases I observe. However, it is also possible that changes in the prices and composition of purchased products may also contribute to the observed reduction. Economic theory does not imply that households reduce purchases only enough to maintain their previous alcohol spending. Depending on their responsiveness to higher prices, households may reduce alcohol purchases sufficiently for spending to decline even after accounting for the additional tax. They may also shift toward less expensive products or purchase alcohol less frequently. Importantly, the expenditure results are not able to distinguish between these mechanisms.

This study has many limitations. Although I use SDiD to construct a comparison group, I still only have one treated state. Having a much larger number of comparison households does not address this limitation, and inference remains challenging. In addition, the NielsenIQ Consumer Panel records the prices of products that households purchase rather than store-level prices for a consistent set of products. Changes in these observed prices may therefore reflect changes in what households buy or where they shop, which limits my ability to estimate price elasticities and examine how retailers adjusted prices after the tax increase. Looking at expenditures provides additional context, but it does not resolve this limitation, and these expenditures do not include sales taxes. Finally, I observe household alcohol purchases rather than individual consumption, and this data does not capture purchases at bars and restaurants. Therefore, the reductions in recorded purchases may not translate into equivalent reductions in overall alcohol consumption.

Overall, these findings suggest that alcohol sales taxes can meaningfully reduce alcohol purchases and may serve as a useful complement to traditional excise taxes. By increasing

[42] Total expenditure includes both alcohol and non-alcohol expenses.

the final retail price across a wide range of products, sales taxes can generate reductions in demand without inducing substantial substitution across beverage types. These results have important implications for alcohol tax policy. Policymakers should consider both excise and sales taxes as part of a comprehensive strategy to reduce excessive alcohol consumption and its associated social harms.

8 Conclusion

The 2011 alcohol sales tax increase in Maryland appears to have been an effective tool for reducing household alcohol purchases. I estimate a 5.5% decline in monthly ethanol ounces purchased per adult in the full sample and a 9.2% decline among households that purchased alcohol in the 12 months preceding the tax increase, equivalent to approximately 0.6 and 1.7 fewer standard alcoholic drinks per month. I observe these reductions across beer, wine, and spirits, with no evidence of substitution towards a specific type of alcohol. I also find small increases in purchases of non-alcoholic beer and wine and no economically meaningful changes in tobacco-related products.

In addition, while households with low, moderate, and high baseline purchasing levels all reduce alcohol purchases following the tax increase, the largest absolute declines happen among pre-tax high-purchasing households. I also conduct subgroup analyses that estimate broadly similar reductions in alcohol purchases across demographic groups defined by income, education, household composition, and race or ethnicity, although these estimates are less precise. I also find consistent reductions in alcohol spending, likely driven by people purchasing less alcohol.

Taken together, the findings indicate that an alcohol sales tax can meaningfully reduce alcohol purchases. Alcohol sales taxes may be a useful complement to traditional excise taxes. Because sales taxes are proportional to retail prices, they adjust as prices rise and do not require repeated legislative action to maintain their value, unlike excise taxes. Reductions in alcohol purchases are a first-order effect of the policy and may lead to broader societal benefits, including reductions in alcohol-impaired driving and other alcohol-related 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.

Acuff, Samuel F, Justin C Strickland, and Derek D Reed, "The Effect of Price on the Behavioural Economic Substitutability of Non-Alcoholic and Alcoholic Beverages," *Drug and Alcohol Review*, 2026, *45* (2), e70121.

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.

Goel, Rajeev K, James E Payne, and James W Saunoris, "Spillovers from the beer market to US cigarette demand," *Economic Modelling*, 2016, *55*, 292–297.

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.

_ , Luke B Wilson, Duncan Gillespie, Colin Angus, Damon Morris, and Alan Brennan, "Estimation of integrated price elasticities for alcohol and tobacco in the United Kingdom using the living costs and food survey 2006–2017," *Drug and Alcohol Review*, 2024, *43* (1), 315–324.

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.

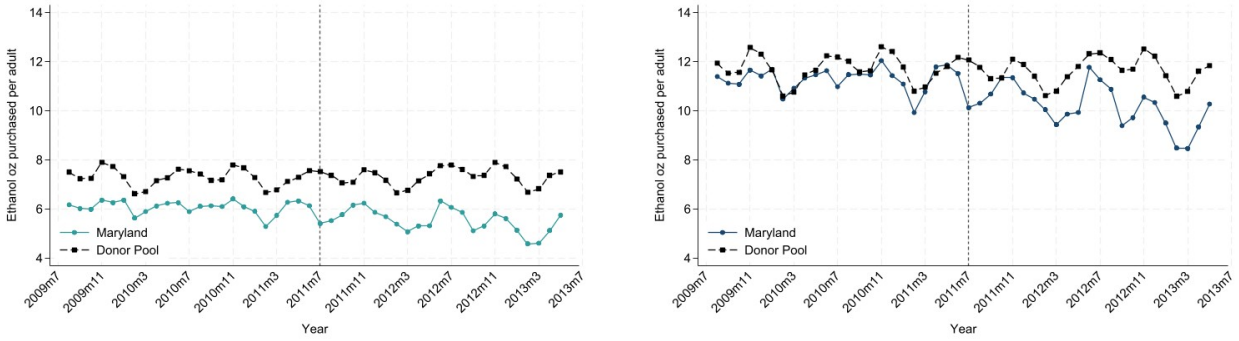
Staras, Stephanie AS, Melvin D Livingston, and Alexander C Wagenaar, "Maryland alcohol sales tax and sexually transmitted infections: A natural experiment," *American journal of Preventive Medicine*, 2016, *50* (3), e73–e80.

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

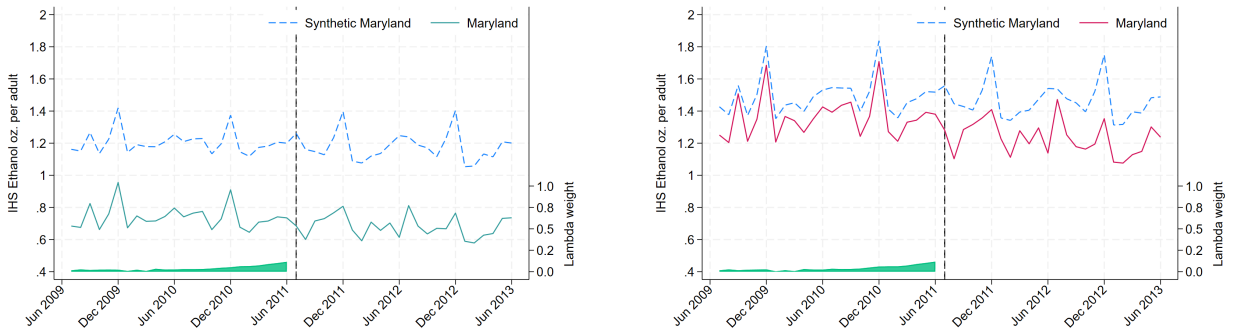


(a) Full sample

(b) Drinking sample

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

Note. Panel (a) presents the full sample, and Panel (b) restricts the sample to households that purchased alcohol at least once during the 12 months before the tax increase. Each panel plots the three-month moving average of mean ethanol ounces purchased per adult in Maryland (solid line) and the donor pool (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.



(a) Full sample

(b) Drinking sample

Figure 2. Observed and synthetic Maryland purchase trends

Note. Panel (a) presents the full sample, and Panel (b) restricts the sample to households that purchased alcohol at least once during the 12 months before the tax increase. Each panel plots the 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.

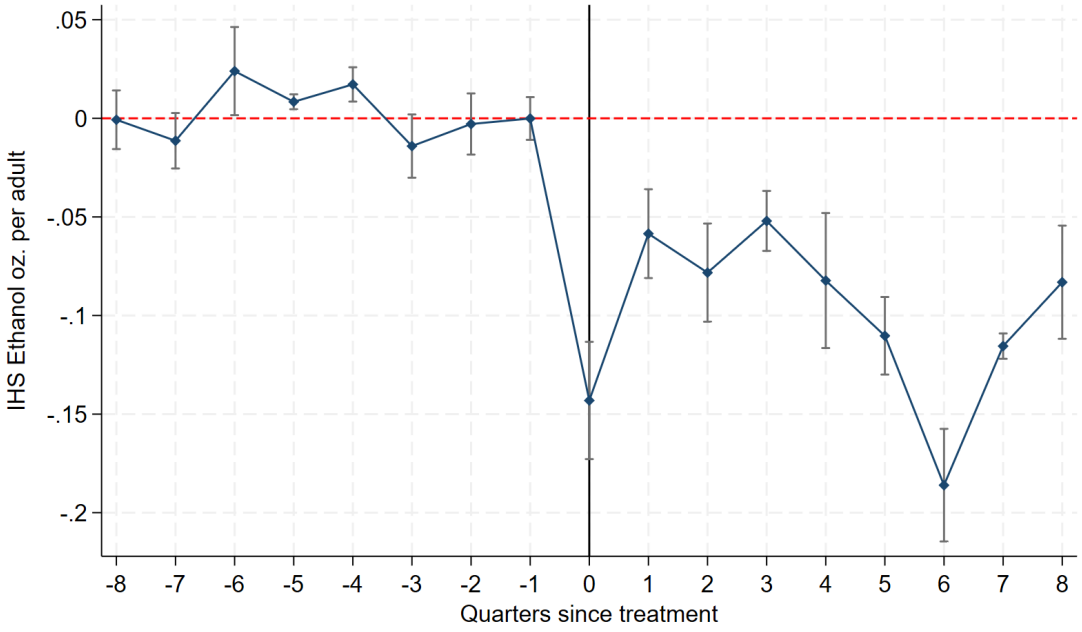


Figure 3. Dynamic Event study estimates for ethanol ounces purchased

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. Sample only includes households who have purchased alcohol at least once during the 12 months prior to the tax increase. Data on alcohol purchases come from the NielsenIQ Consumer panel for the period from July 2009 to June 2013.

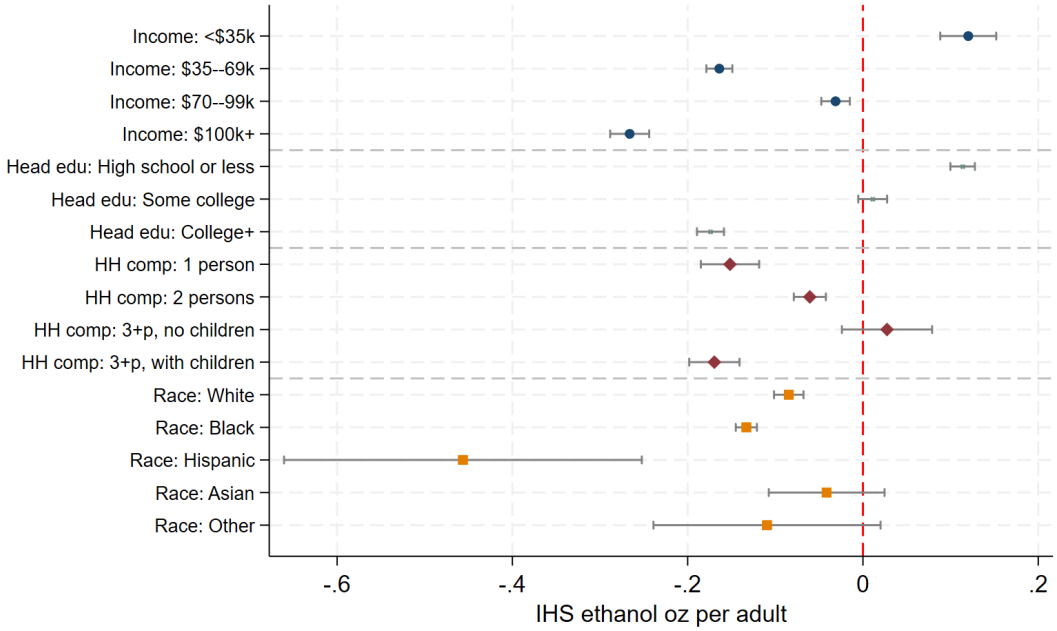


Figure 4. Effect of Maryland's alcohol sales tax on ethanol ounces purchased across demographic groups

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 July 2009 to June 2013.

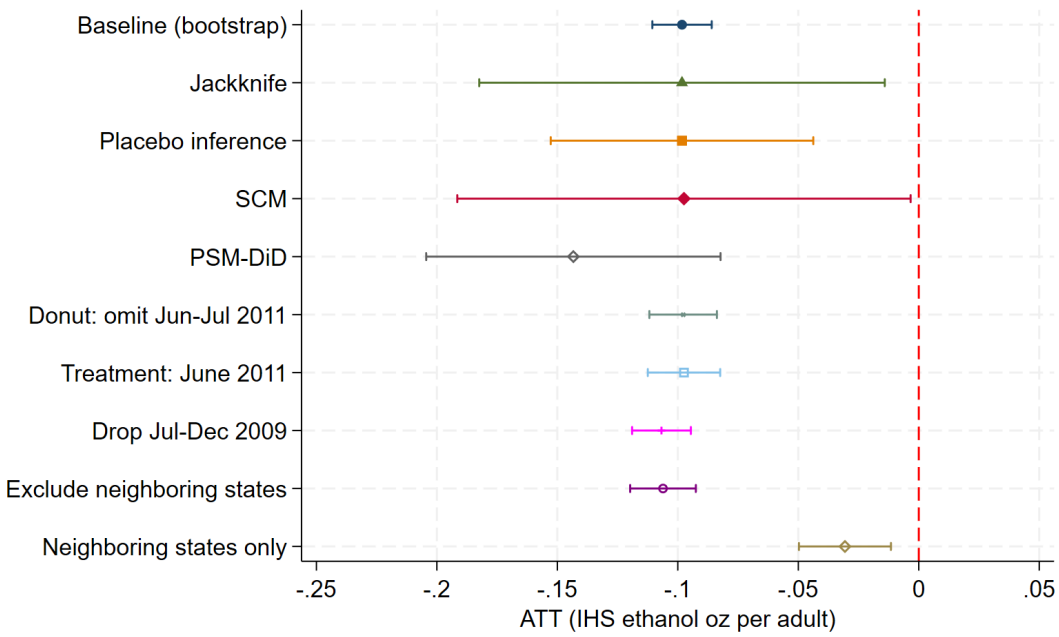
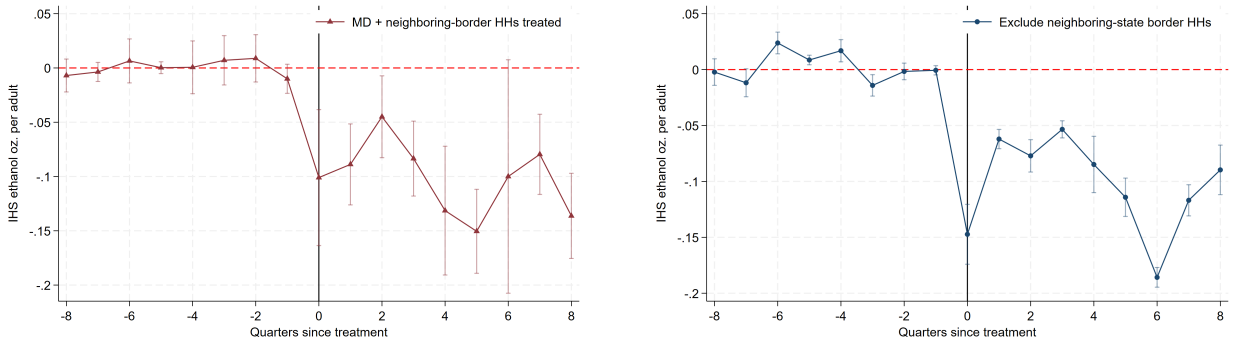


Figure 5. Robustness of estimates to alternative specifications

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 omits June and July 2011. The empty light blue square denotes the SDiD estimate when treatment is defined as June 2011, rather than July. The light pink marker shows the estimate where July through December 2009 are dropped due to potential exposure to the Great Recession. The purple hollow circle plots the SDiD estimate excluding neighboring states, while the hollow olive diamond plots the SDiD estimate using only nearby donor states (Delaware, Pennsylvania, Virginia, and West Virginia). Gray capped lines represent 95% confidence intervals. All models include 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.



(a) Neighboring border households as treated

(b) Neighboring border households excluded

Figure 6. Event study estimates for ethanol ounces purchased for alternative treatment definitions for Maryland border counties

Note.The figure shows synthetic difference-in-differences (SDiD) event-study estimates for two models. Panel 6a presents the estimates when the treatment definition is expanded to include neighboring border households from other states and panel 6b plots the event study estimates when neighboring border household from other states are dropped. The dependent variable is the inverse hyperbolic sine of monthly ethanol ounces purchased per adult. 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. Circles indicate point estimates, and capped lines represent 95% confidence intervals. Standard errors are clustered at the county level. The model includes controls for income, household size, marital status, presence of children, race, education, and age. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.

Tables

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

	Full sample		Drinking sample	
	Donor Pool	Maryland	Donor Pool	Maryland
Alcohol purchases				
Ethanol oz per adult	7.61	6.47	12.35	11.96
Beer	2.76	1.68	4.48	3.12
Wine	1.79	2.15	2.90	3.98
Spirits	3.07	2.64	4.97	4.85
Purchaser intensity				
Low intensity	–	–	0.34	0.34
Moderate intensity	–	–	0.34	0.26
High intensity	–	–	0.32	0.39
Household income				
Less than $35,000	0.34	0.27	0.28	0.20
$35,000–$49,999	0.16	0.10	0.15	0.10
$50,000–$69,999	0.16	0.14	0.17	0.14
$70,000–$99,999	0.16	0.21	0.17	0.26
$100,000 or more	0.20	0.28	0.24	0.30
Age of household head				
25–44 years	0.18	0.15	0.19	0.17
45–64 years	0.55	0.61	0.56	0.59
65 years or older	0.27	0.24	0.25	0.25
Education of household head				
Less than high school	0.03	0.03	0.02	0.00
High school graduate	0.32	0.26	0.28	0.22
Some college	0.30	0.27	0.31	0.28
College graduate or more	0.35	0.44	0.38	0.50
Household composition				
1 person	0.31	0.31	0.28	0.31
2 persons	0.37	0.39	0.39	0.40
3+ persons, no children	0.10	0.08	0.11	0.08
3+ persons, with children	0.22	0.21	0.22	0.22
Married	0.44	0.44	0.46	0.45
With children	0.24	0.24	0.24	0.26
Race and ethnicity				
White	0.81	0.63	0.81	0.58
Asian	0.03	0.03	0.03	0.05
Black	0.10	0.29	0.10	0.32

Continued on next page

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

	Full sample		Drinking sample	
	Donor Pool	Maryland	Donor Pool	Maryland
Hispanic	0.10	0.03	0.11	0.03
Other race	0.06	0.05	0.07	0.06
Observations	404,880	9,240	250,224	4,920
Unique households	16,870	385	10,426	205

Notes. This table presents weighted pre-treatment means for households in Maryland and the donor pool. Each observation is at the household-month level. The full sample includes all households that satisfy the sample restrictions, while the drinking sample is restricted to households with at least one alcohol purchase during the 12 months before the tax increase. Comparison states exclude Connecticut, Illinois, New Jersey, New York, North Carolina, Washington, the District of Columbia, and Kansas because these jurisdictions experienced alcohol tax changes during the sample period. Values for demographic and household characteristics are expressed as proportions. Purchaser-intensity categories are defined only for households in the drinking sample and exclude households classified as stockpilers. The price measure retains observed zeros for household-months without alcohol purchases. 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. Observation and household counts are unweighted.

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

	Full sample		Drinking sample	
	(1) Baseline	(2) Covariates	(3) Baseline	(4) Covariates
Treat x Post	-0.032***	-0.056***	-0.098***	-0.096***
	(0.004)	(0.004)	(0.006)	(0.006)
Percent effect	-3.1	-5.5	-9.4	-9.2
MD Pre-treatment mean	6.06	6.06	11.29	11.29
Observations	828,240	828,240	510,288	510,288
Household FE	YES	YES	YES	YES
Year-month FE	YES	YES	YES	YES
Household controls	NO	YES	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. 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. Effects of the Alcohol Sales Tax on Related Product Purchases.

Outcome	Estimate	SE	Percent effect	MD pre-treatment mean
Panel A: Non-alcoholic beer and wine				
Overall	0.019***	(0.005)	2.0	0.72
Beer	0.022***	(0.005)	2.2	0.37
Wine	0.007**	(0.003)	0.7	0.35
Panel B: Tobacco				
Any purchase	-0.002	(0.001)	-1.4	0.124
Cigarettes	-0.006	(0.005)	-0.6	38.638
Cigars	-0.001	(0.004)	-0.1	0.871
Chewing tobacco	-0.001	(0.001)	-0.1	0.001
Smoking tobacco	-0.006***	(0.002)	-0.6	0.022
Observations	510,288			

Note. The table reports synthetic difference-in-differences (SDiD) estimates for non-alcoholic beer and wine (Panel A), and for tobacco (Panel B) (vertically). All outcome variables were transformed using the the inverse hyperbolic sine (IHS). Panel A presents estimates for overall non-alcoholic ounces, non-alcoholic beer, and non-alcoholic wine per adult per household. Panel B presents the estimates for any tobacco purchase (binary-non-transformed), counts of cigarettes, cigars, and e-cigarettes, and ounces of chewing tobacco, and smoking tobacco per adult per household. All specifications include household fixed effects, month-year fixed effects, and 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 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 5. 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 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.

Table 7. Effects on alcohol expenditure and share spent on alcohol per adult

	(1) Alcohol expenditure per adult	(2) Alcohol share of recorded expenditures (%)
Panel A: Full sample		
Treat × Post	-0.035***	-0.127***
	(0.004)	(0.016)
Percent effect	-3.5	-0.127
MD pre-tax mean	8.04	1.79
Observations	828,240	828,240
Panel B: Drinkers		
Treat × Post	-0.107***	-0.285***
	(0.008)	(0.023)
Percent effect	-10.2	-0.285
MD pre-tax mean	14.95	3.31
Observations	510,288	510,288
Household FE	YES	YES
Year-month FE	YES	YES

Note. The table presents estimates for two outcome variables: (1) the inverse hyperbolic sine (IHS) of alcohol expenditure in a month, and (2) the share of alcohol expenditures (the denominator includes both alcohol and non-alcohol expenditures). I also present the estimates for two samples, the full sample (Panel A), and the drinking sample wich is restricted to households that made at least one alcohol purchases in the month preceding the tax increase (Panel B). All specifications include household and year-month fixed effects, household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. In column (1), percent effect reports the approximate percentage change implied by the IHS coefficient in 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.

A Appendix

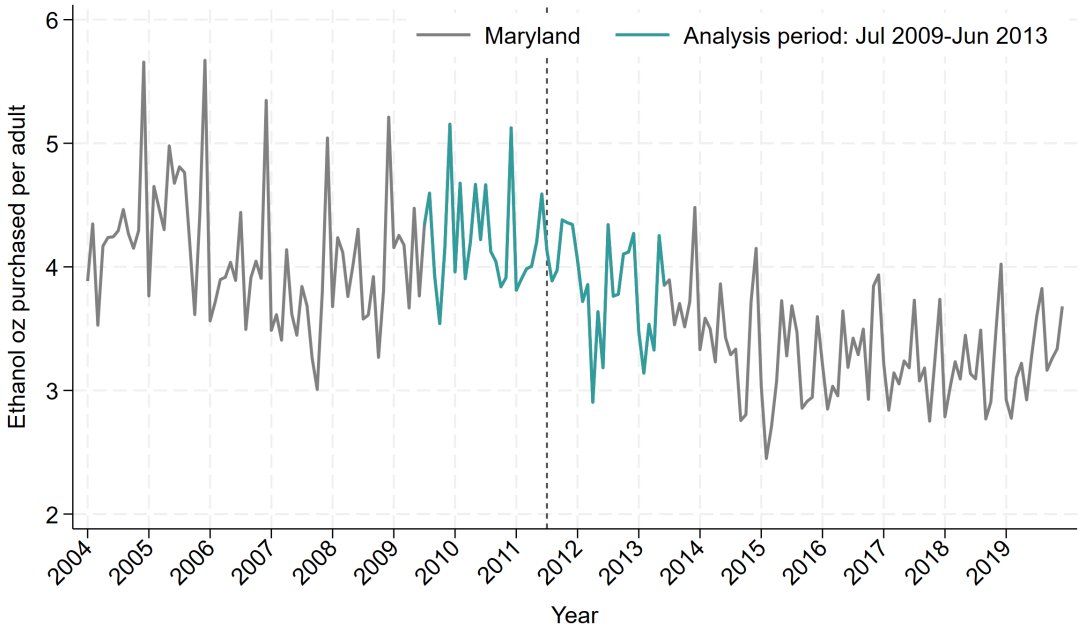


Figure A1. Ethanol ounces purchased per adult for all available time periods in Maryland

Note. The figure plots monthly average ethanol ounces per adult in Maryland households from 2004 to 2019. Monthly ethanol ounces from July 2009 to June 2013 are highlighted in turquoise as that is the sample period used in the preferred specification. Data are from the NielsenIQ Consumer Panel.

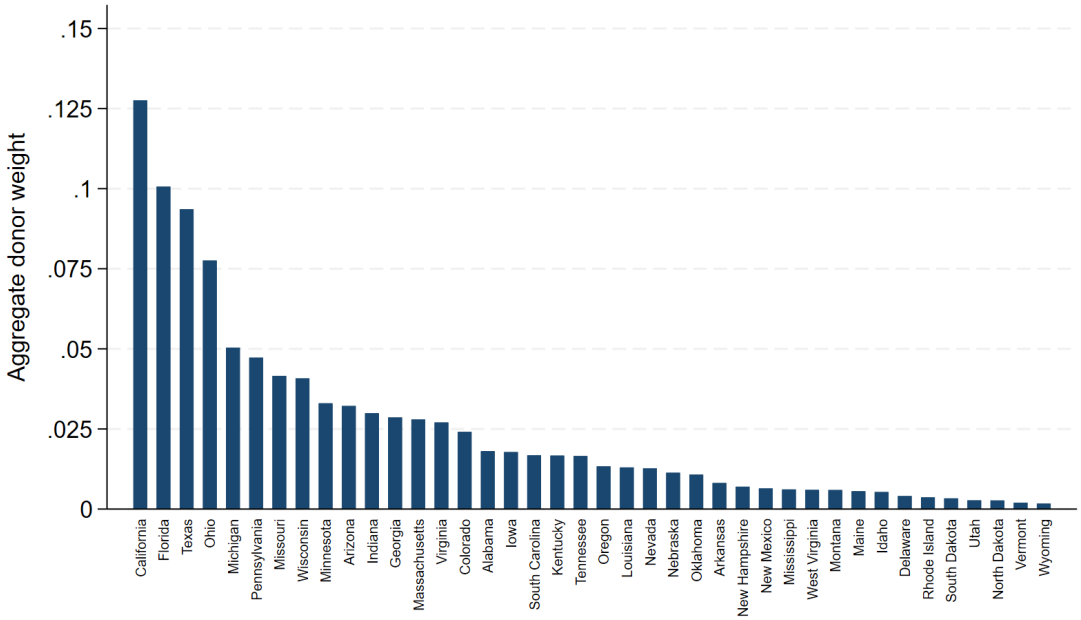


Figure A2. Aggregate SDiD donor weights by state.

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.

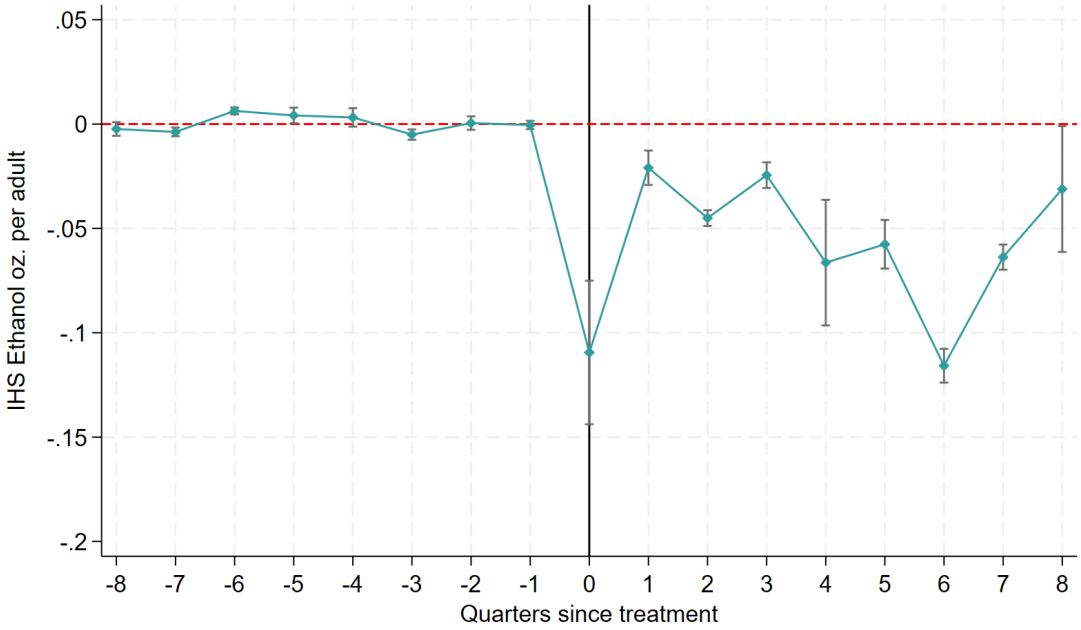


Figure A3. Dynamic Event study estimates for ethanol ounces purchased for the full sample

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 July 2009 to June 2013.

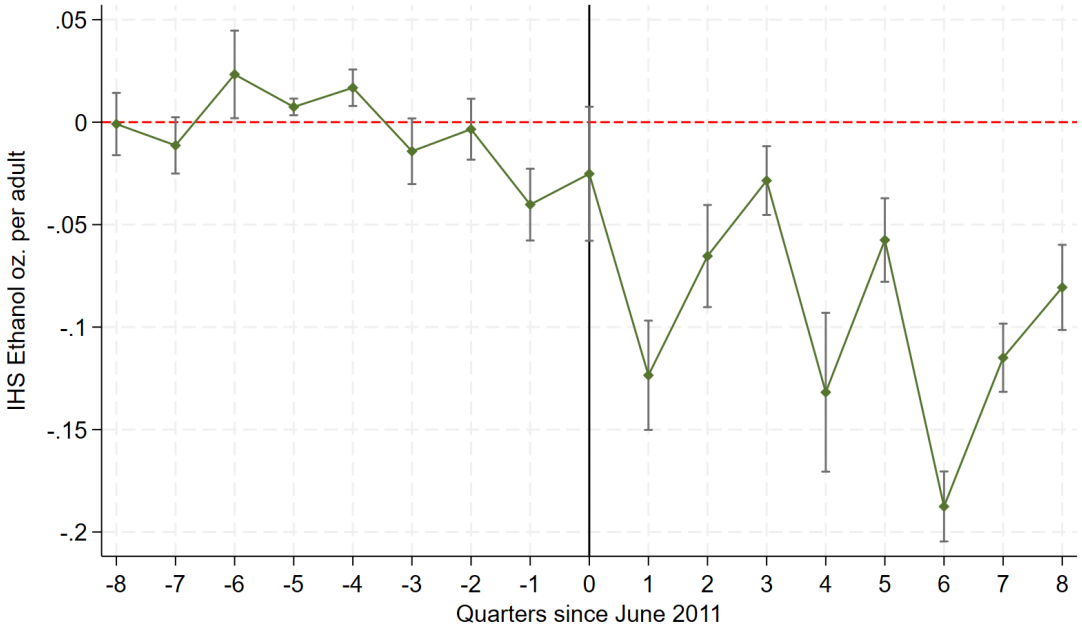
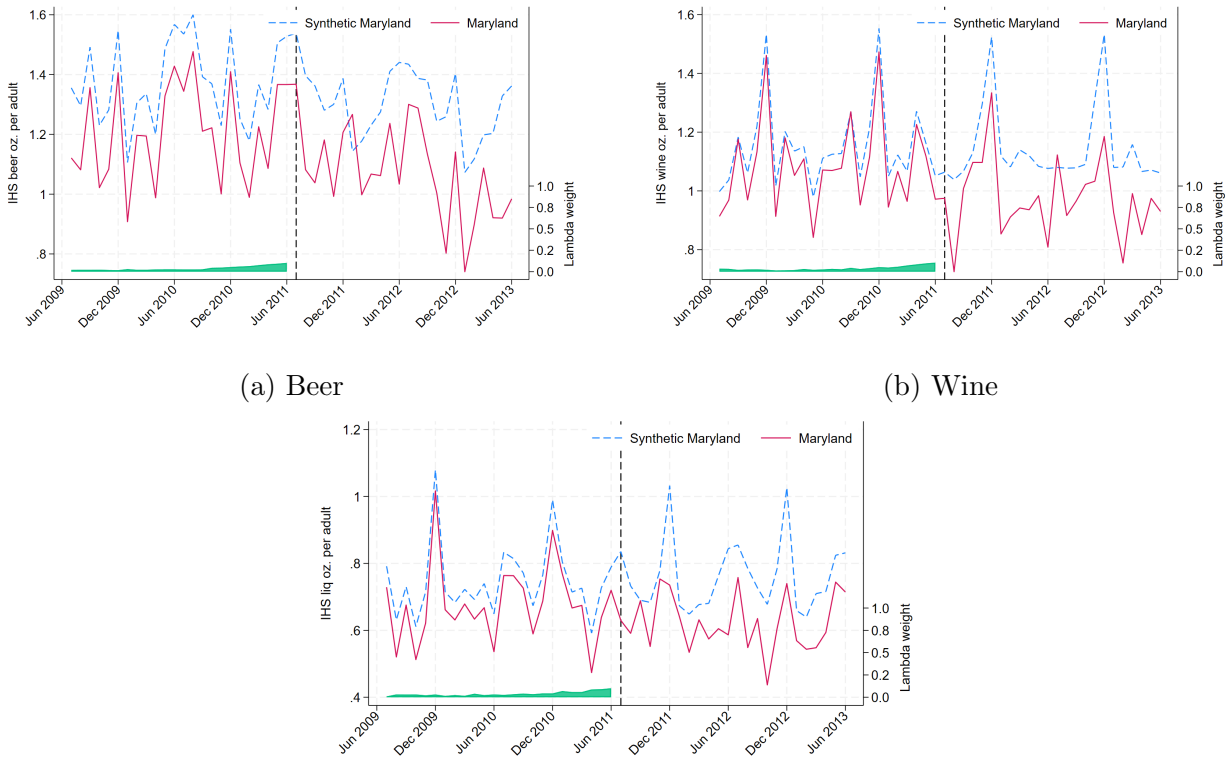


Figure A4. Anticipation event study.

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.



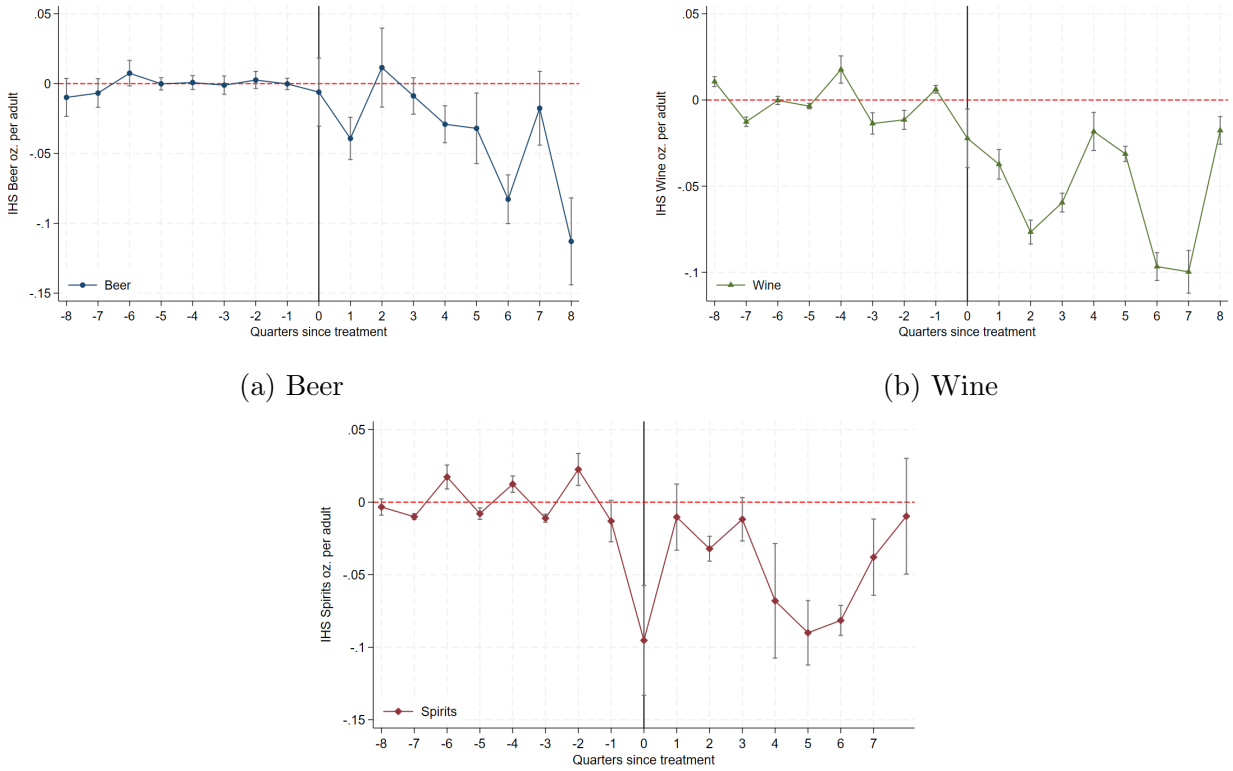
(a) Beer

(b) Wine

(c) Spirits

Figure A5. 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.



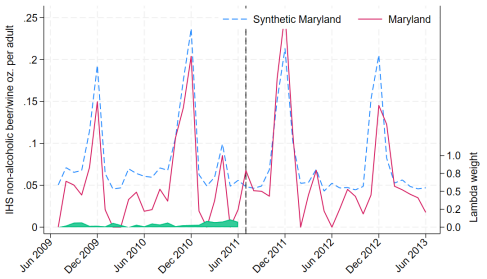
(a) Beer

(b) Wine

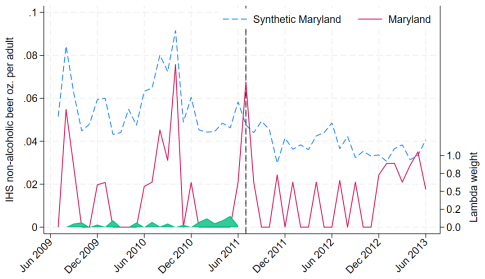
(c) Spirits

Figure A6. Event study estimates by alcohol type: beer, wine, and spirits

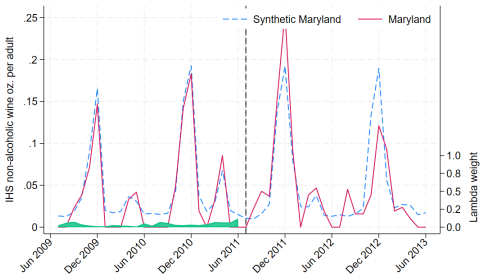
Note. The figure presents event study estimates of the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month by alcohol type, estimated using the Synthetic Difference-in-Differences (SDID) method. Monthly event-time coefficients were averaged to the quarter level for presentation. 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 from July 2009 to June 2013.



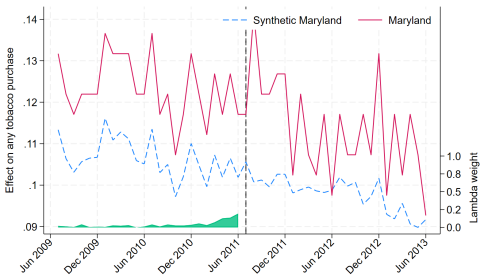
(a) Non-alcoholic beer or wine



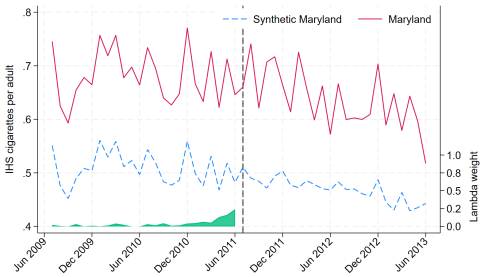
(b) Non-alcoholic beer



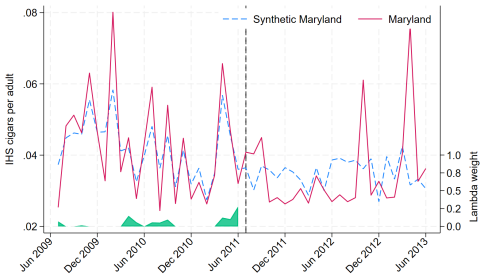
(c) Non-alcoholic wine



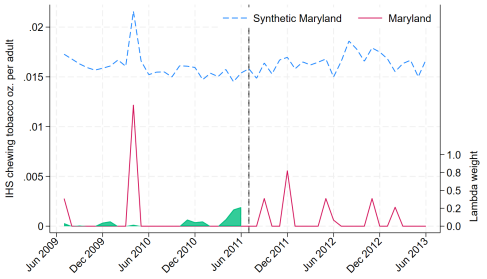
(d) Any tobacco



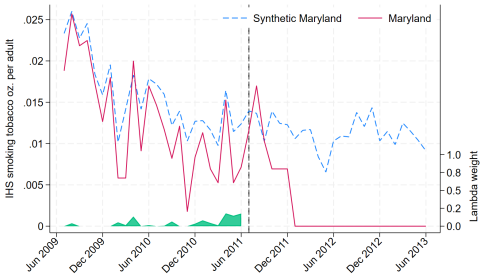
(e) Cigarettes



(f) Cigars



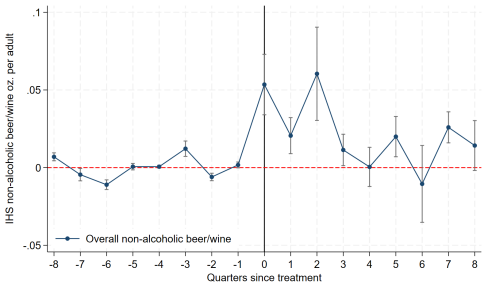
(g) Chewing tobacco



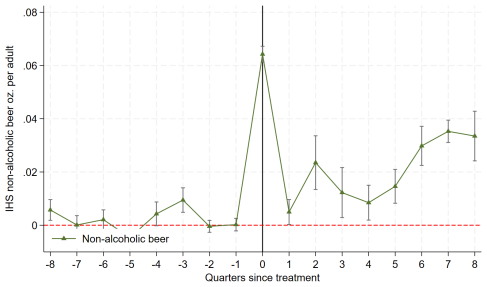
(h) Smoking tobacco

Figure A7. Trends in non-alcoholic beer and wine, and tobacco-related purchases

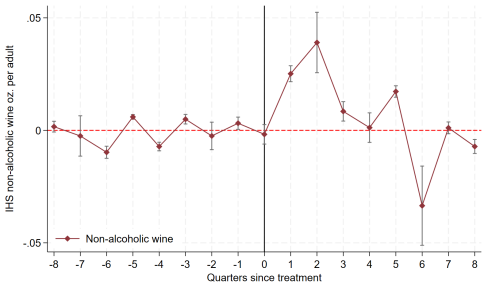
Note. Trends in non-alcoholic (NA) and tobacco-related purchases for Maryland households and the synthetic control constructed using Synthetic Difference-in-Differences (SDiD) weights. Panels (a) −(c) report purchases of NA beer or wine, NA beer, and NA wine, respectively (ounces). Panels (d)−(h) report any tobacco purchase (binary), cigarettes (units), cigars (counts), chewing tobacco (ounces), and smoking tobacco (ounces), respectively. Purchase quantities are measured per adult. All outcomes are transformed using the inverse hyperbolic sine. Time weights (λ_t) minimize the error term in Equation 1. Solid and dashed lines show trends for Maryland and the synthetic control, respectively.



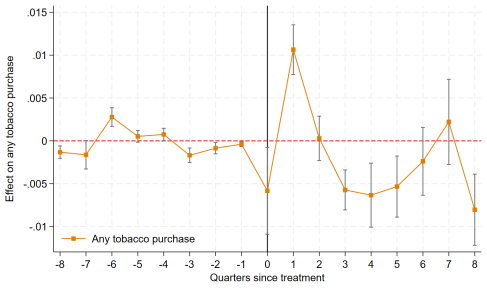
(a) Non-alcoholic beer or wine



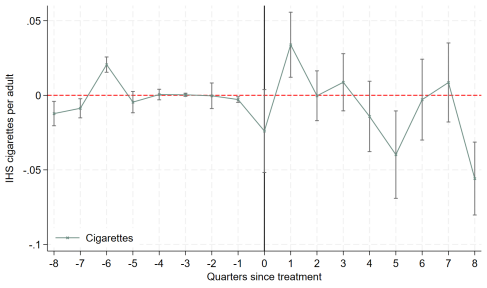
(b) Non-alcoholic beer



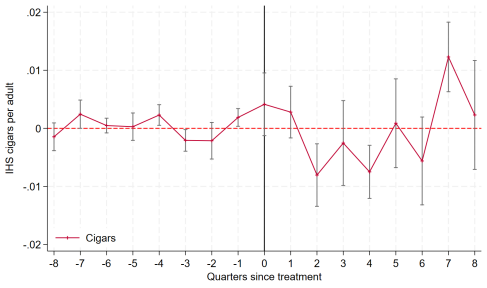
(c) Non-alcoholic wine



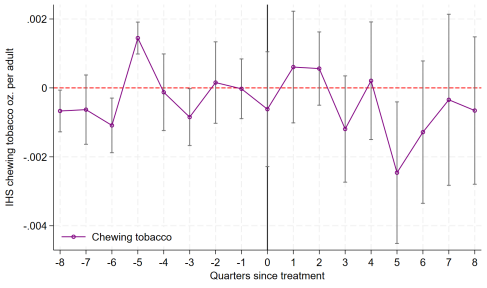
(d) Any tobacco



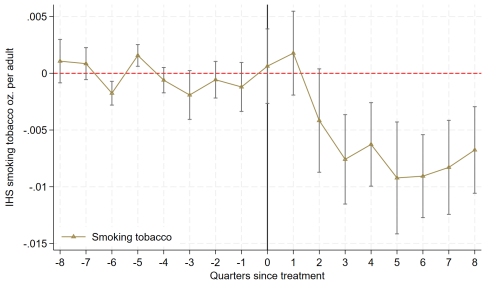
(e) Cigarettes



(f) Cigars



(g) Chewing tobacco



(h) Smoking tobacco

Figure A8. Trends in non-alcoholic beer and wine, and tobacco-related purchases

Note. Event studies for non-alcoholic (NA) and tobacco-related purchases using Synthetic Difference-in-Differences (SDiD). Panels (a)–(c) report NA beer/wine, NA beer, and NA wine purchases (ounces), respectively. Panels (d)–(h) report any tobacco purchase (binary), cigarettes (units), cigars (counts), chewing tobacco (ounces), and smoking tobacco (ounces), respectively. Quantities are measured per adult per household per month. All outcomes are inverse hyperbolic sine transformed. Monthly coefficients are averaged by quarter. The vertical line at zero marks the tax change. Gray capped lines show 95% confidence intervals from 1,000 bootstrap replications. Data comes from NielsenIQ Consumer Panel, July 2009–June 2013.

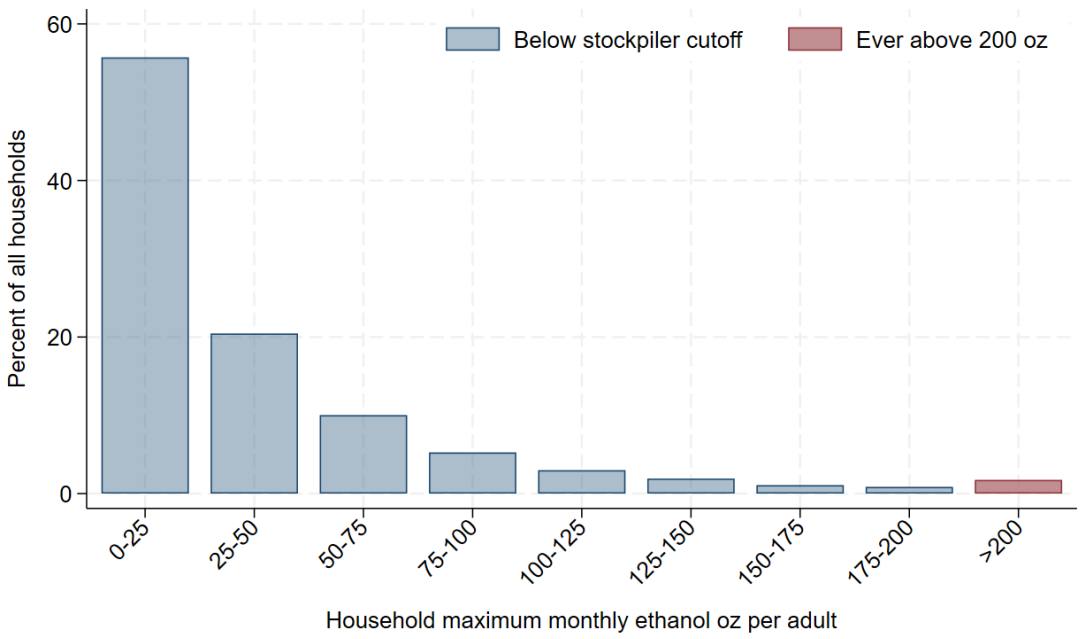
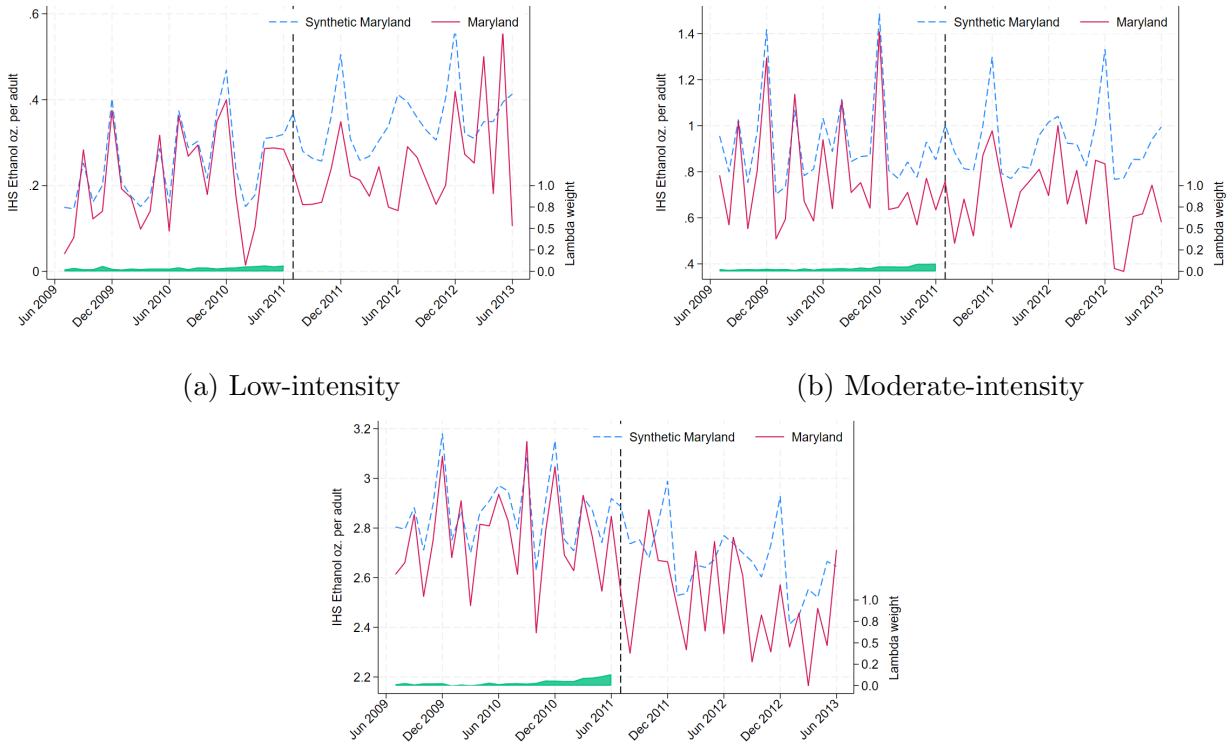


Figure A9. Distribution of maximum monthly ethanol purchases by household.

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.



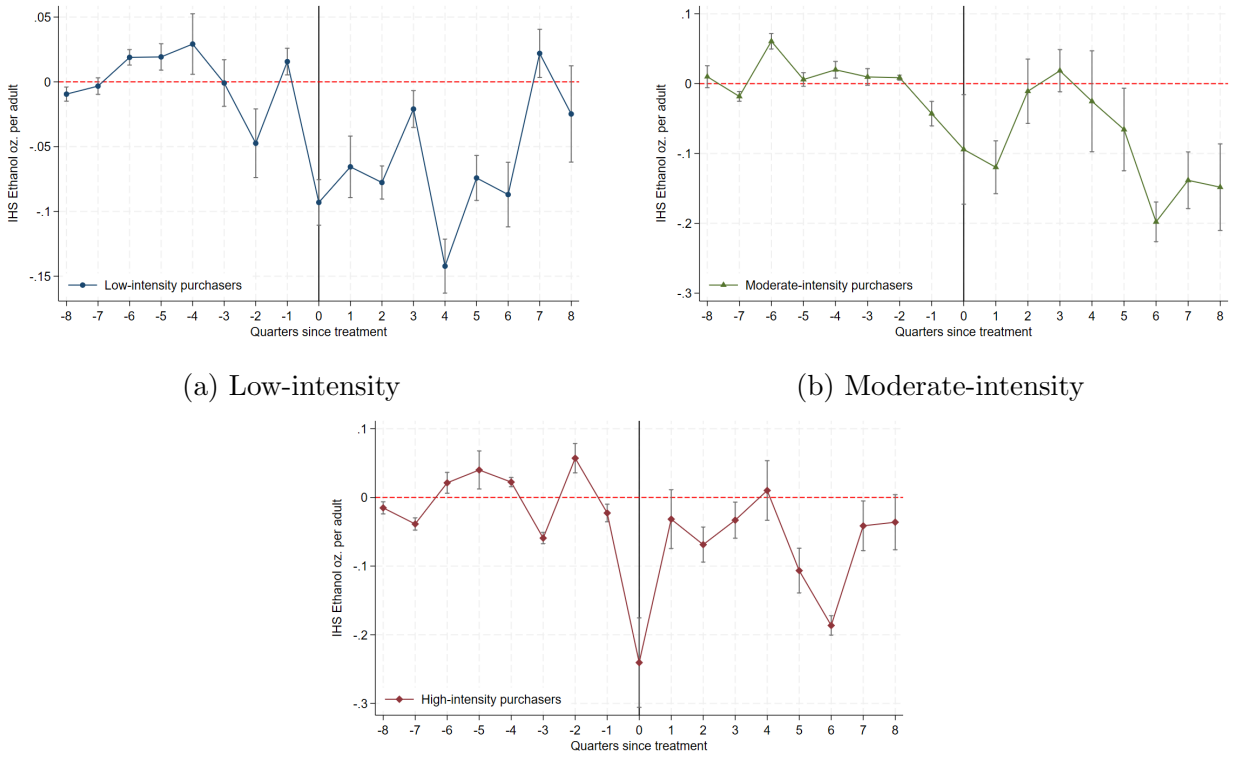
(a) Low-intensity

(b) Moderate-intensity

(c) High-intensity

Figure A10. 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 from July 2009 to June 2013



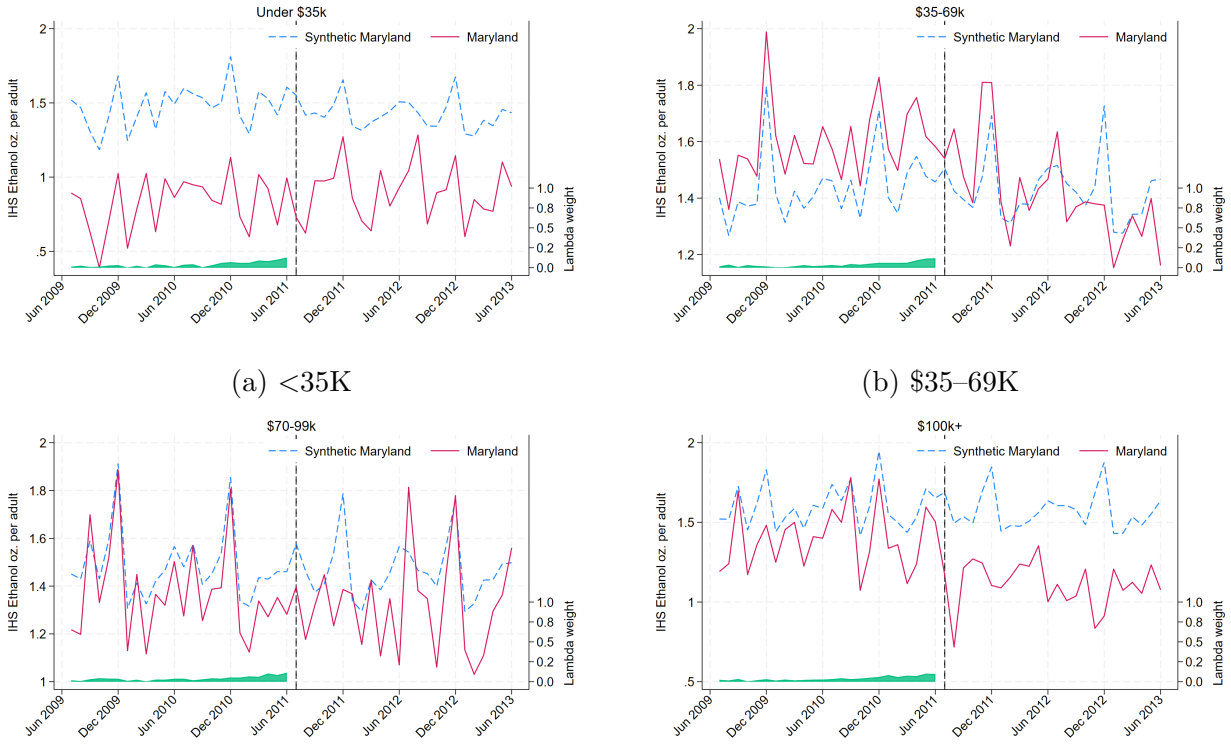
(a) Low-intensity

(b) Moderate-intensity

(c) High-intensity

Figure A11. 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 from period July 2009 to June 2013.



(a) <35K

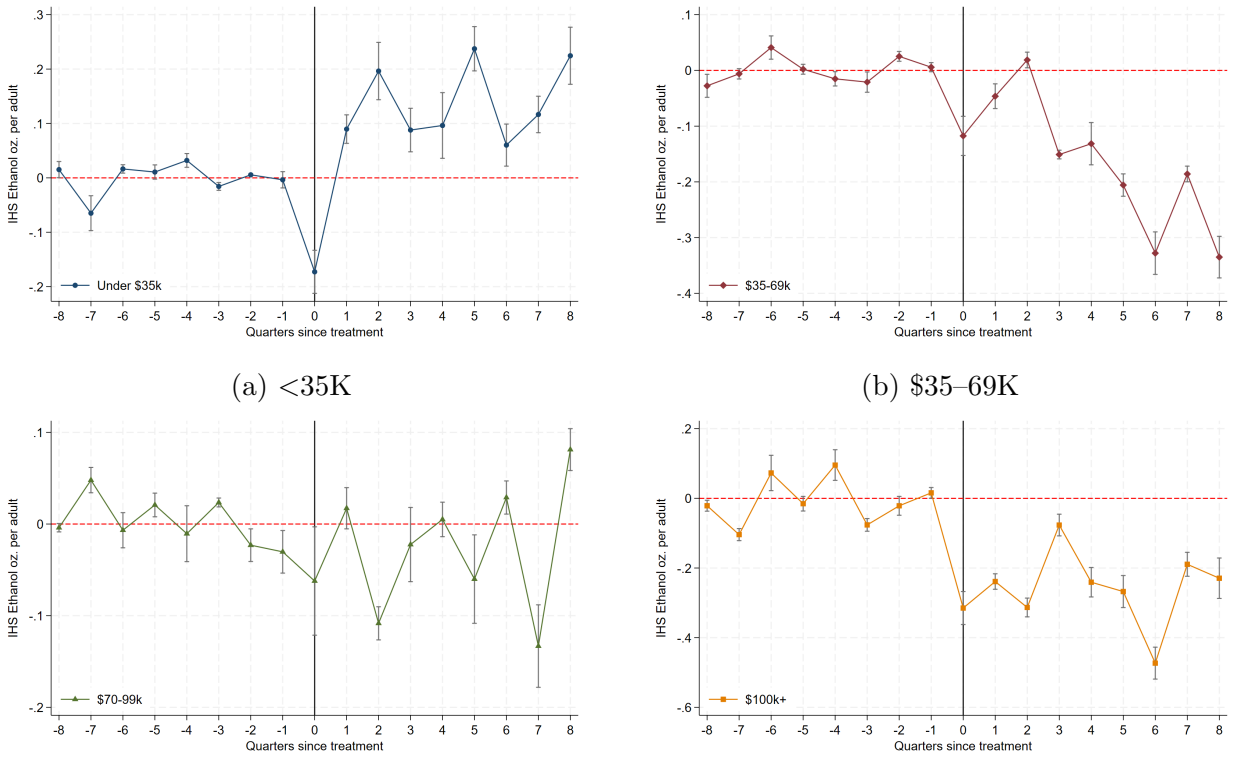
(b) $35–69K

(c) $70–99K

(d) $100K+

Figure A12. SDiD trends in ethanol purchases by household income

Note. Trends in ethanol ounces purchased per adult by household income, using Synthetic Difference-in-Differences (SDiD) weights. Panels (a)–(d) report annual household income groups of less than $35,000, $35,000–$69,999, $70,000–$99,999, and $100,000 or more, respectively. Solid and dashed lines show trends for Maryland and the synthetic control, respectively. The time weights (λ_t) are estimated by minimizing the error term in Equation 1. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



(a) <35K

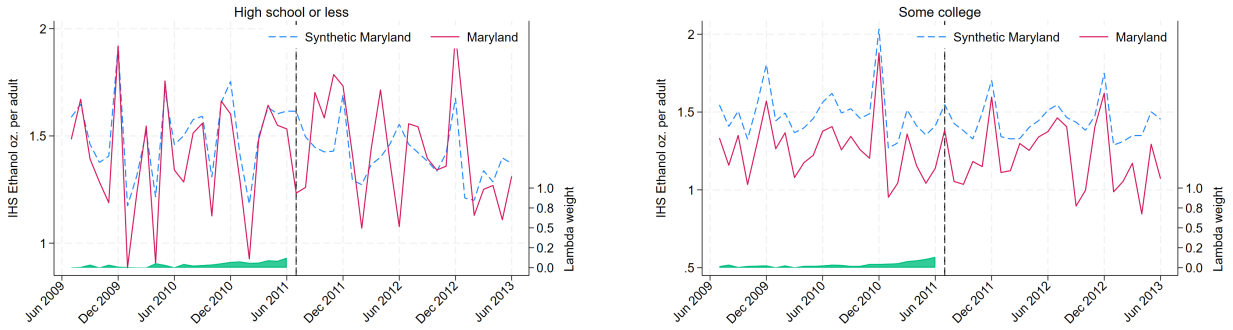
(b) $35–69K

(c) $70–99K

(d) $100K+

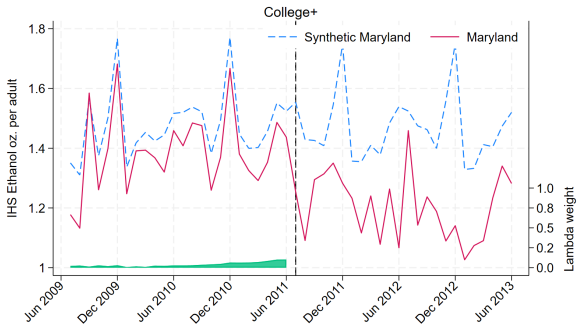
Figure A13. SDiD event study estimates of ethanol purchases by household income

Note. Synthetic Difference-in-Differences (SDiD) event study estimates for the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month, by household income. Panels (a)–(d) report annual household income groups of less than $35,000, $35,000–$69,999, $70,000–$99,999, and $100,000 or more, respectively. Monthly event-time coefficients are averaged by quarter. The vertical line at zero marks the tax change. Gray capped lines show 95% confidence intervals from 1,000 bootstrap replications. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



(a) High school or less

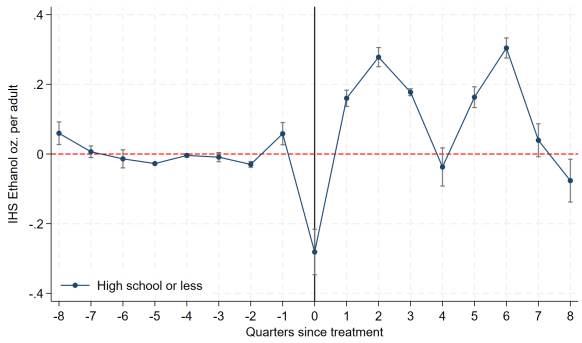
(b) Some college



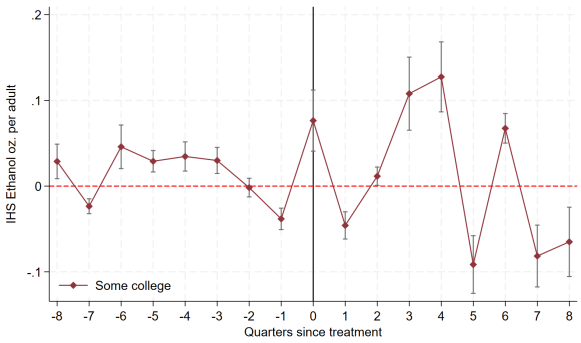
(c) College +

Figure A14. SDiD trends in ethanol purchases by household head's education

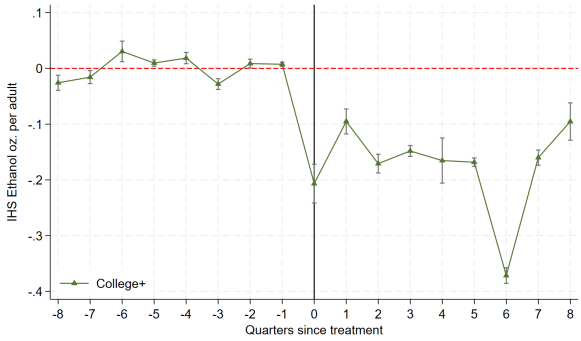
Note. Trends in ethanol ounces purchased per adult by household head's education, using Synthetic Difference-in-Differences (SDiD) weights. Panels (a)–(c) report high school or less, some college, and college or higher, respectively. Solid and dashed lines show trends for Maryland and the synthetic control, respectively. The time weights (λ_t) are estimated by minimizing the error term in Equation 1. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



(a) High school or less



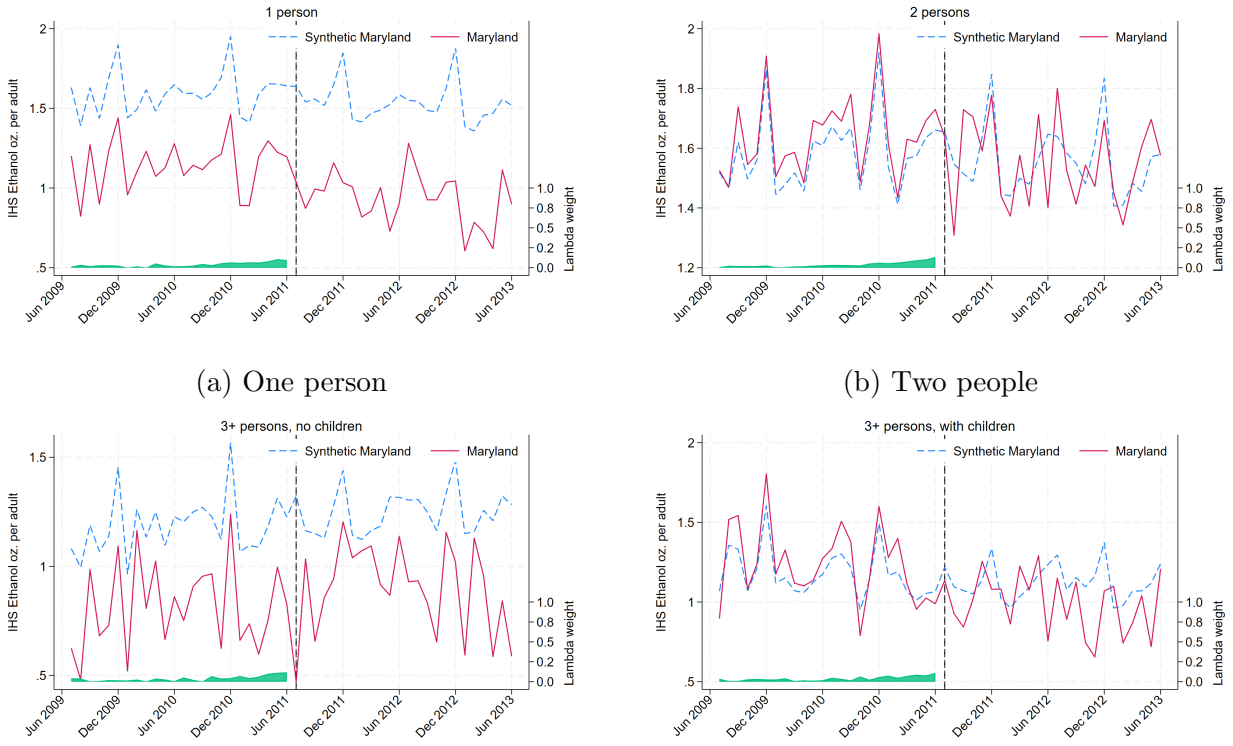
(b) Some college



(c) College +

Figure A15. SDiD event study estimates of ethanol purchases by household head's education

Note. Synthetic Difference-in-Differences (SDiD) event study estimates for the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month, by household head's education. Panels (a)–(c) report high school or less, some college, and college or higher, respectively. Monthly event-time coefficients are averaged by quarter. The vertical line at zero marks the tax change. Gray capped lines show 95% confidence intervals from 1,000 bootstrap replications. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



(a) One person

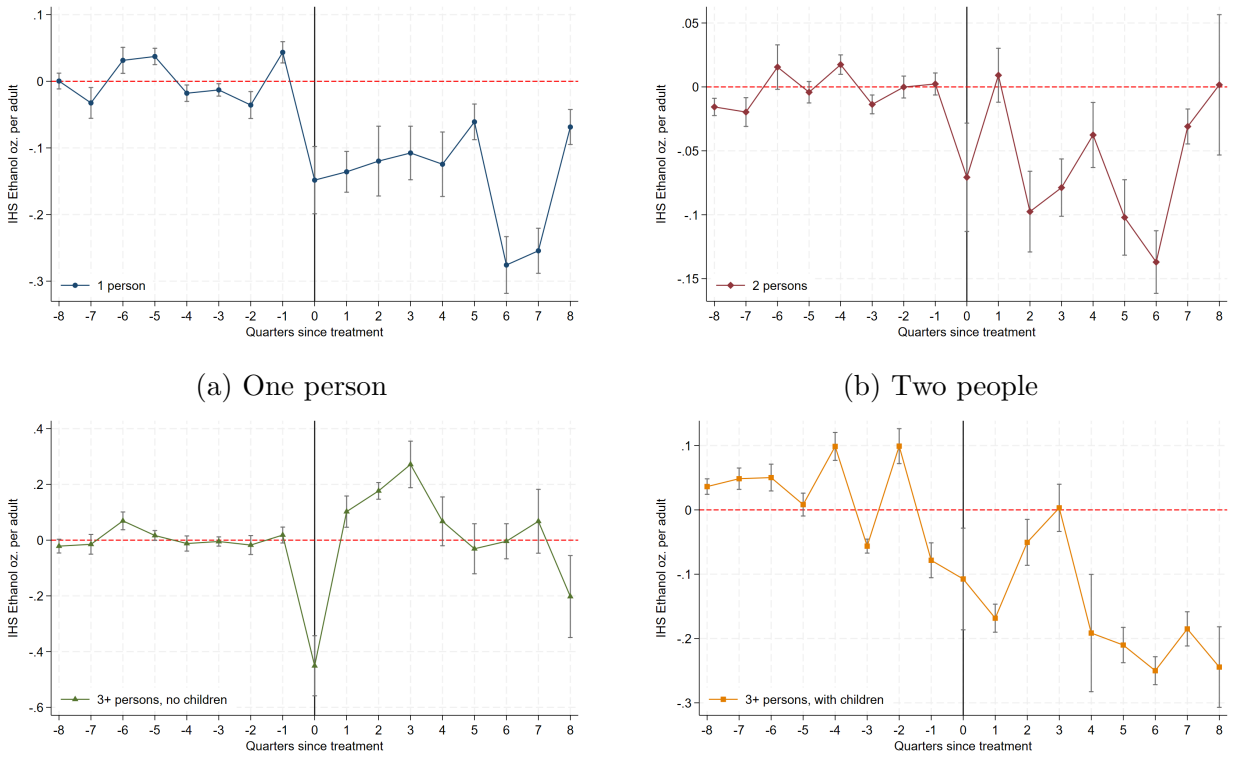
(b) Two people

(c) Three or more people, no children

(d) Three or more people, with children

Figure A16. SDiD trends in ethanol purchases by household composition

Note. Trends in ethanol ounces purchased per adult by household composition, using Synthetic Difference-in-Differences (SDiD) weights. Panels (a)–(d) report households with one person, two people, three or more people without children, and three or more people with children, respectively. Solid and dashed lines show trends for Maryland and the synthetic control, respectively. The time weights (λ_t) are estimated by minimizing the error term in Equation 1. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



(a) One person

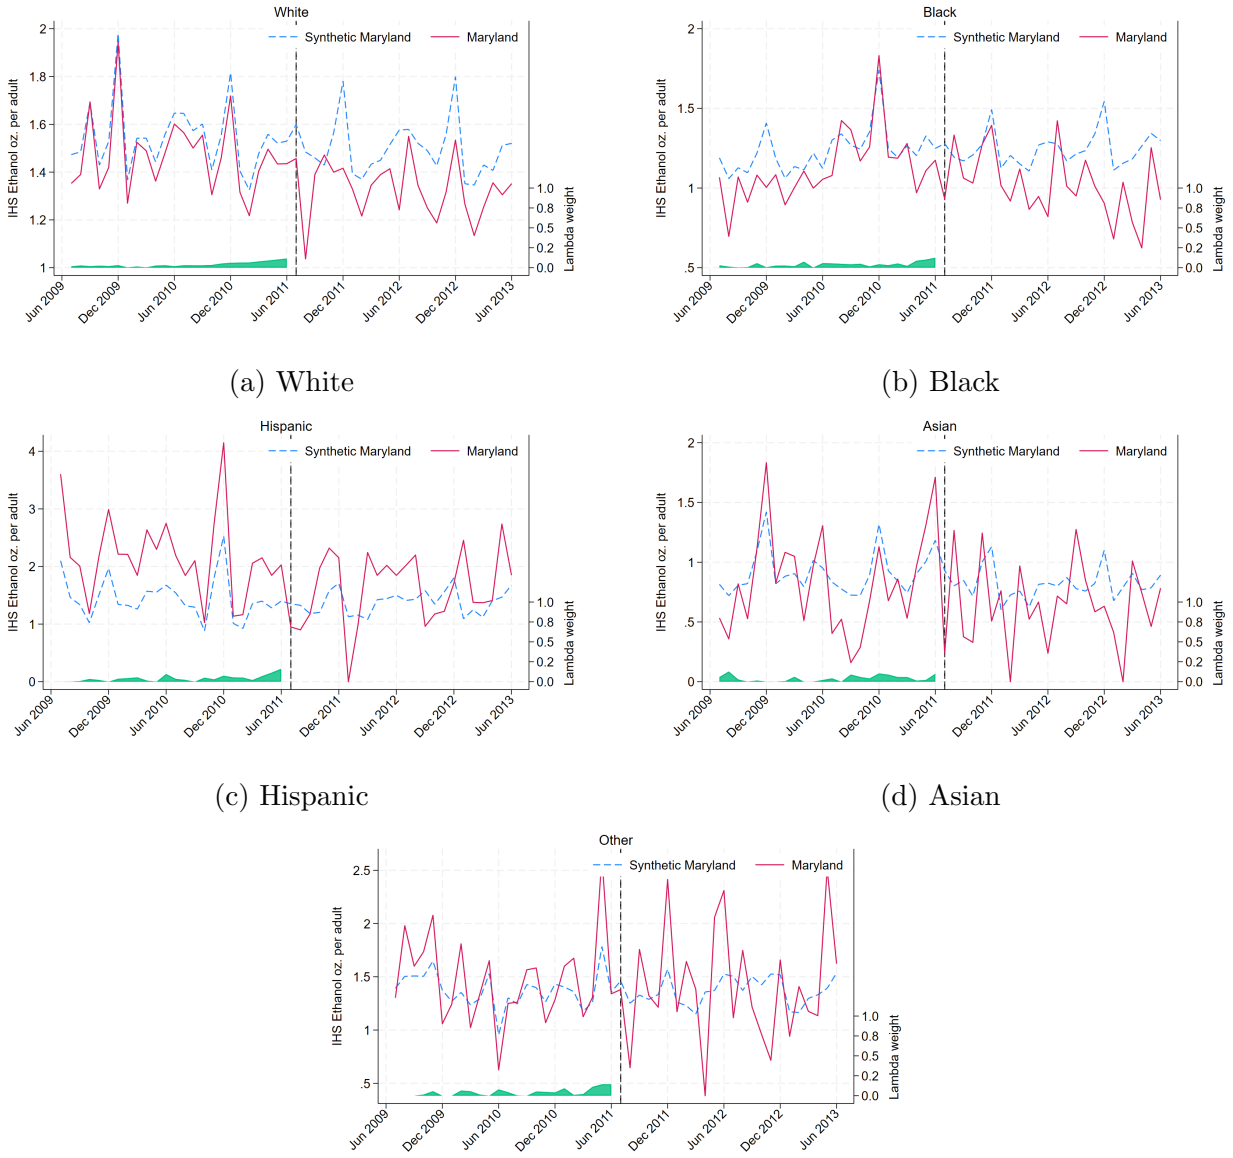
(b) Two people

(c) Three or more people, no children

(d) Three or more people, with children

Figure A17. SDiD event study estimates of ethanol purchases by household composition

Note. Synthetic Difference-in-Differences (SDiD) event study estimates for the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month, by household composition. Panels (a)–(d) report households with one person, two people, three or more people without children, and three or more people with children, respectively. Monthly event-time coefficients are averaged by quarter. The vertical line at zero marks the tax change. Gray capped lines show 95% confidence intervals from 1,000 bootstrap replications. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



(a) White

(b) Black

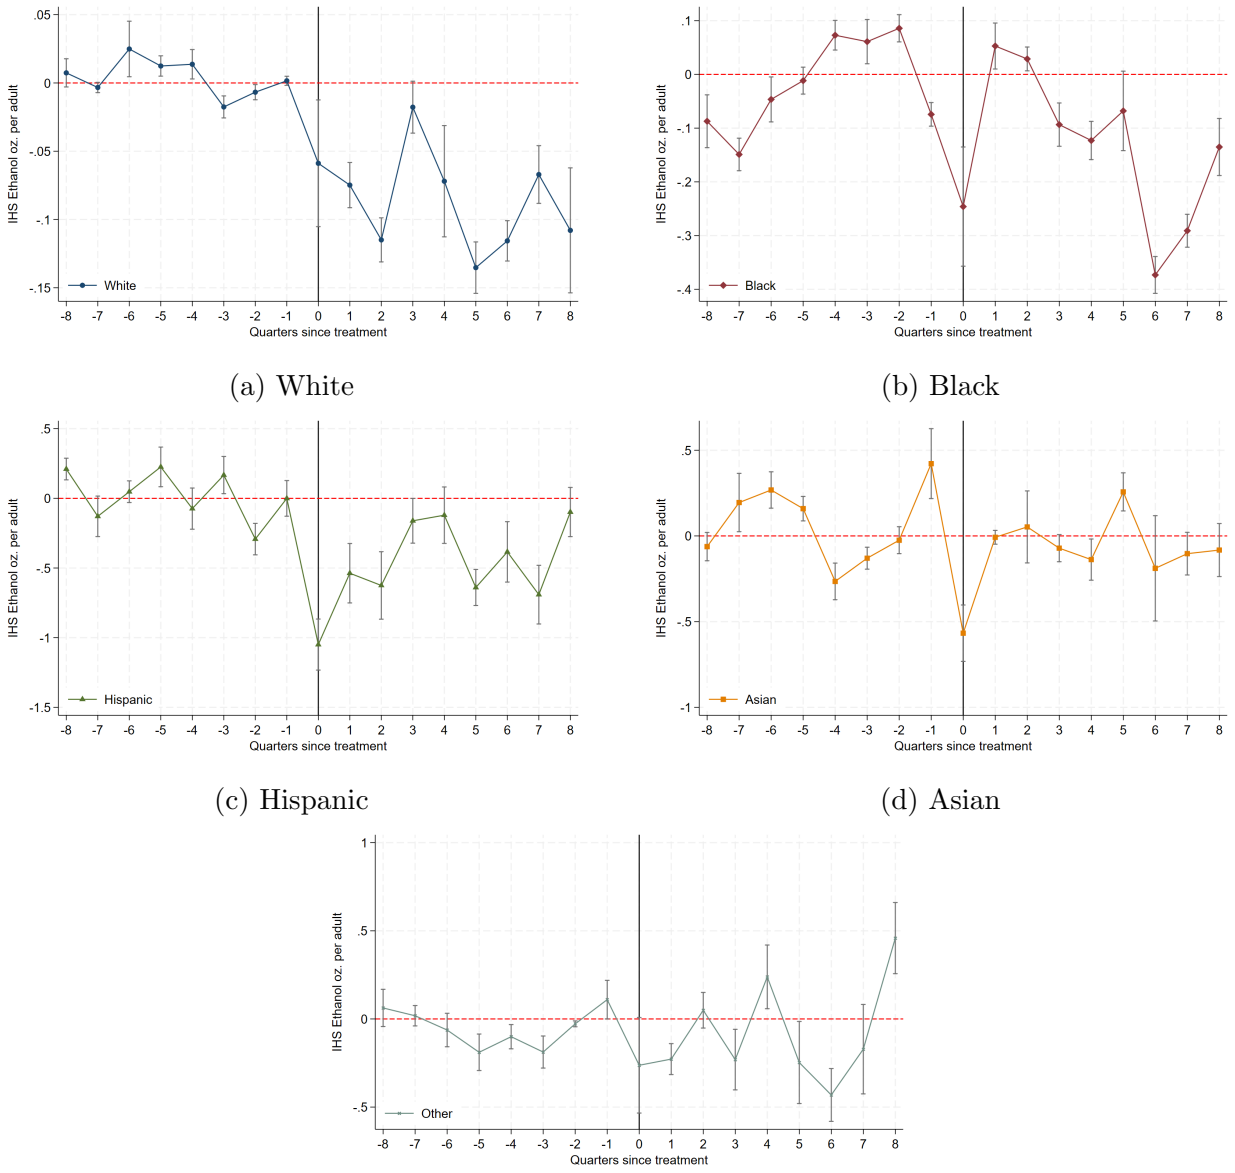
(c) Hispanic

(d) Asian

(e) Other

Figure A18. SDiD trends in ethanol purchases by race and ethnicity of household head

Note. Trends in ethanol ounces purchased per adult by race and ethnicity of household head, using Synthetic Difference-in-Differences (SDiD) weights. Panels (a)–(e) report White, Black, Hispanic, Asian, and Other groups, respectively. Solid and dashed lines show trends for Maryland and the synthetic control, respectively. The time weights (λ_t) are estimated by minimizing the error term in Equation 1. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



(a) White

(b) Black

(c) Hispanic

(d) Asian

(e) Other

Figure A19. SDiD event study estimates of ethanol purchases by race and ethnicity of household head

Note. Synthetic Difference-in-Differences (SDiD) event study estimates for the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month, by race and ethnicity of household head. Panels (a)–(e) report White, Black, Hispanic, Asian, and Other groups, respectively. Monthly event-time coefficients are averaged by quarter. The vertical line at zero marks the tax change. Gray capped lines show 95% confidence intervals from 1,000 bootstrap replications. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.

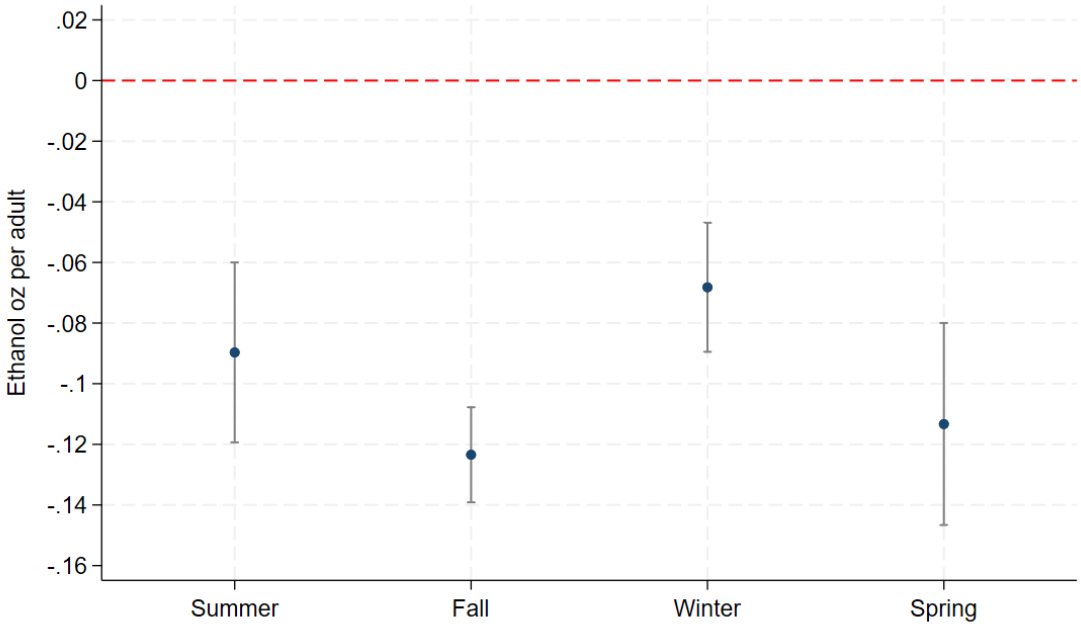


Figure A20. Two-way fixed effects estimates by season

Note. The figure plots two-way fixed effect coefficients by season: summer, fall, winter, and spring. The dependent variable is the inverse hyperbolic sine of ethanol ounces purchased per adult per household per month. 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 July 2009 to June 2013.

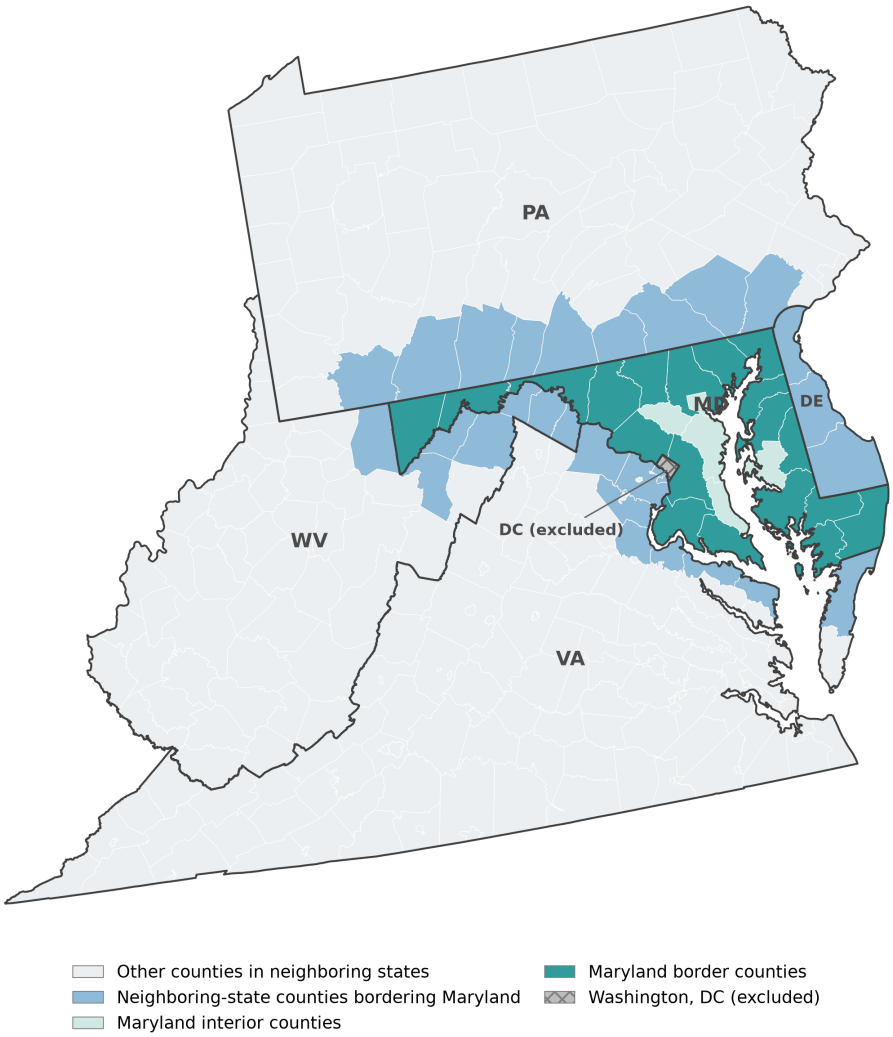
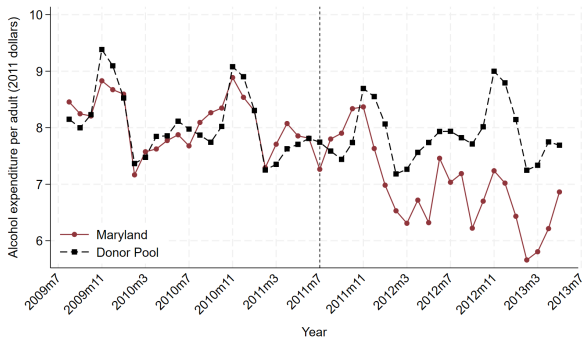
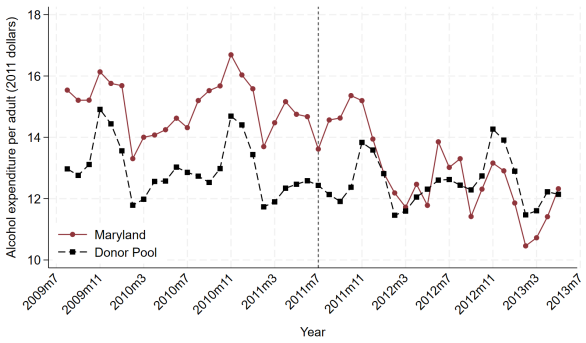


Figure A21. Maryland and surrounding states county borders

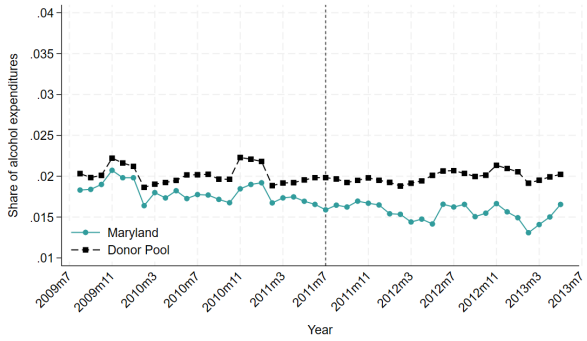
Note. The map shows Maryland, Washington, DC, and the four states bordering Maryland: Delaware, Pennsylvania, Virginia, and West Virginia. Maryland counties are classified as border or interior counties. Counties in neighboring states are classified according to whether they share a border with Maryland. Washington, DC is excluded from the analysis because it experienced a tax change during the sample period. Counties in neighboring states may be classified as adjacent to Maryland when their jurisdictional boundaries meet across a body of water.



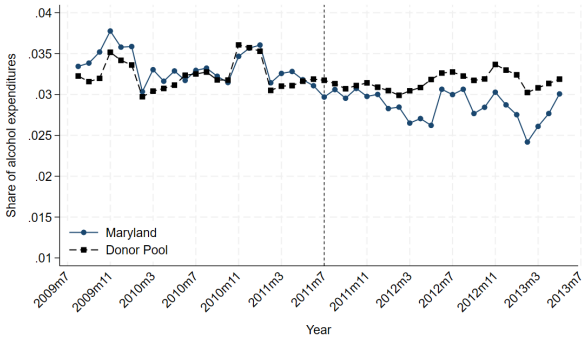
(a) Full sample: alcohol expenditure



(b) Drinking sample: alcohol expenditure



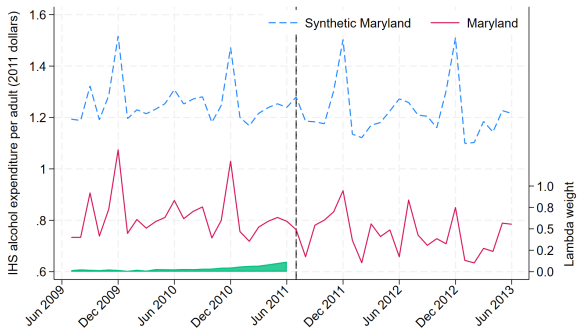
(c) Full sample: alcohol as share of total expenditure



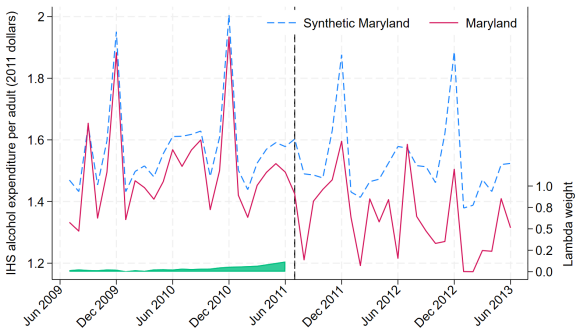
(d) Drinking sample:alcohol as share of total expenditure

Figure A22. Trends in alcohol expenditure and its share of total expenditure in Maryland and the donor pool

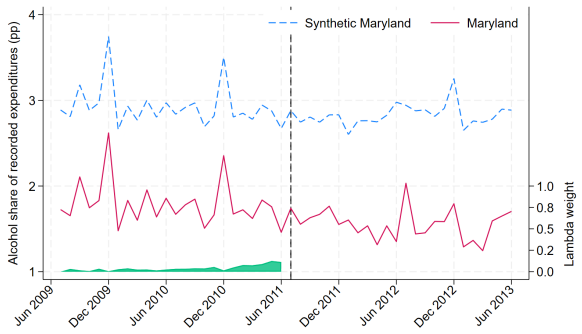
Note. Panels (a) and (b) present alcohol expenditure, and panels (c) and (d) present alcohol expenditure as a share of total expenditure. Panels (a) and (c) use the full sample; panels (b) and (d) restrict the sample to households that purchased alcohol at least once during the 12 months before the tax increase. Each panel plots the three-month moving average of the mean outcome for Maryland (solid line) and the donor pool (dashed line). The expenditure share is calculated as alcohol expenditures divided by total household expenditures recorded in the NielsenIQ Consumer Panel. 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.



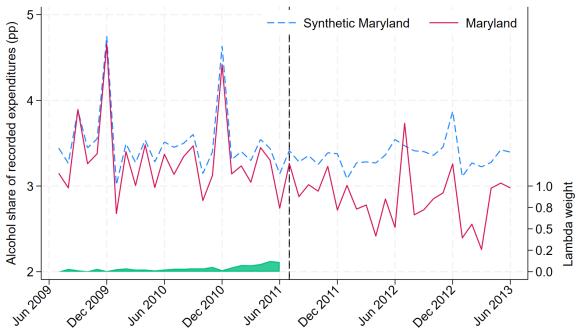
(a) Full sample: alcohol expenditure



(b) Drinking sample: alcohol expenditure



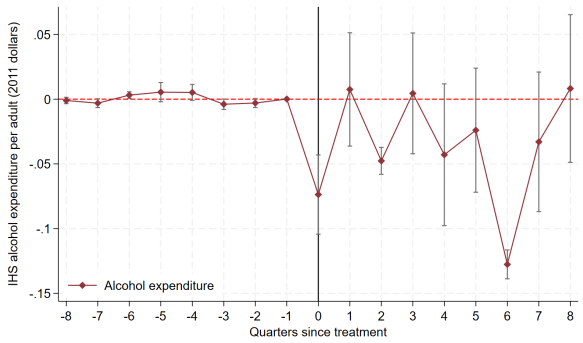
(c) Full sample: alcohol as share of total expenditure



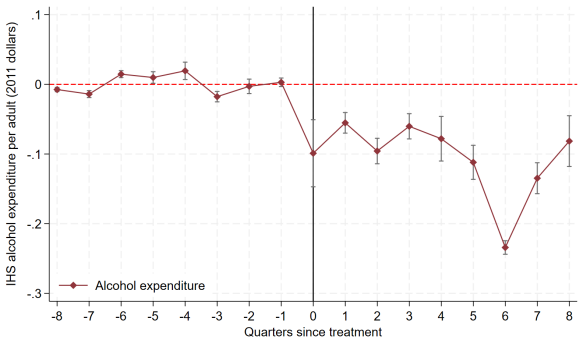
(d) SDiD trends in alcohol expenditure and its share of total expenditure

Figure A23. Synthetic difference-in-difference trends in alcohol expenditures and share of total expenses on alcohol

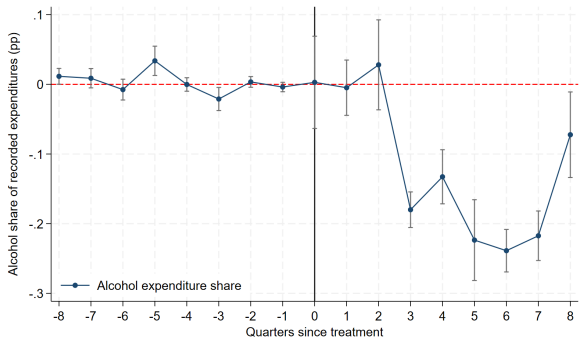
Note. Panels (a) and (b) present alcohol expenditure, and panels (c) and (d) present alcohol expenditure as a share of total expenditure. Panels (a) and (c) use the full sample; panels (b) and (d) restrict the sample to households that purchased alcohol at least once during the 12 months before the tax increase. Solid and dashed lines show trends for Maryland and the synthetic control constructed using Synthetic Difference-in-Differences (SDiD) weights, respectively. The expenditure share is calculated as alcohol expenditures divided by total household expenditures recorded in the NielsenIQ Consumer Panel. The time weights (λ_t) are estimated by minimizing the error term in Equation 1. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.



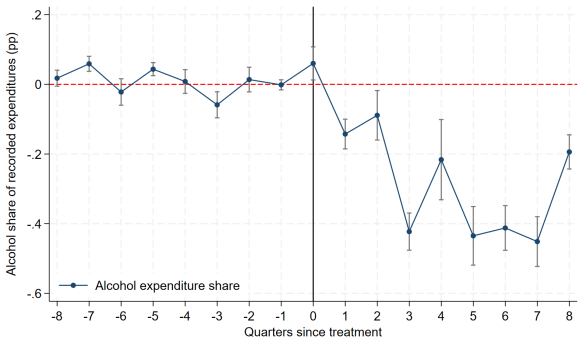
(a) Full sample: alcohol expenditure



(b) Drinking sample: alcohol expenditure



(c) Full sample: alcohol as share of total expenditure



(d) SDiD trends in alcohol expenditure and its share of total expenditure

Figure A24. Event study estimates for alcohol expenditures and share of total expenses on alcohol

Note. Synthetic Difference-in-Differences (SDiD) event study estimates for alcohol expenditure in panels (a) and (b), and alcohol expenditure as a share of total expenditure in panels (c) and (d). Panels (a) and (c) use the full sample; panels (b) and (d) restrict the sample to households that purchased alcohol at least once during the 12 months before the tax increase. The expenditure share is calculated as alcohol expenditures divided by total household expenditures recorded in the NielsenIQ Consumer Panel. Monthly event-time coefficients are averaged by quarter. The vertical line at zero marks the tax change. Gray capped lines show 95% confidence intervals from 1,000 bootstrap replications. Data are from the NielsenIQ Consumer Panel, July 2009–June 2013.

Table A1. States with Alcohol Sales Tax Changes from July 2009 to June 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). These dates represent en 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 July 2009 to June 2013

State	FIPS	Date	Beer tax	Wine tax	Spirit tax
Connecticut	9	5/4/2011	↑ $0.04	↑ $0.12	↑ $1.00
Illinois	17	9/1/2009	↑ $0.04	↑ $0.66	↑ $4.05
New Jersey	9	8/1/2011	↑ $0.04	↑ $0.18	↑ $1.10
New York	36	5/1/2009	↑ $0.03	↑ $0.11	-
North Carolina	37	9/1/2009	↑ $0.09	↑ $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. Robustness to alternative outcome transformations

	(1) Baseline	(2) Covariates
Panel A: Log of ethanol ounces purchased		
Treat × Post	-0.080***	-0.079***
	(0.005)	(0.005)
MD pre-tax mean	11.29	11.29
Panel B: IHS of alcohol ounces purchased		
Treat × Post	-0.150***	-0.153***
	(0.011)	(0.009)
Percent effect	-13.9	-14.2
MD pre-tax mean	113.47	113.47
Panel C: Poisson: ethanol ounces purchased		
Treat × Post	-0.091***	-0.101***
	(0.007)	(0.007)
MD pre-tax mean	11.29	11.29
Household FE	YES	YES
Year-month FE	YES	YES
Household controls	NO	YES

Note. The table reports estimates from alternative outcomes. Panel A presents synthetic difference-in-differences (SDiD) estimates where the outcome is the natural logarithm of one plus monthly ethanol ounces purchased per adult per household. Panel B presents SDiD estimates using the inverse hyperbolic sine of monthly alcohol ounces purchased per adult per household. Panel C presents estimates from a Poisson specification using monthly ethanol ounces purchased per adult per household. Column (1) reports baseline estimates, while column (2) additionally controls for household income, education, age, marital status, presence of children, and race. All specifications include household and year-month fixed effects. MD pre-tax mean refers to the mean of the corresponding untransformed outcome in Maryland prior to July 2011. Standard errors are clustered at the state level and reported in parentheses. For SDiD estimates, standard errors are obtained through bootstrapping with 1,000 replications. Data 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 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 A5. 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 A6. 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 A7. 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 A8. 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 A9. 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 A10. 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 A11. Border-county analysis within Maryland.

	(1) Expand treatment to include neighboring-state border households	(2) Exclude neighboring-state border households from donor pool
Treatment effect	-0.100***	-0.100***
	(0.020)	(0.008)
Observations	510,000	500,256
Treated households	408	205
Control households	10,217	10,217

Note. The table reports synthetic difference-in-differences estimates (SDiD) for the inverse hyperbolic sine (IHS) of monthly ethanol ounces purchased per adult per household. Column (1) expands the treated group to include households in neighboring-state counties that border Maryland. Column (2) retains Maryland households as the treated group and excludes neighboring-state border households from the donor pool. The models include household demographic covariates: income, education, age, marital status, presence of children, and race, as well as household and year-month fixed effects. Standard errors, reported in parentheses, are obtained using a state-clustered bootstrap. 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:

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

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.

Figure B1 presents the SDID event-study estimates from the jackknife analysis. The event-study confidence intervals are computed using a household-level bootstrap. State-level jackknife inference is infeasible in this setting as Maryland is the only treated state, while the jackknife requires at least two treated clusters. The figure's confidence intervals should therefore be interpreted as household-bootstrap intervals, not jackknife standard errors.

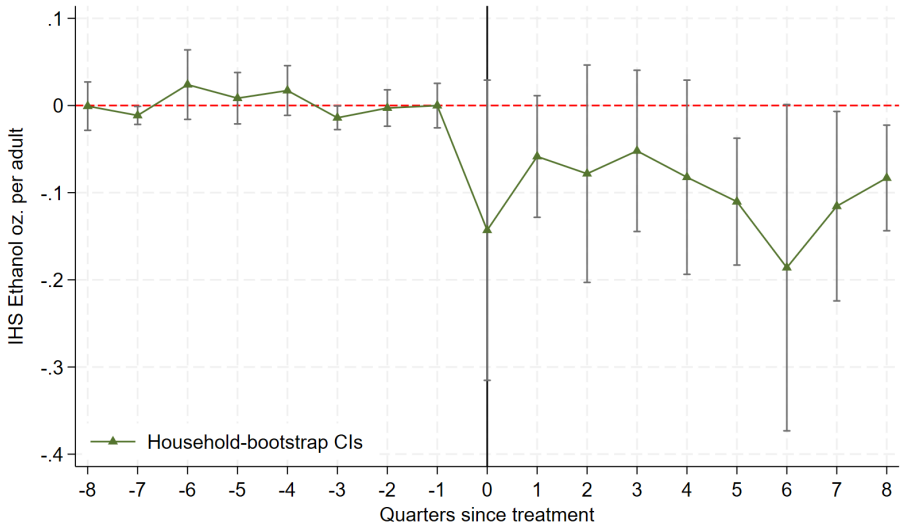


Figure B1. Dynamic Event study estimates for ethanol ounces purchased with Jackknife standard errors

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. Green triangles 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. Sample only includes households who have purchased alcohol at least once during the 12 months prior to the tax increase. Data on alcohol purchases come from the NielsenIQ Consumer panel for the period from July 2009 to June 2013.

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. Figure B2 presents the SDID event-study estimates from the placebo specification.

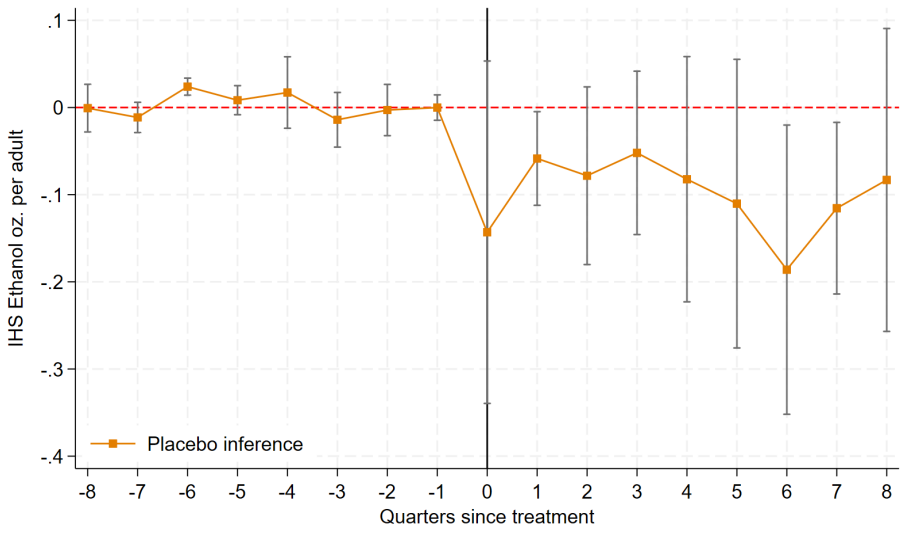


Figure B2. Dynamic Event study estimates for ethanol ounces purchased with Jackknife standard errors

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. Orange squares indicate point estimates, and gray capped lines represent 95% confidence intervals computed from 1,000 placebo 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. Sample only includes households who have purchased alcohol at least once during the 12 months prior to the tax increase. Data on alcohol purchases come from the NielsenIQ Consumer panel for the period from July 2009 to June 2013.

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)):

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

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).

Note. The figure plots standardized percentage bias against the residual variance ratio for covariates

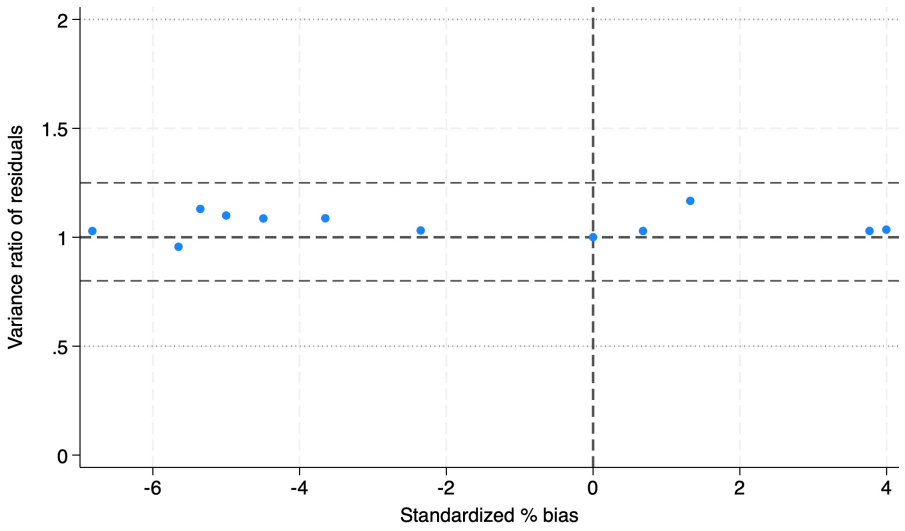


Figure C1. Standardized Bias and Variance ratio of the residuals

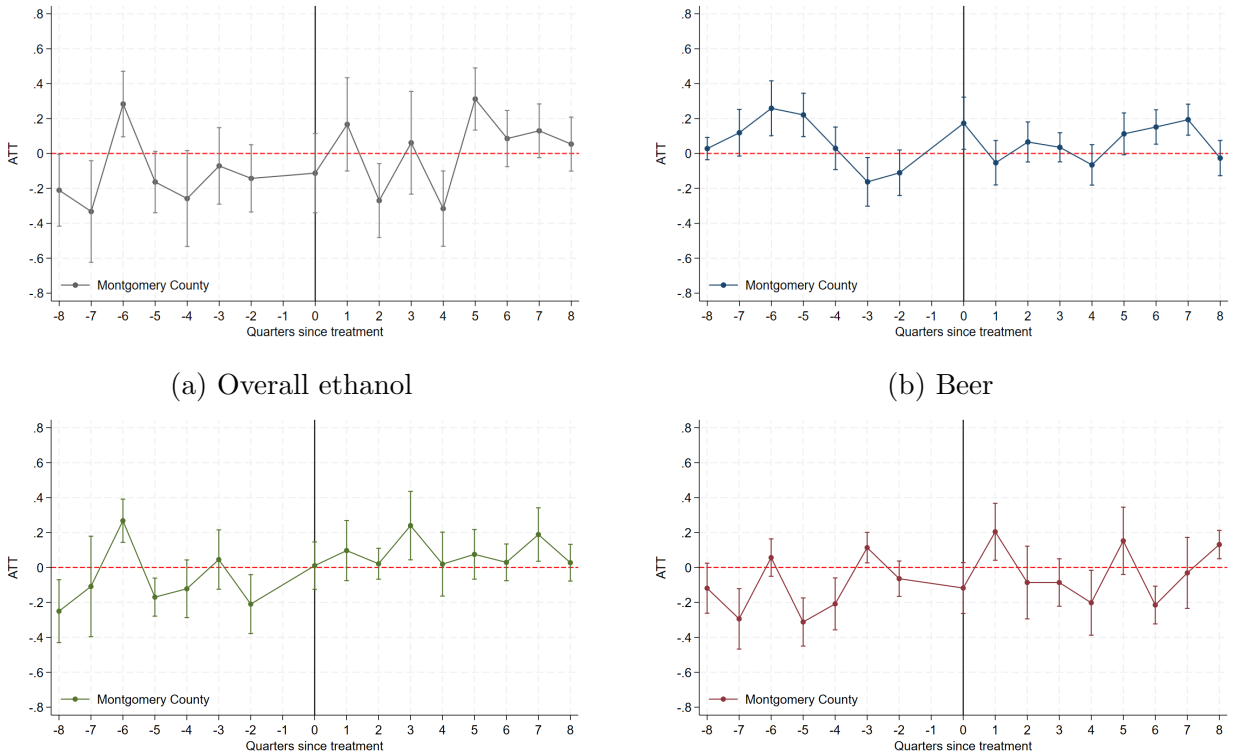
orthogonal to the linear propensity-score index. Standardized differences are expressed as percentage bias. Following Rubin (2001), the variance ratio compares the variance of the residualized covariates between treated and matched control households.

D Appendix: Montgomery County

D.1 Montgomery County

Now, I turn to Montgomery County, which is the only county in Maryland that operates as a control jurisdiction for alcohol distribution. The Montgomery County government controls the wholesale distribution of alcoholic beverages and the retail sale of distilled spirits, while other Maryland counties operate under private licensing systems. However, the statewide tax increase still applied to Montgomery County.

Because this analysis focuses exclusively on Maryland households (only the treated sample), the SDiD framework is not applicable. Instead, I estimate a traditional TWFE model with household and time fixed effects, clustering standard errors at the county level. By interacting event-time indicators with an indicator for Montgomery County, this specification estimates whether changes in ethanol purchases differ for Montgomery County households relative to households in other Maryland counties. Figure D1 presents the event-study estimates from this model for overall ethanol purchases (Panel D1a), beer purchases (Panel D1b), wine purchases (Panel D1c), and spirits purchases (Panel D1d).



(a) Overall ethanol

(b) Beer

(c) Wine

(d) Spirits

Figure D1. Alcohol purchasing patterns in Montgomery County relative to other Maryland counties

Note.The figure shows event-study estimates from a TWFE model restricted to Maryland households, comparing households in Montgomery County with households in other Maryland counties. Panel D1a presents overall ethanol purchases (gray), Panel D1b beer purchases (navy), Panel D1c wine purchases (olive green), and Panel D1d spirits purchases (maroon). The dependent variable is the inverse hyperbolic sine of monthly ethanol ounces purchased per adult. 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. Circles indicate point estimates, and capped lines represent 95% confidence intervals. Standard errors are clustered at the county level. The model includes controls for income, household size, marital status, presence of children, race, education, and age. Data on alcohol purchases come from the NielsenIQ Consumer Panel for the period July 2009 to June 2013.

Given the small sample size, these estimates should be interpreted descriptively rather than causally.[43] Across all outcomes, the pre-treatment coefficients do not have a clear flat pattern, making it difficult to interpret the post-treatment estimates as evidence of differential responses to the tax. Although the post-treatment estimates generally do not suggest large differences in purchasing patterns between Montgomery County and other Maryland counties, the pre-treatment trends prevent drawing strong conclusions from this comparison.

[43] In this sample, there are 135 households in Montgomery County.