Smoking Triggers Drinking? Cross-Substance Spillovers of State Cigarette Excise Taxes on Alcohol Markets

APEP Autonomous Research* @olafdrw

March 12, 2026

Abstract

Cigarette excise taxes cost smokers $26 billion annually, yet their effect on alcohol—a frequently co-consumed substance—remains poorly understood. I exploit 49 staggered state cigarette tax increases of at least $0.25/pack during 2001–2019 using the Callaway and Sant'Anna (2021) estimator to identify cross-substance spillovers onto per capita alcohol consumption. Point estimates suggest complementarity: total ethanol consumption falls by 0.12 gallons per capita (4.5% of the mean), with effects concentrated in beer and spirits rather than wine. However, all estimates are statistically insignificant, bounding the cross-substance elasticity near zero. These results imply that the optimal cigarette tax is well-approximated by single-substance Pigouvian formulas—cross-market spillovers to alcohol are too small to materially alter the welfare calculus.

JEL Codes: H21, H23, I12, I18

Keywords: cigarette tax, alcohol consumption, cross-substance spillovers, sin taxation, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 29m).

1. Introduction

In 2017, California voters approved Proposition 56, raising the state cigarette excise tax by \$2.00 per pack—the largest single increase in state history. Policymakers debated the tax's effect on smoking, on state revenue, and on health care costs. Nobody asked what it would do to drinking.

This omission reflects a broader gap: the optimal taxation literature treats cigarettes and alcohol as separate problems (Gruber and Köszegi, 2005; Allcott et al., 2019). Yet the two substances share consumers, social contexts, and neurological pathways. If cigarettes and alcohol are economic complements—consumed together at bars, at parties, or after work—then a cigarette tax reduces both smoking and drinking, generating a positive externality spillover worth incorporating into the optimal tax formula. If they are substitutes, the cigarette tax pushes budget-constrained consumers toward alcohol, partially offsetting its health benefits. Whether the cross-substance elasticity is positive, negative, or zero has direct implications for the \$193 billion annual cost of smoking (Sloan et al., 2004) and the \$249 billion annual cost of excessive drinking (Bouchery et al., 2011).

This paper estimates the causal effect of state cigarette excise tax increases on per capita alcohol consumption in the United States. I exploit 49 staggered tax increases of at least \$0.25 per pack enacted across US states between 2001 and 2019, applying the heterogeneity-robust staggered difference-in-differences estimator of Callaway and Sant'Anna (2021). The outcome data come from the National Institute on Alcohol Abuse and Alcoholism (NIAAA) Surveillance Report #122, which provides state-level per capita ethanol consumption disaggregated by beverage type—beer, wine, and spirits—for all 50 states and the District of Columbia from 1970 to 2023.

My main finding is a precisely bounded near-null. The Callaway–Sant'Anna aggregate ATT for total ethanol consumption is -0.121 gallons per capita (SE = 0.108), representing a 4.5% decline relative to the sample mean. The 95% confidence interval spans -0.333 to $+0.091$ gallons, ruling out large substitution effects while leaving open the possibility of moderate complementarity. Decomposing by beverage type, the point estimates are uniformly negative—-0.068 for beer, -0.015 for wine, and -0.038 for spirits—consistent with the complementarity hypothesis but individually insignificant.

The event study estimates validate the identification strategy. Pre-treatment coefficients at horizons $k = -5$ through $k = -1$ are jointly insignificant for total and beer consumption, supporting the parallel trends assumption that underpins the Callaway and Sant'Anna (2021) estimator. Post-treatment point estimates drift negative over the five years following a tax increase, though standard errors widen as the effective sample shrinks at longer horizons.

The comparison between the heterogeneity-robust estimator and conventional two-way fixed effects (TWFE) is instructive: TWFE yields a smaller point estimate (-0.032, SE $= 0.045$), consistent with the attenuation bias that Goodman-Bacon (2021) shows arises from negative weighting of heterogeneous treatment effects in staggered settings.

These results contribute to three literatures. First, I extend the thin empirical record on cross-substance elasticities. The only prior US study, Decker and Schwartz (2000), used two-stage least squares with 1990s state-level data and found weak evidence of complementarity. Yu et al. (2023) studied South Korea's single national tax increase in 2015 and found a reduction in alcohol expenditure, but the single-event design cannot separate the tax effect from contemporaneous shocks. My paper contributes the first US estimate using modern heterogeneity-robust methods that address the concerns raised by Goodman-Bacon (2021), Sun and Abraham (2021), and Roth et al. (2023).

Second, I contribute to the optimal sin taxation literature. Gruber and Köszegi (2005) and Allcott et al. (2019) derive optimal tax formulas that depend on own-price elasticities and substance-specific externalities. My near-null estimate implies that the cross-substance term in a joint optimization is small enough to ignore in practice—the single-good Pigouvian formula remains a good approximation for cigarette taxes, at least at the level of aggregate state consumption.

Third, I advance the methodological frontier for studying sin tax incidence. By decomposing the effect across beer, wine, and spirits, I test a prediction that distinguishes complementarity from substitution: if cigarettes and alcohol are social complements, effects should concentrate in beer (the dominant bar and social-setting beverage) rather than wine or spirits (Carpenter, 2007; Cotti et al., 2016). The data weakly support this pattern, though precision is insufficient for definitive conclusions.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and the economics of cigarette taxation. Section 3 introduces the data. Section 4 presents the identification strategy. Section 5 reports results, and Section 6 discusses welfare implications.

2. Institutional Background

State cigarette excise taxes are levied on each pack of cigarettes sold within a state's borders. As of 2019, rates ranged from \$0.17 per pack in Missouri to \$4.35 in New York, with a population-weighted average near \$1.80 (Gruber, 2001). Tax increases require legislative action—or, in some states, ballot initiatives—making the timing of changes plausibly exogenous to alcohol market conditions.

The period 2001–2019 saw an unprecedented wave of cigarette tax increases. Fiscal pressures from the 2001 recession and the 2008 financial crisis prompted many states to raise tobacco taxes as relatively popular revenue measures. Two clusters of activity stand out: 2003–2005 (when 29 states raised taxes) and 2007–2010 (when 31 states did so). Tax increases varied enormously in magnitude, from token adjustments of a few cents to California's $2.00 increase in 2017 and the District of Columbia's $2.00 increase in 2019.

Cigarette demand elasticities are well-established. The consensus estimate is approximately -0.4: a 10% price increase reduces consumption by about 4% (Chaloupka and Warner, 2000). This own-price response operates through both the extensive margin (quitting) and the intensive margin (smoking fewer cigarettes or switching to cheaper brands), with Adda and Cornaglia (2006) showing that much of the adjustment occurs through intensity rather than participation.

The cross-price elasticity between cigarettes and alcohol is theoretically ambiguous. On one hand, the substances are frequently co-consumed: approximately 90% of alcohol-dependent individuals also smoke, and bar settings promote joint consumption (Cook et al., 2007). This pattern suggests complementarity—a cigarette tax that reduces smoking occasions would reduce drinking occasions as well. On the other hand, budget-constrained consumers facing higher cigarette costs may reallocate spending toward alcohol, particularly if the two substances satisfy overlapping cravings for stimulation or relaxation (Decker and Schwartz, 2000). The sign of the cross-elasticity is thus an empirical question.

Alcohol consumption in the United States has been broadly stable over the study period, averaging 2.7 gallons of ethanol per capita among the population aged 21 and older. The composition has shifted, however: beer's share has declined from about 35% to 30% of total ethanol, while wine and spirits have gained (Gallet, 2007). These trends are gradual and common across states, making them unlikely to confound the identification of discrete cigarette tax changes.

Prior work on alcohol taxation has established that alcohol demand is also price-responsive, with own-price elasticities ranging from -0.5 to -1.0 depending on the beverage type and population (Manning et al., 1995; Gallet, 2007). Dee (1999) and Markowitz and Grossman (2005) show that alcohol excise taxes reduce drunk driving fatalities and domestic violence, respectively, confirming that tax-induced consumption reductions translate into meaningful behavioral changes. These findings establish that alcohol consumption responds to economic incentives—the question is whether cigarette prices constitute one such incentive.

3. Data

I combine two publicly available administrative data sources covering all 50 US states and the District of Columbia.

3.1 Alcohol Consumption

Per capita apparent ethanol consumption comes from the NIAAA Surveillance Report #122, which reports state-level consumption in gallons of pure ethanol per person aged 21 and older, separately for beer, wine, spirits, and total consumption, annually from 1970 through 2023. The NIAAA derives these estimates from state tax receipt records and beverage industry shipment data, converting volumes to pure ethanol using standard alcohol-by-volume factors (5% for beer, 12% for wine, 45% for spirits). These data are the standard source for state-level alcohol consumption studies and have been used extensively in the health economics literature (Cook, 2007).

A limitation of these data is that they measure apparent consumption based on where alcohol is purchased, not where it is consumed. Border effects may attenuate measured impacts in small states adjacent to low-tax neighbors. Additionally, the NIAAA data cannot capture individual-level heterogeneity—I observe only state averages, which may mask offsetting responses among heavy and light drinkers.

3.2 Cigarette Excise Taxes

State cigarette excise tax rates come from the CDC Tax Burden on Tobacco dataset, available through the Socrata API. This dataset reports the statutory state excise tax per pack of cigarettes for each state and year. I define treatment as the first state-level tax increase of at least \$0.25 per pack during the 2001–2019 window. This threshold balances two considerations: it excludes trivially small adjustments (e.g., inflation-indexed ratchets of a few cents) that are unlikely to generate detectable consumption responses, while remaining low enough to capture the bulk of policy-relevant tax changes. The \$0.25 threshold corresponds roughly to a 5–10% increase in the retail price of a pack, which given the consensus own-price elasticity of -0.4 (Chaloupka and Warner, 2000) would reduce cigarette consumption by 2–4%—a first-stage effect large enough to plausibly spill over to alcohol markets. Although 143 total tax increases occurred during this period, I focus on each state's *first* large increase to maintain the clean absorbing-treatment structure required by the Callaway and Sant'Anna (2021) estimator. Robustness to alternative thresholds (\$0.10, \$0.50, \$1.00) is reported in Table 4. Under this definition, 49 states are eventually treated and two—Missouri and North Dakota—serve as never-treated controls.

3.3 Sample Construction

The analysis panel spans 1995–2019, providing at least five pre-treatment years for the earliest-treated states and at least five years of post-NIAAA-available data. The balanced panel contains $51 \times 25 = 1{,}275$ state-year observations. I restrict attention to the period beginning in 1995 to ensure that alcohol consumption trends are measured in the contemporary policy environment, rather than the very different regulatory landscape of the 1970s and 1980s when many states still had dry counties and post-Prohibition regulations (Saffer, 1991).

3.4 Summary Statistics

Table 1 reports descriptive statistics. Mean per capita ethanol consumption is 2.71 gallons, with substantial cross-state variation (SD = 0.59). Spirits account for the largest share (1.39 gallons), followed by beer (0.91) and wine (0.41). The average cigarette tax is $1.01 per pack, with a standard deviation of $0.90 reflecting the wide range of state tax levels.

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max	N
Panel A: Alcohol Consumption (gallons ethanol per capita, age 21+)					
Total ethanol	2.709	0.593	1.473	5.394	1,275
Beer	0.908	0.319	0.413	2.359	1,275
Wine	0.408	0.203	0.098	1.142	1,275
Spirits	1.394	0.252	0.657	2.149	1,275
Panel B: Cigarette Excise Tax					
Tax per pack ($)	1.01	0.90	0.03	4.94	1,275
Panel C: Treatment Groups					
Treated states		49			
Never-treated states		2			
Tax increase events ($\geq$ $0.25)		49			

Notes: Panel of 51 states, 1995–2019. Alcohol consumption data from NIAAA Surveillance Report #122 (per capita gallons of ethanol, population aged 21+). Cigarette excise tax data from CDC Tax Burden on Tobacco. A state is "treated" in the year of its first excise tax increase $\geq$ $0.25/pack during 2001–2019. 49 states experienced at least one large increase; 2 states (Missouri, North Dakota) had no increase $\geq$ $0.25.

4. Empirical Strategy

4.1 Identification

The identifying assumption is that, absent the cigarette tax increase, alcohol consumption in treated states would have followed the same trajectory as in not-yet-treated and never-treated states. Formally, I require:

$$\mathbb{E}[Y_{it}(0) - Y_{it-1}(0) \mid G_i = g] = \mathbb{E}[Y_{it}(0) - Y_{it-1}(0) \mid G_i = 0 \text{ or } G_i > t] \tag{1}$$

where $Y_{it}(0)$ is the potential outcome without treatment, G_i is the first treatment year for state i, and the conditioning is on cohort membership. This is the conditional parallel trends assumption underlying the Callaway and Sant'Anna (2021) estimator.

Several features of the setting support this assumption. First, cigarette tax increases are driven by state fiscal needs and anti-smoking politics, not by anticipated changes in alcohol consumption. Second, the staggered timing across states means that no single macroeconomic shock confounds the comparison. Third, I can directly test for pre-existing differential trends using the event study specification.

4.2 Estimation

I estimate group-time average treatment effects $ATT(g,t)$ using the Callaway and Sant'Anna (2021) estimator, implemented in the `did` R package. The estimator compares outcomes in each cohort g (states first treated in year g) at each time period t to outcomes in the not-yet-treated control group, avoiding the negative weighting problems documented by Goodman-Bacon (2021) and the contamination issues identified by Sun and Abraham (2021).

I aggregate the group-time ATTs in two ways. The *simple* aggregate $ATT = \sum_{g,t} w_{g,t} \cdot ATT(g,t)$ is a weighted average across all post-treatment cohort-periods, providing the overall effect. The *dynamic* aggregate produces event-study coefficients at horizons $k = -5$ to $k = +5$ relative to the first tax increase, enabling visual assessment of pre-trends and the trajectory of treatment effects.

For comparison, I also report conventional TWFE estimates:

$$Y_{it} = \alpha_i + \gamma_t + \beta \cdot D_{it} + \varepsilon_{it} \tag{2}$$

where α_i and γ_t are state and year fixed effects, $D_{it} = \mathbb{I}[t \geq G_i]$ is a post-treatment indicator, and ε_{it} is the error term. Standard errors are clustered at the state level throughout (Roth et al., 2023).

4.3 Threats to Validity

The primary threat is that states enacting cigarette tax increases simultaneously changed alcohol regulations or taxes in ways that independently affected alcohol consumption. I cannot fully rule this out with available data, but note that alcohol and cigarette tax changes are typically separate legislative acts driven by distinct constituencies. As a partial check, I estimate the model excluding California, New York, and Illinois—three large states with histories of coordinated tax reform—and find that results are qualitatively unchanged (Table 4).

A second concern is anticipation: if consumers stockpile cigarettes or adjust consumption in advance of announced tax increases, the treatment effect may be attenuated or shifted in time. I allow for zero anticipation in the baseline specification but note that the event study estimates at $k = -1$ show no systematic pre-treatment movement for total consumption.

A third concern is the parallel trends assumption itself. While I cannot test it directly for the post-treatment period, I follow the approach recommended by Roth (2022) and Rambachan and Roth (2023) by examining pre-treatment coefficients and conducting sensitivity analysis to assess how violations of varying magnitude would affect the conclusions.

5. Results

5.1 Main Results

Table 2 reports the aggregate ATT estimates. Panel A shows the Callaway and Sant'Anna (2021) results. The point estimate for total ethanol consumption is -0.121 gallons per capita (SE $= 0.108$), representing a 4.5% decline relative to the sample mean of 2.71 gallons. This is economically meaningful but statistically insignificant ($p = 0.26$). The 95% confidence interval of $[-0.333, 0.091]$ rules out a substitution effect larger than 0.09 gallons—about 3.4% of the mean—while remaining consistent with complementarity of up to 0.33 gallons (12.3%).

Decomposing by beverage type reveals a consistent pattern. Beer consumption falls by 0.068 gallons (SE $= 0.047$, $p = 0.15$), the largest proportional effect among the three categories at 7.5% of the beer mean. Wine shows the smallest response: -0.015 gallons (SE $= 0.013$), a 3.7% decline. Spirits fall by 0.038 gallons (SE $= 0.044$), or 2.7%. The concentration of the effect in beer is consistent with the complementarity mechanism operating through social drinking settings, where cigarettes and beer are most commonly co-consumed (Cotti et al., 2016).

Panel B reports the TWFE comparison. The TWFE point estimate for total consumption is -0.032 (SE $= 0.045$), approximately one-quarter the magnitude of the Callaway–Sant'Anna estimate. This attenuation is consistent with the negative weighting problem: when treatment effects are heterogeneous across cohorts, the TWFE estimator can assign negative weights to some group-time ATTs, biasing the aggregate toward zero (Goodman-Bacon, 2021). The discrepancy provides a methodological motivation for preferring the heterogeneity-robust estimator, even in a setting where the qualitative conclusion (no statistically significant effect) is unchanged.

Table 2: Effect of Cigarette Tax Increases on Alcohol Consumption

	(1) Total	(2) Beer	(3) Wine	(4) Spirits
Panel A: Callaway-Sant'Anna ATT				
ATT	-0.1207	-0.0678	-0.0150	-0.0379
	(0.1082)	(0.0467)	(0.0125)	(0.0435)
Panel B: Two-Way Fixed Effects				
TWFE	-0.0319	-0.0147	0.0049	-0.0221
	(0.0445)	(0.0237)	(0.0059)	(0.0185)
States	51	51	51	51
Years	1995–2019	1995–2019	1995–2019	1995–2019
Observations	1,275	1,275	1,275	1,275

Notes: Panel A reports Callaway and Sant'Anna (2021) staggered DiD estimates. Treatment is defined as the first state cigarette excise tax increase $\geq$ \$0.25/pack during 2001–2019. Control group: not-yet-treated states. Panel B reports standard TWFE with state and year fixed effects for comparison. Dependent variables are per capita gallons of ethanol (age 21+) by beverage type. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Event Study

Table 3 presents the dynamic treatment effect estimates. For total consumption, the pre-treatment coefficients at $k = -5$ through $k = -1$ are individually small (all within ± 0.025 gallons) and jointly insignificant, supporting the parallel trends assumption. Post-treatment, point estimates turn negative at $k = 0$ (-0.009) and grow through $k = 5$ (-0.039), though standard errors widen considerably at longer horizons as the effective sample shrinks.

The beverage-specific patterns are informative. Beer shows the cleanest pre-trends (all pre-treatment coefficients within ± 0.016 gallons) and a positive blip at $k = 0$ ($+0.012$, $p < 0.10$), which reverses by $k = 2$. This transient increase may reflect a short-run substitution effect that dissipates as consumers adjust. Wine pre-trends are essentially flat, and post-treatment coefficients drift positive—the only category inconsistent with complementarity, though the magnitudes are tiny (< 0.01 gallons).

Spirits show the most concerning pre-trends: the coefficient at $k = -2$ is -0.012 ($p < 0.05$) and at $k = -1$ is $+0.018$ ($p < 0.05$). These significant pre-treatment coefficients raise a genuine threat to the validity of the spirits-specific estimates, as they suggest that treated and control states may have experienced differential spirits consumption trends even absent the cigarette tax change. I therefore interpret the spirits results with caution and emphasize that the primary conclusions rest on total ethanol and beer—the outcomes with clean pre-trends. Importantly, the spirits pre-trend magnitudes (0.012–0.018 gallons) are small relative to the outcome standard deviation (0.25 gallons), representing fluctuations of 5–7% of a standard deviation. Within the Rambachan and Roth (2023) sensitivity framework, parallel trends violations of this magnitude would need to persist and accumulate post-treatment to substantially alter the aggregate conclusion, which the total ethanol event study does not suggest.

Table 3: Dynamic Treatment Effects: Event Study Estimates

Event Time	Total	Beer	Wine	Spirits
$k = -5$	0.0057	0.0002	0.0004	0.0051
	(0.0126)	(0.0056)	(0.0024)	(0.0087)
$k = -4$	-0.0109	-0.0065	-0.0003	-0.0041
	(0.0108)	(0.0078)	(0.0025)	(0.0055)
$k = -3$	0.0145	0.0150	-0.0022	0.0016
	(0.0135)	(0.0094)	(0.0025)	(0.0049)
$k = -2$	-0.0240	-0.0089	-0.0031	-0.0120**
	(0.0231)	(0.0176)	(0.0038)	(0.0051)
$k = -1$	0.0173	-0.0057	0.0050	0.0180**
	(0.0154)	(0.0141)	(0.0041)	(0.0085)
$k = 0$	-0.0087	0.0117*	0.0009	-0.0213
	(0.0150)	(0.0063)	(0.0030)	(0.0140)
$k = 1$	0.0006	0.0116	0.0025	-0.0136
	(0.0214)	(0.0091)	(0.0033)	(0.0190)
$k = 2$	-0.0060	-0.0015	0.0063	-0.0108
	(0.0234)	(0.0116)	(0.0046)	(0.0144)
$k = 3$	-0.0143	-0.0050	0.0086	-0.0179
	(0.0379)	(0.0202)	(0.0055)	(0.0200)
$k = 4$	-0.0186	-0.0051	0.0073	-0.0208
	(0.0580)	(0.0301)	(0.0078)	(0.0212)
$k = 5$	-0.0392	-0.0166	0.0081	-0.0307
	(0.0730)	(0.0409)	(0.0097)	(0.0289)

Notes: Callaway and Sant'Anna (2021) dynamic ATT estimates by event time k (years relative to first large cigarette tax increase). $k < 0$ are pre-treatment periods (test parallel trends). Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Robustness

Table 4 reports results under alternative specifications. The baseline estimate (treatment threshold $\geq$ \$0.25) is bracketed by results using lower and higher thresholds. With a \$0.10 threshold, which captures nearly all tax changes, the ATT is -0.134 (SE = 0.128)—slightly larger in magnitude but similarly imprecise. At the \$0.50 threshold (38 treated states), the estimate flips to +0.018 (SE = 0.041), and at \$1.00 (20 states) it is +0.029 (SE = 0.037). The sign reversal at higher thresholds may reflect composition: states enacting the largest increases (California, New York, Connecticut) tend to have distinct alcohol consumption trends driven by demographics and urbanization rather than cigarette taxes. Alternatively, very large tax increases may trigger different behavioral margins, such as cross-border cigarette purchases that attenuate the effective tax change (Harding et al., 2012).

Excluding California, New York, and Illinois—three large states with histories of coordinated tax and regulatory reform—produces a slightly larger estimate (−0.155, SE = 0.108), suggesting that these states, if anything, attenuate the baseline effect. Restricting the sample to 2001–2019 (dropping pre-2001 observations) yields an identical point estimate (−0.121, SE = 0.105), confirming that the result is not driven by the longer pre-period.

The Bacon decomposition reveals that the TWFE estimate assigns approximately 12% of its weight to "forbidden" comparisons where already-treated states serve as controls for later-treated states—further justifying the use of the Callaway and Sant'Anna (2021) estimator that avoids these problematic contrasts.

Table 4: Robustness: Alternative Specifications

	ATT	SE	States
Baseline ($\geq$ \$0.25)	-0.1207	0.1082	49
Threshold $\geq$ \$0.10	-0.1335	0.1277	49
Threshold $\geq$ \$0.50	0.0177	0.0405	38
Threshold $\geq$ \$1.00	0.0288	0.0370	20
Excl. CA, NY, IL	-0.1552	0.1079	46
2001–2019 only	-0.1207	0.1048	51

Notes: All specifications use Callaway and Sant'Anna (2021) with not-yet-treated controls. Dependent variable: total per capita ethanol consumption (gallons, age 21+). Baseline defines treatment as first increase $\geq$ \$0.25/pack. Alternative rows vary the threshold or sample.

6. Discussion

6.1 Welfare Implications

The central welfare question is whether the cigarette tax should be adjusted for cross-substance spillovers. In a multi-good Pigouvian framework, the optimal cigarette tax equals the marginal external cost of smoking *plus* an adjustment for cross-market effects: $\tau^*_{\text{cig}} = e_{\text{cig}} + \eta_{\text{cross}} \cdot e_{\text{alc}}$, where e_{cig} and e_{alc} are the per-unit external costs of smoking and drinking, respectively, and η_{cross} is the cross-substance elasticity (Gruber and Köszegi, 2005). If alcohol is a complement ($\eta_{\text{cross}} < 0$), the optimal cigarette tax should be *higher* than the single-substance formula suggests, because the tax also reduces alcohol-related externalities.

To translate my estimates into policy-relevant magnitudes, I compute the implied cross-price elasticity. The average cigarette tax increase in the sample is approximately \$0.60 per pack, representing roughly a 12% increase in the retail price. The point estimate of -0.121 gallons (a 4.5% decline in total ethanol) implies a cross-price elasticity of approximately -0.38 ($= -4.5\%/12\%$). The 95% confidence interval spans from -1.03 to $+0.28$, encompassing both economically meaningful complementarity and modest substitution but ruling out large positive cross-elasticities.

Using Bouchery et al.'s (2011) estimate of \$2.05 per drink in external alcohol costs and approximately 500 drinks per gallon of ethanol, I can bound the welfare-relevant spillover. At the point estimate, the alcohol externality reduction per capita is $0.121 \times 500 \times \$2.05 = \$124$. At the upper bound of the confidence interval (0.333 gallons), it reaches \$341 per capita. At the lower bound ($+0.091$ gallons of *increased* consumption), the spillover *adds* \$93 in alcohol externality costs. The confidence interval thus encompasses both a meaningful welfare bonus and a meaningful welfare penalty, preventing a definitive correction to the Pigouvian formula. The single-substance benchmark remains the pragmatic choice—not because cross-spillovers are known to be zero, but because their sign and magnitude remain too uncertain to improve upon the simpler formula.

6.2 Interpretation of the Near-Null

Three candidate explanations account for the imprecise null result. First, the cross-substance elasticity may genuinely be near zero at the population level—cigarettes and alcohol serve sufficiently different consumption motives that tax-induced changes in one do not spill over to the other. Diamond and Sekhon (2009) argue that addictive substances interact primarily through income effects, which are small for moderate tax changes.

Second, complementarity and substitution may coexist in the population, with the aggregate effect masking offsetting responses. Heavy co-users (for whom the substances are complements) may reduce both, while budget-constrained light smokers (for whom they are substitutes) may shift spending toward alcohol. State-level data cannot distinguish these margins, and the net effect may cancel to near zero. Individual-level data from the Behavioral Risk Factor Surveillance System (BRFSS) could potentially decompose these heterogeneous responses, though such data are not available at the annual state level with sufficient precision.

Third, the measurement may be too coarse to detect real effects. State-level apparent consumption, derived from sales and tax data, reflects where alcohol is purchased rather than consumed. Border effects, informal markets, and measurement error in the NIAAA data all contribute noise that attenuates estimated treatment effects (Kenkel, 2005). Classical measurement error in the outcome would bias the ATT toward zero without inflating standard errors, meaning the true effect could be larger than the point estimate suggests. However, measurement error also limits the design's statistical power: with 51 states and 25 years, the minimum detectable effect at 80% power (assuming the observed standard errors) is approximately 0.21 gallons, or 7.8% of the mean—roughly double the point estimate. The design is thus powered to detect moderate-to-large spillovers but may miss economically meaningful effects in the range of 2–5% of consumption.

6.3 Comparison with Prior Work

My results are broadly consistent with Decker and Schwartz (2000), who found weak evidence of complementarity using 1990s data and two-stage least squares. Their preferred specification yielded a cross-price elasticity of approximately -0.1, comparable in sign and magnitude to the implied elasticity from my point estimate. The consistency across methods (2SLS vs. staggered DiD), time periods (1990s vs. 2000s–2010s), and data sources provides convergent evidence that the cross-substance relationship, if it exists, is small.

My findings contrast with Yu et al. (2023), who found a statistically significant reduction in alcohol expenditure following South Korea's 2015 cigarette tax increase. The difference likely reflects the identification setting: South Korea's single national tax change cannot be separated from contemporaneous trends using only within-country variation, while my staggered US design leverages cross-state timing differences. The Korean result may also reflect institutional differences in drinking culture and tax salience (Chetty et al., 2009).

7. Conclusion

State cigarette excise taxes do not generate large spillovers to alcohol markets. Point estimates suggest modest complementarity—a pattern consistent with joint consumption in social settings—but the effects are too imprecisely estimated to distinguish from zero. The practical implication is reassuring for policymakers: the single-substance Pigouvian framework provides an adequate approximation for setting cigarette tax rates, without requiring a complex cross-market correction.

This conclusion comes with a caveat. State-level aggregate data may lack the power to detect economically meaningful effects at the individual level. A natural extension would use individual-level panel data to decompose responses by smoking intensity, income, and co-use patterns. The question of whether sin taxes in one market affect behavior in adjacent markets remains an open frontier for both public economics and behavioral health—one where the answer, even if near zero in the aggregate, has large stakes for optimal policy design.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @olafdrw
First Contributor: https://github.com/olafdrw

References

Adda, Jérôme and Francesca Cornaglia, "Taxes, Cigarette Consumption, and Smoking Intensity," *American Economic Review*, 2006, *96* (4), 1013–1028.

Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky, "Regressive Sin Taxes, with an Application to the Optimal Soda Tax," *Quarterly Journal of Economics*, 2019, *134* (3), 1557–1626.

Bouchery, Ellen E., Henrick J. Harwood, Jeffrey J. Sacks, Carol J. Simon, and Robert D. Brewer, "Economic Costs of Excessive Alcohol Consumption in the U.S., 2006," *American Journal of Preventive Medicine*, 2011, *41* (5), 516–524.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Carpenter, Christopher, "Heavy Alcohol Use and Crime: Evidence from Underage Drunk-Driving Laws," *Journal of Law and Economics*, 2007, *50* (3), 539–557.

Chaloupka, Frank J. and Kenneth E. Warner, "The Economics of Smoking," *Handbook of Health Economics*, 2000, *1*, 1539–1627.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Cook, Michael B., Sanford M. Dawsey, Neal D. Freedman, and Peter D. Inskip, "The Relationship of Cigarette Smoking and Alcohol Use to Cancer Incidence," *Cancer Epidemiology, Biomarkers & Prevention*, 2007, *16* (4), 663–673.

Cook, Philip J., "Paying the Tab: The Costs and Benefits of Alcohol Control," *Princeton University Press*, 2007.

Cotti, Chad, Erik Nesson, and Nathan Tefft, "The Effect of Smoking Bans on Costs of Alcohol-Related Hospitalizations," *Health Economics*, 2016, *25* (3), 373–390.

Decker, Sandra L. and Amy Ellen Schwartz, "Cigarettes and Alcohol: Substitutes or Complements?," Working Paper 7535, National Bureau of Economic Research 2000.

Dee, Thomas S., "State Alcohol Policies, Teen Drinking and Traffic Fatalities," *Journal of Public Economics*, 1999, *72* (2), 289–315.

Diamond, Alexis and Jasjeet S. Sekhon, "The Identification and Estimation of Simultaneous Effects among Addictive Commodities," *The Review of Economics and Statistics*, 2009, *91* (2), 375–396.

Gallet, Craig A., "The Demand for Alcohol: A Meta-Analysis of Elasticities," *Australian Journal of Agricultural and Resource Economics*, 2007, *51* (2), 121–135.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gruber, Jonathan, "Tobacco at the Crossroads: The Past and Future of Smoking Regulation in the United States," *Journal of Economic Perspectives*, 2001, *15* (2), 193–212.

_ **and Botond Köszegi**, "Sin Taxes," *Journal of Public Economics*, 2005.

Harding, Matthew, Ephraim Leibtag, and Michael F. Lovenheim, "The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data," *American Economic Journal: Economic Policy*, 2012, *4* (4), 169–198.

Kenkel, Donald S., "Are Alcohol Tax Hikes Fully Passed Through to Prices? Evidence from Alaska," *American Economic Review*, 2005, *95* (2), 273–277.

Manning, Willard G., Linda Blumberg, and Lawrence H. Moulton, "The Demand for Alcohol: The Differential Response to Price," *Journal of Health Economics*, 1995, *14* (2), 123–148.

Markowitz, Sara and Michael Grossman, "Alcohol Regulation and Domestic Violence Towards Children," *Contemporary Economic Policy*, 2005, *23* (1), 103–117.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Roth, Jonathan, "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–322.

_ **, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe**, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Saffer, Henry, "Alcohol Consumption and Alcohol Advertising Bans," *American Economic Review*, 1991, *81* (2), 95–99.

Sloan, Frank A., Jan Ostermann, Gabriel Picone, Christopher Conover, and Donald H. Taylor, "The Price of Smoking," *MIT Press*, 2004.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Yu, Shijin, Beomsoo Kim, and Byounghoon Kim, "Cigarette Tax Increase and Changes in Alcohol Consumption: A Korean Case Study," *Health Policy*, 2023, *128*, 59–65.

A. Standardized Effect Sizes

To facilitate cross-study comparison, Table 5 reports standardized effect sizes for each main outcome. The SDE expresses the treatment effect in standard deviation units of the outcome variable, enabling comparison with effects measured on different scales across studies and meta-analyses.

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Total ethanol	-0.1207	0.1082	0.593	-0.2034	0.1823	Large negative
Beer	-0.0678	0.0467	0.319	-0.2127	0.1465	Large negative
Wine	-0.0150	0.0125	0.203	-0.0739	0.0616	Moderate negative
Spirits	-0.0379	0.0435	0.252	-0.1506	0.1729	Large negative

Notes: Standardized effect sizes (SDE = $\hat{\beta}$/SD(Y)) for binary treatment. SD(Y) is the unconditional standard deviation of per capita ethanol consumption (gallons, age 21+). **Research question:** Do state cigarette excise tax increases cause cross-substance spillovers to alcohol consumption? **Treatment:** Binary — first state cigarette tax increase ≥ \$0.25/pack, 2001–2019. **Data:** NIAAA Surveillance Report #122 and CDC Tax Burden on Tobacco, 51 states, 1995–2019. **Method:** Callaway–Sant'Anna staggered DiD, not-yet-treated controls, state-clustered inference. **Sample:** 1,275 state-year observations. Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. "Null" denotes a near-zero effect size ($|\text{SDE}| < 0.005$), not a failure to reject a null hypothesis.