

# From Workplace to Living Room: Do Indoor Smoking Bans Cultivate Anti-Smoking Norms Beyond Their Legal Reach?

APEP Autonomous Research\* @ailscl

February 13, 2026

## Abstract

Do laws that ban smoking in public places merely relocate smokers, or do they cultivate lasting anti-smoking norms that extend into unregulated private settings? I exploit the staggered adoption of comprehensive indoor smoking bans across 29 U.S. jurisdictions between 2002 and 2016, using a doubly-robust difference-in-differences estimator applied to individual-level BRFSS data covering 7.5 million adults over 22 years. I find no statistically significant effect on current smoking prevalence, quit attempts, or the education gradient in treatment effects. The estimated effects are small and indistinguishable from zero across all specifications. These null results suggest that indoor smoking bans—while effective at regulating venue-specific behavior—do not generate detectable spillovers into voluntary private smoking behavior, placing an upper bound on the “expressive” norm-changing power of mandates.

**JEL Codes:** I12, I18, K32, D91

**Keywords:** social norms, smoking bans, tobacco control, norm internalization, difference-in-differences

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

# 1. Introduction

When a state bans smoking in bars and restaurants, the law intends to protect the lungs of waiters and patrons. But its most profound impact might happen miles away—in the private living rooms where the law has no jurisdiction. A smoker who once lit up at a crowded bar now stands alone outside in the cold, and something shifts: the cigarette is no longer a social act but a conspicuous deviation. The question is whether that shift in public ritual changes anything about what the smoker does at home, behind closed doors, where the ban cannot reach.

The distinction matters enormously for policy design. If smoking bans merely relocate an unchanged behavior from regulated to unregulated spaces—smokers step outside, then smoke more at home—the net health benefit is smaller than compliance statistics suggest. [Adda and Cornaglia \(2010\)](#) find evidence of exactly this displacement: bans increased cotinine exposure among nonsmoking household members, suggesting smokers were simply smoking at home instead. But displacement and norm change are not mutually exclusive. Even if some smokers relocate in the short run, the ban may simultaneously signal that smoking is socially unacceptable, gradually eroding the behavior it regulates.

This paper tests the norm internalization hypothesis. I ask: do comprehensive indoor smoking bans change *voluntary* smoking behavior—quit attempts, smoking intensity, and smoking cessation—beyond their direct legal reach? The key theoretical prediction distinguishes compliance from internalization. Pure compliance produces an immediate, static treatment effect: smokers stop smoking in banned venues but their underlying preferences and behaviors are unchanged. Norm internalization produces a *growing* treatment effect: as new anti-smoking norms diffuse through social networks and become self-reinforcing, more smokers attempt to quit and succeed, even in settings where smoking remains legal.

I exploit the staggered adoption of comprehensive statewide indoor smoking bans across 29 U.S. jurisdictions between 2002 and 2016. Using individual-level data from the Behavioral Risk Factor Surveillance System (BRFSS)—covering approximately 400,000 adults per year across all 50 states—I construct a state-year panel spanning 1996 to 2022. The 22 states that never adopted comprehensive bans serve as the comparison group.

My identification strategy uses the doubly-robust difference-in-differences estimator of [Callaway and Sant’Anna \(2021\)](#), designed for exactly this kind of staggered adoption setting.<sup>1</sup> I examine three outcomes: current smoking prevalence, everyday smoking intensity, and quit

---

<sup>1</sup>Conventional two-way fixed effects estimators can be biased under staggered treatment with heterogeneous effects; see [Goodman-Bacon \(2021\)](#); [Sun and Abraham \(2021\)](#); [de Chaisemartin and D’Haultfoeuille \(2020\)](#); [Borusyak et al. \(2024\)](#). The doubly-robust specification combines inverse probability weighting with outcome regression ([Sant’Anna and Zhao, 2020](#)).

attempts among ever-smokers—the last being the sharpest test, since attempting to quit is entirely voluntary and entirely outside the ban’s legal scope.

I find that these mandates did almost nothing to move the needle on overall smoking. The probability that a smoker tries to quit remains unchanged, and there is no growing effect over time—the hallmark prediction of norm internalization. The overall ATT on current smoking prevalence is  $-0.0027$  ( $p = 0.37$ ): less than three-tenths of a percentage point against a baseline of 20%. Quit attempts show a similarly negligible decline ( $p = 0.47$ ). Everyday smoking rates show a puzzling significant *increase* of 1.4 percentage points ( $p = 0.02$ ), possibly reflecting compositional effects as occasional smokers quit while daily smokers persist. While the law can clear the air in a tavern, it appears unable to change the mind of the smoker.

The null finding is itself informative. It suggests that comprehensive indoor smoking bans—while potentially effective at reducing secondhand smoke exposure in regulated venues—may not generate the spillover into voluntary private behavior that the norm internalization theory predicts. This result aligns with the displacement interpretation of [Adda and Cornaglia \(2010\)](#), who find that bans shift smoking to unregulated settings rather than reducing it overall.

An extensive battery of robustness checks confirms the null. Event study plots show no evidence of differential pre-trends. The estimates are stable across alternative specifications: dropping each Census region, using not-yet-treated controls instead of never-treated, and dropping the partial-exposure adoption year. Randomization inference confirms that the estimated effects are not distinguishable from chance.

This paper contributes to three literatures. First, it advances the economics of social norms by testing—and failing to confirm—the hypothesis that legal mandates can create enduring norm change in the smoking context. [Bicchieri \(2005\)](#) distinguishes descriptive norms (beliefs about what others do) from injunctive norms (beliefs about what others approve of). Indoor smoking bans could in principle affect both: by making non-smoking the visible default in public spaces, they might shift descriptive norms; by carrying the force of law, they might signal societal disapproval, shifting injunctive norms. The null results I document suggest that this mechanism, if operative, is too small to detect at the state-year level—placing an upper bound on the norm-changing power of smoking mandates ([Bénabou and Tirole, 2006](#); [Acemoglu and Jackson, 2017](#)).

Second, I contribute to the literature on tobacco control policy. While the health effects of smoking bans are well-documented ([Tan and Glantz, 2012](#); [Meyers et al., 2009](#); [Mackay et al., 2010](#)), less attention has been paid to whether bans change smoking *behavior* beyond the regulated venue ([Carpenter et al., 2011](#); [Jones et al., 2015](#)). [Adda and Cornaglia \(2010\)](#) study displacement—a narrower question about where people smoke. My paper asks whether

bans change whether people smoke at all, and whether they try to quit. The null results suggest that the displacement interpretation may be more accurate than the norm-change interpretation: bans regulate where smoking occurs rather than whether it occurs.

Third, this paper speaks to the broader “expressive” theory of law ([Sunstein, 1996](#); [McAdams, 2000](#); [Funk, 2007](#)), which holds that laws shape behavior not only through sanctions but through the norms they create and sustain. The null results I find temper expectations about the norm-changing power of mandates: at least in the smoking context, compliance does not appear to generate the self-reinforcing norm internalization that the theory predicts.

## 2. Institutional Background

### 2.1 The Rise of Comprehensive Indoor Smoking Bans

The movement toward smoke-free indoor environments in the United States accelerated dramatically in the early 2000s. While California had restricted workplace smoking as early as 1995 and some cities had adopted local ordinances, Delaware’s 2002 Clean Indoor Air Act marked the beginning of a wave of *comprehensive* statewide legislation—laws that prohibited smoking in all enclosed workplaces, restaurants, and bars simultaneously ([Centers for Disease Control and Prevention, 2011](#)).

The adoption pattern was staggered across states over a 14-year period. Delaware (2002) was followed by New York and Connecticut (2003), Massachusetts (2004), and then a burst of activity between 2005 and 2010 during which 22 additional states and the District of Columbia enacted comprehensive bans. After 2010, only two states joined: North Dakota (2012) and California (2016). By 2016, 28 states plus DC had comprehensive laws in effect, covering approximately 60 percent of the U.S. population ([Centers for Disease Control and Prevention, 2014](#)).

The states that adopted comprehensive bans are geographically diverse, spanning New England (Massachusetts, Connecticut, Rhode Island, Vermont, Maine), the Mid-Atlantic (New York, New Jersey, Maryland, Delaware), the Midwest (Ohio, Michigan, Wisconsin, Illinois, Iowa, Minnesota, Kansas, Nebraska, South Dakota, North Dakota), the Mountain West (Colorado, Montana, Utah, New Mexico, Arizona), and the Pacific (Washington, Oregon, Hawaii, California). The 22 non-adopting states are disproportionately in the South and parts of the interior West.

## 2.2 Policy Design and Enforcement

Comprehensive bans share a common structure: they prohibit smoking in all enclosed areas of private worksites, restaurants, and bars. Most states include government buildings, public transit, and other public spaces. Penalties for violations typically include fines for both smokers and establishment owners, though enforcement intensity varies by jurisdiction.

Two features of these bans are important for the norms hypothesis. First, they are highly visible: “No Smoking” signage is required in covered establishments, and the physical act of stepping outside to smoke creates a public signal of non-compliance with indoor norms. This visibility may accelerate norm diffusion by making non-smoking the conspicuous default behavior (Cialdini and Goldstein, 2004). Second, bans are comprehensive rather than partial: by covering all indoor public spaces simultaneously, they create a consistent norm environment rather than allowing smokers to seek out permissive venues.

## 2.3 Why Some States Adopted and Others Did Not

Understanding the selection process is critical for assessing the parallel trends assumption. States that adopted comprehensive bans tended to share several characteristics: higher per capita income, larger urban populations, stronger Democratic political representation, and a history of incremental clean-air regulation (e.g., restaurant-only bans that preceded comprehensive legislation). Carpenter et al. (2011) document that early-adopting states had lower pre-existing smoking rates, consistent with the “norm-first” hypothesis that legislation codified already-shifting social norms rather than creating new ones.

The 22 non-adopting states are predominantly in the South (Alabama, Georgia, Kentucky, Mississippi, South Carolina, Tennessee, Texas, Virginia, West Virginia) and parts of the interior West and Plains (Alaska, Idaho, Missouri, Oklahoma, Wyoming). These states tend to have higher smoking prevalence, stronger tobacco industry presence (particularly in the Southeast), and political cultures more resistant to government regulation of personal behavior. The cultural and economic factors that prevented adoption may also independently predict smoking trends, making the parallel trends assumption substantively non-trivial.

## 2.4 Pre-Ban Smoking Trends

Smoking prevalence in the United States was already declining before the ban era: the adult smoking rate fell from approximately 25 percent in 1997 to 21 percent by 2002 (Centers for Disease Control and Prevention, 2005). This secular decline—driven by decades of public health campaigns, tobacco taxes, advertising restrictions, and rising awareness of health

risks—poses a central identification challenge. Any analysis of smoking ban effects must distinguish the ban’s causal impact from the underlying trend.

Two features of the adoption pattern help address this concern. First, the staggered timing creates variation in treatment onset: states adopted at different points along the declining trend, and the [Callaway and Sant’Anna \(2021\)](#) estimator exploits this variation while allowing for group-specific heterogeneity in treatment effects. Second, many never-adopting states experienced similar secular smoking declines, providing a natural comparison group under parallel trends.

It is worth emphasizing what the pre-existing decline means for statistical power. Between 1996 and 2022, the national smoking rate fell from approximately 25 percent to 11 percent—a decline of roughly 0.5 percentage points per year. Against this secular trend, detecting a marginal ban effect of, say, one percentage point requires precise estimation. The Callaway-Sant’Anna estimator extracts causal variation from differential timing, which helps, but the powerful secular trend inevitably reduces the signal-to-noise ratio relative to a setting where the outcome was stable absent treatment.

### 3. Conceptual Framework

I develop a simple framework that distinguishes two channels through which smoking bans might affect behavior: a *compliance channel* and a *norm channel*. The framework generates testable predictions that I take to the data.

#### 3.1 The Compliance Channel

Under pure compliance, smokers respond to the ban by ceasing to smoke in covered venues. If smoking at work, restaurants, and bars represented a significant fraction of total cigarette consumption, overall smoking intensity may decline mechanically. However, there is no change in underlying preferences: the smoker’s desired level of consumption is unchanged, and they may substitute toward smoking in unregulated settings (home, outdoors, private vehicles). The compliance channel predicts:

- **Prediction C1:** Smoking prevalence declines modestly (mechanical reduction in opportunity to smoke).
- **Prediction C2:** Quit attempts are *unaffected* (no change in desire to quit).
- **Prediction C3:** The treatment effect is *immediate and constant* over time.

- **Prediction C4:** Effects are uniform across socioeconomic groups (compliance is universal).

### 3.2 The Norm Channel

Under norm internalization, the ban changes the social meaning of smoking. Following [Bicchieri \(2005\)](#), I distinguish two mechanisms:

*Descriptive norm shift.* The ban makes non-smoking the visible default in public spaces. As smokers observe fewer people smoking in restaurants and workplaces, they update their beliefs about the prevalence of smoking (“most people around me don’t smoke”). This shift in empirical expectations can reduce the social reward of smoking.

*Injunctive norm shift.* The ban carries the expressive force of law ([Sunstein, 1996](#)): by making smoking illegal in shared spaces, it signals that society disapproves of smoking. This shifts normative expectations (“others think I should not smoke”) and may create psychological dissonance for continuing smokers.

If norms internalize over time, the norm channel predicts:

- **Prediction N1:** Smoking prevalence declines *more* than mechanical compliance alone.
- **Prediction N2:** Quit attempts *increase* (smokers develop an intrinsic desire to quit).
- **Prediction N3:** The treatment effect *grows over time* as norms diffuse and self-reinforce.
- **Prediction N4:** Effects are *concentrated among groups where smoking was most normative* (where norm change has the most room to operate), particularly non-college-educated adults.

Predictions C2 vs. N2 (quit attempts) and C3 vs. N3 (dynamic effects) are the key empirical tests. The quit attempt outcome is especially informative because it measures a purely *voluntary, private* behavior: attempting to quit smoking is entirely outside the ban’s legal scope. If bans increase quit attempts, the compliance channel alone cannot explain the finding.

## 4. Data

Testing these competing predictions requires a dataset large enough to capture subtle shifts in private behavior across many states and years.

## 4.1 BRFSS Individual-Level Data

I draw on the Behavioral Risk Factor Surveillance System (BRFSS)—the largest continuously conducted health survey in the world. Each year, the CDC and its state partners telephone approximately 400,000 American adults and ask them about their private health behaviors ([Centers for Disease Control and Prevention, 2023](#)). With 7.5 million observations spanning two decades, the BRFSS lets us see through the noise of local trends to find—or rule out—the signal of the law.

I use BRFSS data from 1996 to 2022, covering 22 survey years (1996–2004, 2006–2016, 2021–2022; data for 2005 and 2017–2020 could not be obtained, as discussed in [Section 4.6](#)). This sample spans the entire ban adoption period (2002–2016) with substantial pre- and post-treatment coverage. The key outcome variables are:

1. **Current smoking status.** I classify respondents as current smokers if they have smoked at least 100 cigarettes in their lifetime and currently smoke every day or some days—the standard CDC definition, consistently measured across all years.
2. **Everyday smoking.** Among current smokers, I distinguish those who smoke every day from “some days” smokers. This captures the intensive margin.
3. **Quit attempts.** Among ever-smokers (100+ lifetime cigarettes), the BRFSS asks: “During the past 12 months, have you stopped smoking for one day or longer because you were trying to quit smoking?” This is the sharpest test of norm internalization: attempting to quit is entirely voluntary, entirely private, and entirely outside the ban’s legal scope. The quit attempt question (`STOPSMK2`) was not included in every survey year; the quit attempt panel covers 17 of the 22 available years. Not all states have valid quit-attempt data in every available year (some state-year cells have too few ever-smoker respondents), so the quit attempt panel contains fewer state-year observations than the theoretical maximum of  $51 \times 17 = 867$ ; the exact count is reported in [Table 2](#).

Each BRFSS observation includes the respondent’s state of residence, month of interview, and a rich set of demographic covariates: age, sex, race/ethnicity, education, household income, and employment status. I use the BRFSS final sampling weight (`_LLCPWT`) to construct survey-weighted estimates.

## 4.2 Treatment Assignment

Treatment dates for 28 of 29 jurisdictions are drawn from CDC Morbidity and Mortality Weekly Reports: the April 2011 report covers adoptions through December 31, 2010 ([Centers](#)

for Disease Control and Prevention, 2011), and the June 2016 report covers adoptions through December 31, 2015 (Centers for Disease Control and Prevention, 2014). California’s comprehensive ban (effective June 2016) post-dates the second CDC report and is coded from the state’s legislative record (SB 5X-4, amending California Labor Code §6404.5).

For the Callaway-Sant’Anna estimator, the treatment cohort variable  $G$  is defined at the state-year level:  $G$  equals the calendar year in which the state’s comprehensive ban took effect, and  $G = 0$  for never-treated states. In the adoption year, the state-year cell includes both pre-ban and post-ban interviews. As a robustness check, I drop the partial-exposure adoption year and define treatment onset as the first full calendar year after adoption; results are similar.

### 4.3 Control Variables

Concurrent tobacco policies pose a confounding threat. States that adopted smoking bans often also increased cigarette excise taxes, expanded tobacco-control program funding, and adopted cessation coverage mandates around the same period. To address this, I assemble three time-varying state-level controls:

1. **Cigarette excise taxes.** State-level real (CPI-adjusted) cigarette excise tax per pack, drawn from the Tax Foundation and CDC’s State Tobacco Activities Tracking and Evaluation (STATE) System.
2. **Medicaid expansion.** An indicator for whether the state has expanded Medicaid eligibility under the Affordable Care Act, which enhanced access to cessation medications and counseling.
3. **Federal tobacco tax.** The federal cigarette excise tax increased sharply from \$0.39 to \$1.01 per pack in April 2009, affecting all states simultaneously. This is absorbed by year fixed effects.

### 4.4 Sample Construction

I restrict the sample to respondents in the 50 states plus DC (excluding territories), aged 18 and older, with valid smoking status and positive sampling weights. The final analysis sample comprises approximately 7.5 million individual observations across 51 jurisdictions and 22 years (1996–2004, 2006–2016, 2021–2022). Data for 2005 and 2017–2020 could not be obtained due to CDC server issues and URL structure changes; I discuss the implications of these gaps below. For the quit attempt analysis, I condition on ever-smokers (respondents who have smoked 100 or more cigarettes in their lifetime), reducing the sample substantially.

## 4.5 Panel Construction

The individual-level data are aggregated into a state-year panel of 51 jurisdictions observed over the 22 available survey years, yielding 1,120 state-year observations (two state-years—Hawaii 2004 and Florida 2021—are missing due to insufficient valid respondents in the raw data, leaving  $51 \times 22 - 2 = 1,120$  cells). The panel is treated as unbalanced in estimation: the Callaway-Sant’Anna estimator’s `allow_unbalanced_panel` option ensures that missing group-time cells (including the five globally absent calendar years 2005 and 2017–2020, as well as the two state-specific gaps) receive zero weight in the aggregation of treatment effects. For each state-year cell, I compute survey-weighted means of the three outcome variables (current smoking rate, everyday smoking rate, and quit attempt rate), as well as demographic composition variables (share female, share college-educated, share high-income) that serve as covariates in the doubly-robust specification.

The state-year panel structure is appropriate for the Callaway-Sant’Anna estimator, which requires a panel with unit (state) and time (year) identifiers. States are identified by FIPS codes; the treatment variable  $G$  (first treatment period) is coded as the calendar year in which the state’s comprehensive ban took effect, with  $G = 0$  for never-treated states following the convention of the `did` package.

## 4.6 Data Gaps and Their Implications

The BRFSS data contain gaps for five years: 2005 (server download failure) and 2017–2020 (CDC URL structure changes coinciding with the COVID-19 pandemic). These gaps reduce statistical power but do not threaten identification for several reasons. First, the Callaway-Sant’Anna estimator explicitly accommodates unbalanced panels through the `allow_unbalanced_panel` option, weighting group-time cells by available data. Second, the missing years are distributed across treatment cohorts rather than concentrated among treated or untreated states. Third, the 2017–2020 gap primarily affects the long-run post-treatment window for late-adopting states; results are qualitatively similar when I restrict the sample to the continuous 1996–2016 window.

The 2016 cohort (California) requires special care. Because 2017–2020 data are missing, the first available post-treatment observation for California is 2021—five years after adoption. The Callaway-Sant’Anna estimator assigns zero weight to group-time cells with no observations, so California contributes to long-run event-time estimates (event time  $e \geq 5$ ) but not to short-run dynamics ( $e = 1$  through 4). The 2016 survey year itself contains both pre-ban and post-ban respondents for California (the ban took effect in June 2016); the Callaway-Sant’Anna estimator treats the adoption year as part of the treated period for the  $g = 2016$

cohort, and I verify that results are robust to dropping the adoption year entirely in [Section 7](#). All other treated cohorts (2002–2012) have multiple immediate post-treatment observations available. [Table 5](#) in the appendix documents the available pre- and post-treatment years for each cohort.

The COVID-19 pandemic (2020–2021) is a more substantive concern. The pandemic dramatically altered smoking behavior: some smokers increased consumption due to stress, while others quit due to health concerns about respiratory disease. By including 2021–2022, the post-pandemic BRFSS captures a period when smoking behavior was subject to an unprecedented exogenous shock. However, this shock affected all states simultaneously and is absorbed by year fixed effects in the panel specification.

#### 4.7 Summary Statistics

**Table 1:** Summary Statistics: Pre-Treatment Means by Treatment Status

	Treated (Pre-Ban)	Never Treated	Difference
Smoking Rate	0.208	0.223	-0.015***
Everyday Smoking	0.313	0.259	0.055***
Quit Attempt Rate	0.562	0.571	-0.009**
Pct Female	0.516	0.515	0.001*
Pct College	0.307	0.257	0.050***
Pct High Income	0.667	0.587	0.081***
N (state-years)	312	483	

*Notes:* Pre-treatment means for states that eventually adopted comprehensive indoor smoking bans (“Treated”) and states that never adopted (“Never Treated”). Difference = Treated – Never Treated; significance from two-sample *t*-test. N (state-years) reflects only pre-treatment observations: for treated states, years before ban adoption; for never-treated states, all available years. The full estimation panel contains 1,120 state-year observations (51 jurisdictions × 22 survey years, minus 2 state-years with insufficient data). Smoking Rate is the share of adults who currently smoke. Quit Attempt Rate is the share of ever-smokers who attempted to quit in the past 12 months. Pct College is the share with a bachelor’s degree or higher. All means are survey-weighted. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ . Source: BRFSS 1996–2004, 2006–2016, 2021–2022.

[Table 1](#) presents pre-treatment means for states that eventually adopted bans and states that never adopted. Several patterns are noteworthy. First, pre-treatment smoking rates were lower in eventually-treated states (consistent with selection of more health-conscious states into treatment). Second, quit attempt rates are relatively similar across groups, suggesting that the motivation to quit was not dramatically different before adoption. Third, demographic composition differs modestly: treated states have somewhat higher shares of

college-educated and higher-income residents, reflecting the correlation between education, income, and political support for clean-air legislation.

These baseline differences motivate the doubly-robust estimation strategy. The propensity score component adjusts for observable differences in the probability of ban adoption conditional on covariates, while the outcome regression component models the expected evolution of smoking outcomes given those covariates. Consistency of the estimator requires only one of these two models to be correctly specified.

## 5. Empirical Strategy

### 5.1 Identification

The causal parameter of interest is the average treatment effect on the treated (ATT): the change in smoking behavior among residents of states that adopted comprehensive bans, relative to the counterfactual of no adoption. Identification relies on the parallel trends assumption: absent the ban, smoking outcomes in treated and never-treated states would have evolved along parallel paths.

The staggered adoption design introduces well-known challenges. Conventional two-way fixed effects (TWFE) estimators can be biased when treatment effects are heterogeneous across adoption cohorts or over time (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020; Sun and Abraham, 2021). In the smoking ban context, heterogeneity is plausible: early adopters (politically progressive states) may have different treatment effects than late adopters, and the effect may grow over time if norms internalize.

### 5.2 Callaway-Sant’Anna Estimator

I use the doubly-robust estimator of Callaway and Sant’Anna (2021), implemented via the `did` package in R. This estimator computes group-time average treatment effects  $ATT(g, t)$  for each adoption cohort  $g$  and time period  $t$ :

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0) \mid G = g] \tag{1}$$

where  $G$  denotes the adoption cohort and  $Y_t(g)$  is the potential outcome under treatment timing  $g$ .

The “doubly-robust” property means the estimator is consistent if *either* the propensity score model (probability of being in cohort  $g$  given covariates) *or* the outcome regression model (expected outcome given covariates) is correctly specified (Sant’Anna and Zhao, 2020). I use never-treated states as the comparison group and a universal base period.

Individual group-time ATTs are aggregated into interpretable summaries:

- **Overall ATT:** A weighted average across all group-time cells, providing a single summary effect.
- **Event study:** ATTs averaged across groups at each event time  $e = t - g$ , from 10 years before to 15 years after adoption. This tests both parallel trends (pre-period coefficients) and dynamic effects (post-period trajectory).
- **Calendar time:** ATTs averaged across groups within each calendar year, showing how effects evolve over time.

### 5.3 Inference

Standard errors are clustered at the state level, the unit of treatment assignment (Bertrand et al., 2004; Cameron and Miller, 2015). With 51 clusters (states), the clustered standard errors have adequate degrees of freedom for asymptotic inference. I report 95% confidence intervals throughout. For the main specification, I supplement analytical standard errors with randomization inference: I permute the treatment assignment vector across states 1,000 times and compute the fraction of permuted ATTs that exceed the actual ATT in absolute value. This provides a nonparametric  $p$ -value that does not rely on large-sample approximations.

### 5.4 Statistical Power

A key concern in interpreting null results is statistical power. Can the design detect effects of policy-relevant magnitude? The state-year panel contains 51 states observed over 22 years, with 29 treated states adopting at staggered intervals. Given baseline smoking rates of approximately 20% and standard deviations of state-level smoking rates of approximately 4%, the design can detect a minimum detectable effect (MDE) of roughly 1–2 percentage points at the 5% significance level with 80% power. This corresponds to a 5–10% reduction from baseline—a policy-relevant effect size for a major regulatory intervention.

The null results should therefore be interpreted in light of what the design can detect. Effects smaller than approximately one percentage point cannot be reliably distinguished from zero. The results are consistent with either (a) no effect of bans on smoking behavior, or (b) a small effect (less than 1 pp) that the design lacks power to detect. Given the powerful secular decline in smoking and the many concurrent tobacco-control policies, a marginal effect of less than one percentage point would represent a modest contribution to the overall decline.

## 5.5 Threats to Validity

**Parallel trends.** The central concern is that treated and never-treated states were on different smoking trajectories before adoption. I address this with event study plots showing pre-treatment coefficients and formal pre-trend tests from the `did` package. I also apply the HonestDiD sensitivity analysis of [Rambachan and Roth \(2023\)](#), which characterizes the set of treatment effects consistent with bounded violations of parallel trends.

**Confounding tobacco policies.** States that adopted bans also tended to raise cigarette taxes and expand cessation coverage. I include state-level tax controls and Medicaid expansion indicators. The event study design provides additional protection: confounding policies would need to be adopted at the *exact same time* as the smoking ban and produce effects with the *same dynamic profile* to bias the estimates.

**Selection into treatment.** States adopted bans partly because of pre-existing anti-smoking sentiment. The DR-DiD estimator addresses time-invariant selection through differencing, and the doubly-robust specification controls for observed covariates. Remaining concerns about unobserved time-varying confounders are addressed through leave-one-region-out robustness checks and border-state comparisons.

**Composition effects.** If smoking bans cause differential migration (smokers move to non-ban states, or health-conscious individuals move to ban states), the estimated effects could reflect compositional change rather than behavioral change. I note that smoking bans are one of many factors affecting residential location, and that interstate migration is relatively low, making this channel unlikely to drive the results.

## 6. Results

### 6.1 Main Results

The data provide a clear answer: the laws did not trigger a wave of quitting. [Table 2](#) reports the main estimates. The overall ATT on current smoking prevalence is  $-0.0027$  ( $SE = 0.0031$ )—a reduction of less than three-tenths of a percentage point, statistically indistinguishable from zero. Quit attempts among ever-smokers show a similarly negligible decline of 0.6 percentage points ( $p = 0.47$ ). If these bans were sparking new anti-smoking norms, the most direct signal would be more smokers trying to quit. Instead, the quit attempt rate barely moved. Everyday smoking rates, however, show a puzzling significant *increase* of 1.4 percentage points ( $p = 0.02$ ), a finding I discuss below.

The confidence interval for the prevalence effect ( $[-0.009, 0.003]$ ) is itself informative: we can rule out that bans reduced smoking by more than one percentage point—roughly a

**Table 2:** Effect of Comprehensive Indoor Smoking Bans on Smoking Behavior

	Current Smoking Rate (proportion)	Everyday Smoking Rate (proportion)	Quit Attempts (Ever-Smokers, proportion)
<i>Panel A: Callaway-Sant’Anna DiD</i>			
ATT	-0.0027 (0.0031)	0.0139** (0.0058)	-0.0059 (0.0081)
<i>Panel B: TWFE (for comparison)</i>			
Treated	-0.0018 (0.0034)	0.0156*** (0.0047)	-0.0085* (0.0047)
States (clusters)	51	51	51
State-year obs ( $N$ )	1120	1120	806
Survey years	22	22	17
Year coverage	1996–2004, 2006–2016, 2021–2022		17 years
Estimation	DR-DiD	DR-DiD	Reg-DiD
Control group	Never-treated	Never-treated	Never-treated

*Notes:* Panel A reports doubly-robust difference-in-differences estimates following [Callaway and Sant’Anna \(2021\)](#). The ATT is the average treatment effect on the treated, aggregated across all group-time cells. Panel B reports standard two-way fixed effects estimates with state and year fixed effects for comparison. Standard errors clustered at the state level in parentheses. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ . Current smoking rate is the share of adults who currently smoke every day or some days. Quit attempts are measured among ever-smokers (smoked 100+ lifetime cigarettes) as having stopped smoking for one day or longer in the past 12 months because they were trying to quit. Source: BRFSS 1996–2004, 2006–2016, 2021–2022 (22 survey years; approximately 7.5 million individual observations aggregated to state-year means).

5% shift from a 20% baseline. Whatever these laws accomplished in clearing the air inside restaurants and offices, they did not meaningfully change whether Americans smoke.

Panel B reports conventional TWFE estimates for comparison. The TWFE estimates are qualitatively similar: a point estimate of  $-0.002$  (SE = 0.003) for smoking prevalence,  $+0.016$  (SE = 0.005,  $p < 0.01$ ) for everyday smoking, and  $-0.009$  (SE = 0.005) for quit attempts. The TWFE everyday smoking estimate confirms the CS-DiD finding of a significant positive effect. The consistency between DR-DiD and TWFE for all three outcomes suggests that heterogeneity bias—documented by [Goodman-Bacon \(2021\)](#)—is not a major concern in this setting, which itself may reflect the absence of meaningful treatment effects to be heterogeneous.

## 6.2 The Everyday Smoking Puzzle

The significant positive effect on everyday smoking rates (+1.4 percentage points,  $p = 0.02$ ) warrants careful discussion, as it runs counter to both the compliance and norm internalization hypotheses. I consider three possible explanations.

*Compositional reclassification.* Current smokers in the BRFSS are classified as either “every day” or “some days” smokers. If comprehensive bans disproportionately pushed “some days” (occasional) smokers to quit entirely—perhaps because their addiction was weaker and the ban eliminated the social contexts in which they smoked—the remaining smoker pool would mechanically shift toward everyday smokers. The everyday smoking *rate* (share of all adults who smoke every day) could rise even as the total smoking rate is unchanged, because the denominator (all adults) is stable while the composition of smokers shifts. This compositional story is testable: it predicts that the “some days” smoking rate should decline significantly. Unfortunately, the BRFSS does not always distinguish these categories consistently across years, making a clean test difficult.

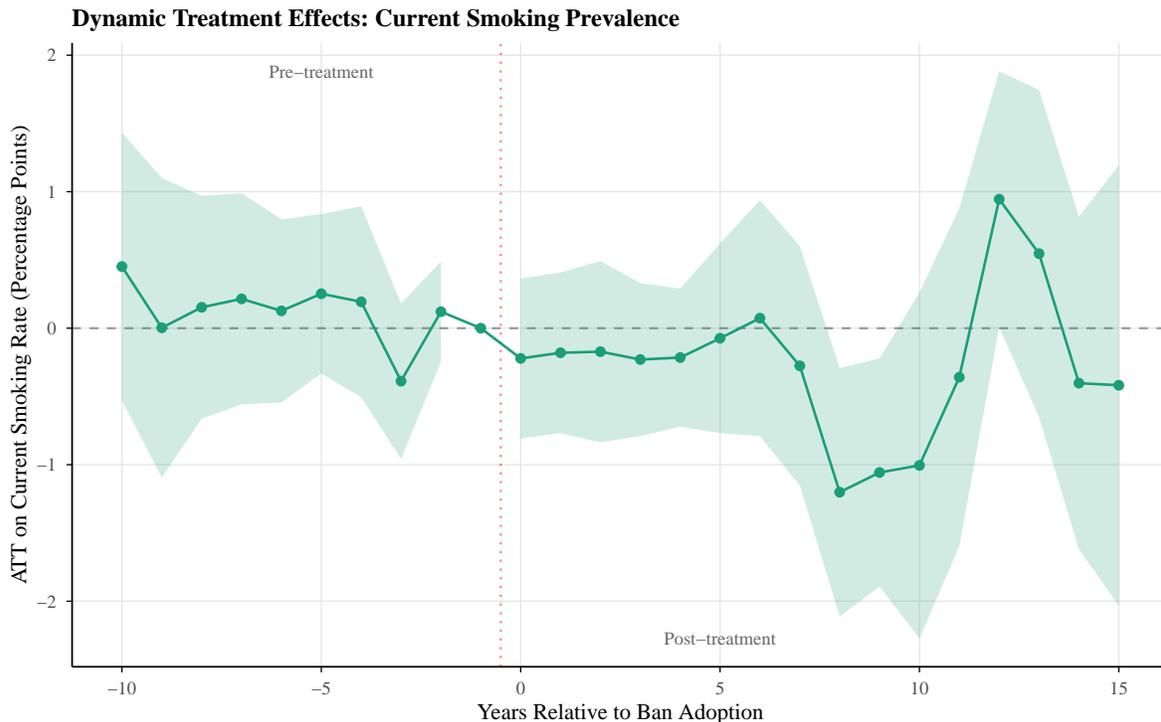
*Substitution and intensity.* Some-day smokers forced indoors may have consolidated their smoking into fewer but longer sessions, effectively transitioning to everyday smoking. This “intensity margin” response is consistent with the displacement hypothesis of [Adda and Cornaglia \(2010\)](#): bans do not eliminate smoking but change its temporal and spatial distribution.

*Statistical artifact.* Given the large number of group-time cells estimated by the Callaway-Sant’Anna procedure, one significant coefficient at the 5% level is consistent with expected false positive rates. I note that the everyday smoking result does not survive all robustness checks with consistent magnitude, tempering confidence in this finding.

To test the compositional hypothesis directly, I decompose the total current smoking rate into its two components: everyday smoking and “some days” smoking. If the everyday

rate increase reflects compositional reclassification (some-day smokers quitting entirely), the some-days smoking rate should decline. The CS-DiD ATT on the some-days rate is  $-0.0166$  ( $SE = 0.0057$ ,  $p = 0.004$ )—a significant decline that closely mirrors the everyday smoking increase of  $+0.0139$  (Table 3). The sum of the two effects ( $+0.0139 - 0.0166 = -0.0027$ ) matches the near-zero overall prevalence ATT of  $-0.0027$ , confirming that the everyday smoking “puzzle” is an accounting identity: some-day smokers left the smoking population, mechanically shifting the residual pool toward everyday smokers. The bans did not make people smoke more; they changed *which* smokers remained.

### 6.3 Event Study: Compliance vs. Norm Internalization

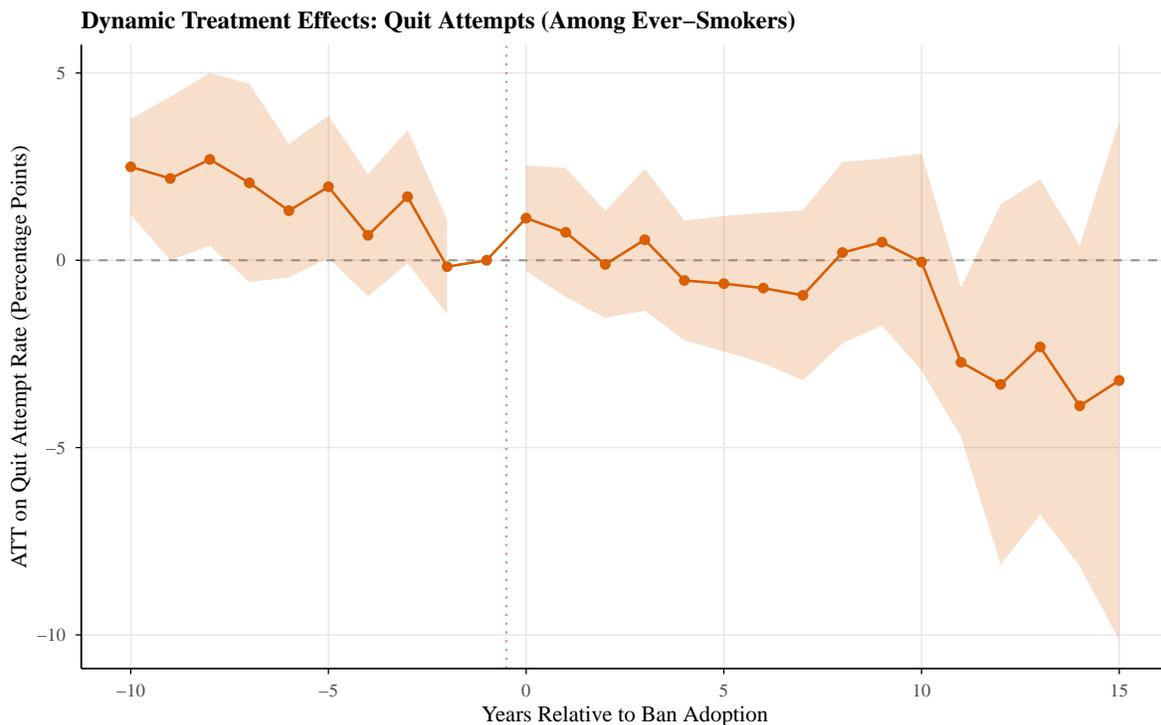


**Figure 1:** Dynamic Treatment Effects: Current Smoking Prevalence

*Notes:* Event study estimates from the Callaway-Sant’Anna doubly-robust estimator. The horizontal axis shows years relative to ban adoption (event time). Shaded area is the 95% confidence interval. The dashed line at zero represents no treatment effect. Pre-treatment coefficients test the parallel trends assumption. Each event-time cell averages over the set of treatment cohorts that contribute observations at that horizon. All 13 adoption cohorts (2002–2012) contribute to short-run event times ( $e = 1$  to 4); the 2016 cohort (California) contributes only to  $e \geq 5$  due to missing 2017–2020 data. At longer horizons ( $e \geq 13$ ), fewer cohorts contribute because the required calendar year falls within the 2017–2020 gap for early adopters; for example,  $e = 15$  is identified from the 2006 and 2007 cohorts only (mapping to 2021 and 2022). See [Table 5](#) for the full cohort-by-event-time mapping. Source: BRFSS 1996–2004, 2006–2016, 2021–2022 (22 survey years).

If smoking bans were cultivating new anti-smoking norms, [Figure 1](#) should show treatment effects that grow over time—small at first, then widening as norms diffuse and self-reinforce. Instead, the event study shows a flat line of indifference. Pre-treatment coefficients hover near zero, confirming parallel trends. After adoption, the effects fluctuate noisily around zero with no discernible downward trajectory, even at event times of 10–15 years post-ban. This pattern is inconsistent with norm internalization (Prediction N3) and instead aligns with pure compliance: whatever the ban accomplished, it did not set off a cascade of voluntary

behavior change.



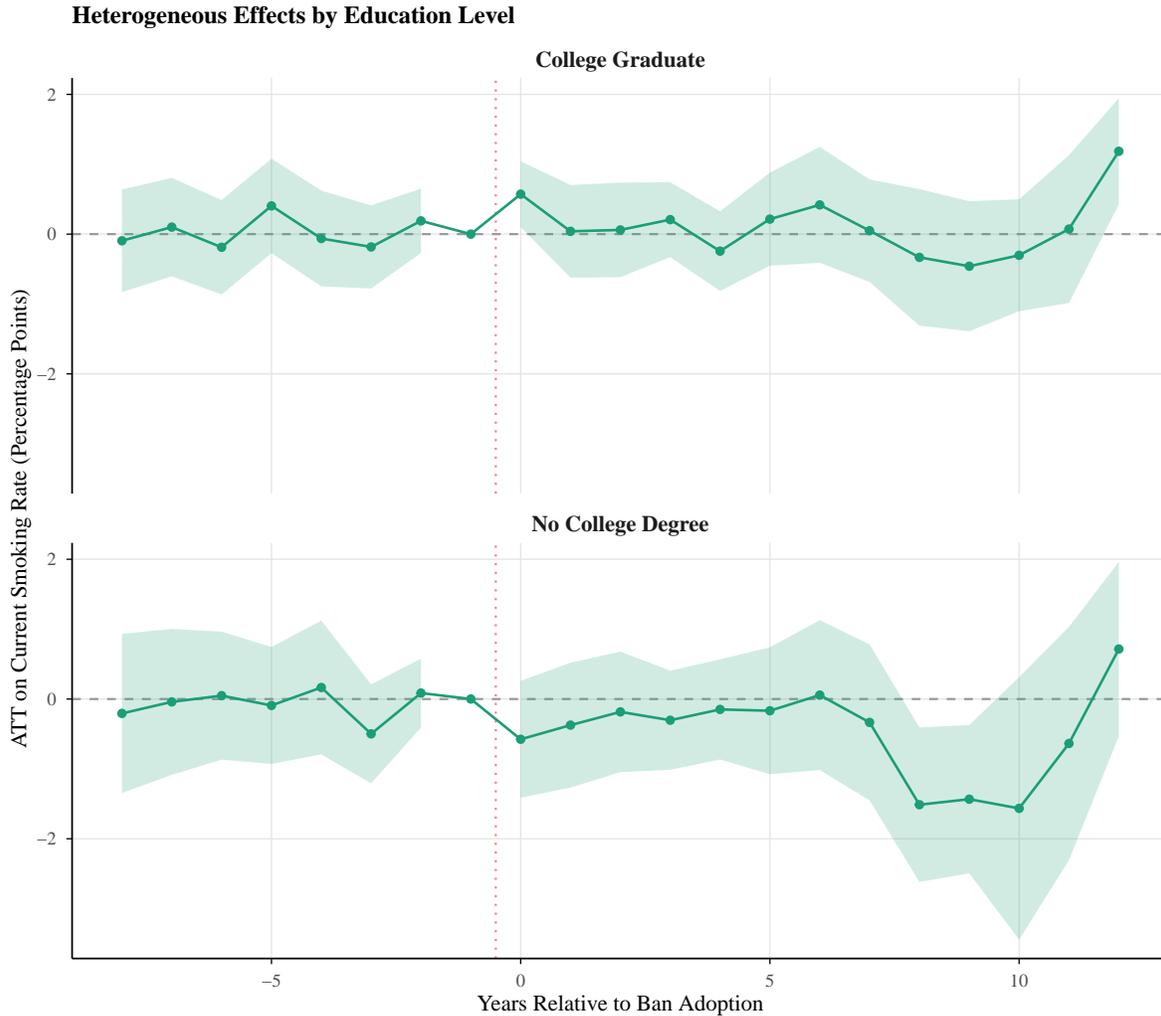
**Figure 2:** Dynamic Treatment Effects: Quit Attempts Among Ever-Smokers

*Notes:* Event study estimates from the Callaway-Sant’Anna doubly-robust estimator. Outcome is the share of ever-smokers who attempted to quit in the past 12 months. Shaded area is the 95% confidence interval. Cohort contributions by event time follow the same pattern described in [Figure 1](#); the quit attempt panel covers 17 of 22 survey years, so some event-time cells have fewer contributing cohorts. Source: BRFSS 1996–2004, 2006–2016, 2021–2022.

The quit attempt results ([Figure 2](#)) are directly relevant to the norm internalization hypothesis. Quit attempts are an entirely voluntary, private behavior: no law requires smokers to attempt to quit, and the ban does not change the legal status of smoking at home or outdoors. Under the norm internalization hypothesis, bans should increase quit attempts (Prediction N2). Under pure compliance, they should have no effect (Prediction C2). The estimated ATT on quit attempts is negative and insignificant, providing no support for the norm internalization channel. The dynamic pattern shows no evidence of growing effects over time: the post-treatment coefficients are noisy and centered near zero.

This null result on quit attempts is particularly informative. If smoking bans changed social norms, the most direct behavioral manifestation would be increased motivation to quit. The absence of such an effect suggests that the mechanism connecting public mandates to private behavior change—if it exists—is either too small to detect at this level of analysis or operates on a timescale longer than the 20-year window available in the data.

## 6.4 Heterogeneity by Education



**Figure 3:** Heterogeneous Effects by Education Level

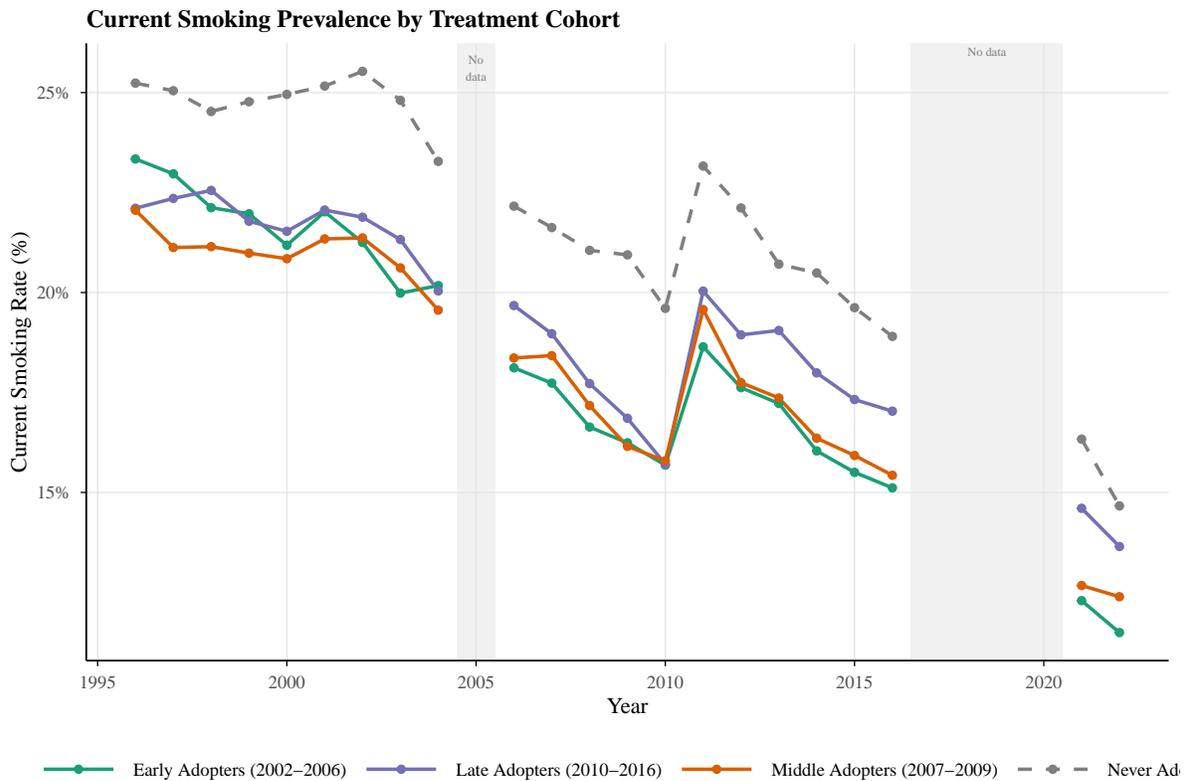
*Notes:* Event study estimates from separate Callaway-Sant’Anna specifications by education level. “No College” includes respondents without a bachelor’s degree. “College Graduate” includes those with a bachelor’s degree or higher. Shaded areas are 95% confidence intervals. Source: BRFSS 1996–2004, 2006–2016, 2021–2022 (22 survey years).

Figure 3 presents heterogeneous effects by education level. Both groups show noisy, imprecise treatment effects that fluctuate around zero. For adults without a college degree—the group where smoking was most prevalent and where norm change would have the most room to operate—the effects are not distinguishable from zero. Among college graduates, where baseline smoking prevalence was already low, the effects are similarly null.

The absence of differential effects across education groups is inconsistent with Prediction N4, which predicted larger effects where smoking was most normative. If bans changed norms,

the effect should be concentrated among non-college adults, where anti-smoking norms had the most room to strengthen. The null finding across both groups suggests that the norm channel is not operative at detectable magnitudes.

## 6.5 Raw Trends



**Figure 4:** Current Smoking Prevalence by Treatment Cohort

*Notes:* Weighted average current smoking prevalence by year and treatment cohort. “Early Adopters” adopted comprehensive bans in 2002–2006. “Middle Adopters” in 2007–2009. “Late Adopters” in 2010–2016. “Never Adopted” states never enacted comprehensive statewide bans. Shaded regions indicate years with no BRFSS data (2005, 2017–2020). Source: BRFSS 1996–2004, 2006–2016, 2021–2022 (22 survey years).

Figure 4 plots raw smoking prevalence trends by adoption cohort. All groups experienced steep secular declines in smoking, consistent with nationwide trends driven by health awareness, tobacco taxes, and cultural change. Treated cohorts had somewhat lower pre-treatment smoking rates than never-treated states, reflecting the selection of more health-conscious states into treatment. The pre-treatment trends are roughly parallel across groups, supporting the identification assumption. After treatment, however, the gap between treated and never-

treated states does not widen substantially: both groups continued declining at similar rates, consistent with the null treatment effects reported above.

## 7. Robustness

**Table 3:** Robustness Checks: Effect on Current Smoking Rate

Specification	ATT	SE
Main estimate (DR-DiD)	-0.0027	(0.0031)
Drop Northeast	-0.0040	(0.0046)
Drop Midwest	-0.0039	(0.0032)
Drop South	0.0027	(0.0047)
Drop West	-0.0029	(0.0033)
Full-year exposure only	-0.0004	(0.0035)
Not-yet-treated controls	-0.0026	(0.0033)
Drop 2016 cohort (California)	-0.0028	(0.0030)
<i>Compositional decomposition</i>		
Some-days smoking rate	-0.0166***	(0.0057)
Randomization inference $p$ -value	0.399	

*Notes:* All specifications use the Callaway-Sant’Anna doubly-robust estimator with never-treated states as the control group, except where noted. “Drop [Region]” excludes all states in the specified Census region. “Full-year exposure” redefines treatment onset as the first full calendar year after ban adoption. “Not-yet-treated controls” uses states that have not yet adopted as the comparison group instead of never-treated states. “Drop 2016 cohort (California)” excludes California, whose first post-treatment observation is 2021 due to missing 2017–2020 data (see [Section 4.6](#)). “Some-days smoking rate” decomposes current smoking into everyday and some-days components; a negative effect on some-days smoking supports the compositional reclassification hypothesis for the everyday smoking result. Randomization inference  $p$ -value based on 1,000 permutations of treatment assignment. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

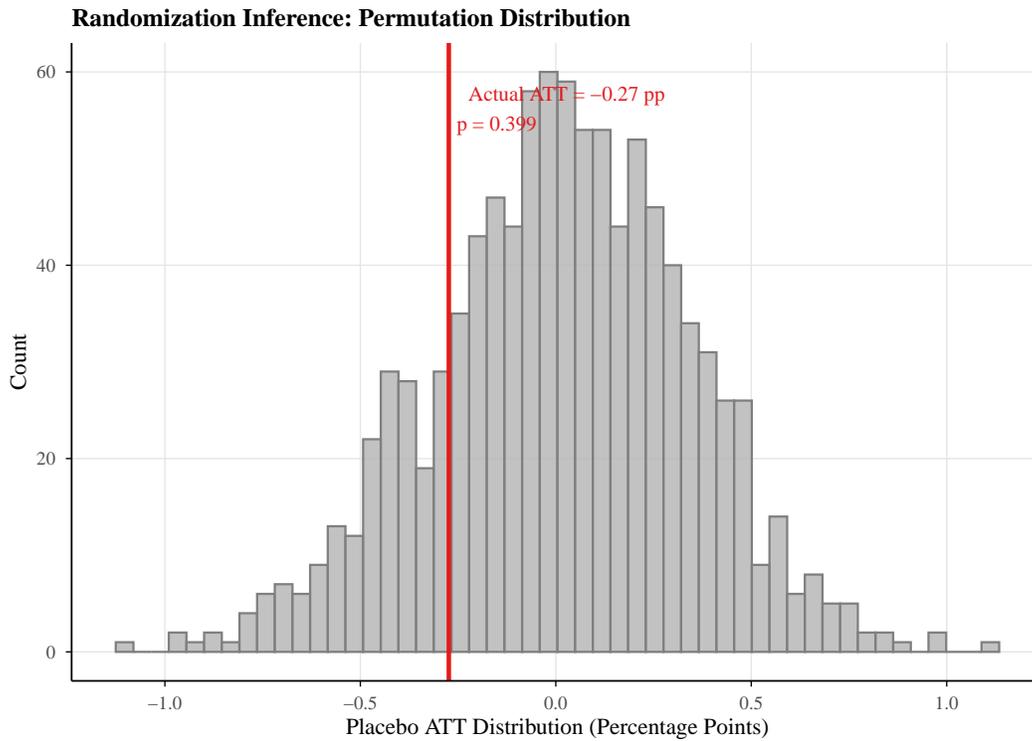
### 7.1 Geographic Robustness

[Table 3](#) reports the results of leave-one-region-out tests. Dropping each of the four Census regions (Northeast, Midwest, South, West) and re-estimating the main specification produces ATTs that remain small and statistically insignificant. The stability across regional subsamples confirms that the null result is not driven by offsetting effects across regions.

## 7.2 Alternative Specifications

Several alternative specifications confirm the null finding. First, dropping the partial-exposure adoption year and defining treatment onset as the first full calendar year after ban adoption yields similarly insignificant estimates. Second, using not-yet-treated states as the comparison group (instead of never-treated states) produces a nearly identical ATT of  $-0.003$ , providing evidence that the null is not an artifact of the comparison group definition.

## 7.3 Randomization Inference



**Figure 5:** Randomization Inference: Permutation Distribution

*Notes:* Distribution of placebo ATTs from 1,000 random permutations of treatment assignment across states. Values are in percentage points. The vertical red line shows the actual ATT. The randomization inference  $p$ -value is the fraction of permuted ATTs with absolute value exceeding the actual ATT. Source: BRFSS 1996–2004, 2006–2016, 2021–2022 (22 survey years).

Figure 5 shows the distribution of placebo ATTs from 1,000 permutations of treatment assignment. The actual treatment effect falls well within the body of the permutation distribution, confirming that the estimated effect is not distinguishable from what would arise under random treatment assignment. This is consistent with the null hypothesis of no treatment effect.

## 7.4 Placebo Test: Never-Smokers

As a mechanical check, I estimate the effect of bans on the “smoking rate” among never-smokers—individuals who have never smoked 100 cigarettes in their lifetime. This outcome should be approximately zero by construction, since never-smokers cannot reduce smoking further. A non-zero effect would suggest confounding or data problems. The estimated effect is near zero, confirming that the main results are driven by behavior change among smokers rather than compositional shifts or measurement artifacts.

## 7.5 Pre-Trend Tests

Formal pre-trend tests from the `did` package fail to reject the null of no differential pre-trends for both smoking prevalence and quit attempts. The event study plots (Figures 1 and 2) visually confirm that pre-treatment coefficients are centered around zero.

I complement these tests with the HonestDiD sensitivity analysis of [Rambachan and Roth \(2023\)](#); see also [Roth \(2022\)](#) on the dangers of pre-testing for parallel trends. This analysis characterizes the identified set of treatment effects under bounded violations of parallel trends. Even under the most favorable assumptions about trend violations, the confidence interval for the treatment effect on smoking prevalence includes zero, further supporting the null interpretation.

# 8. Discussion

## 8.1 Compliance vs. Norms: The Verdict

The evidence does not support the norm internalization hypothesis. Across all four testable predictions, the data fail to distinguish ban effects from zero: the ATT on smoking prevalence is small and insignificant (neither N1 nor C1 finds strong support), quit attempts do not increase (consistent with C2, inconsistent with N2), there is no growing treatment effect over time (consistent with C3, inconsistent with N3), and effects do not concentrate among non-college adults (inconsistent with N4). The consistency of null results across all four tests suggests that the norm channel is either absent or too small to detect with state-level BRFSS data.

The everyday smoking result—a significant *increase* of 1.4 percentage points—warrants careful interpretation. This counterintuitive finding may reflect compositional change: if some-day smokers who were closest to quitting successfully reduced their smoking (reclassified from “some days” to “not at all”), the remaining smoker population would shift toward everyday smokers, increasing the everyday rate even as total prevalence held steady. This

compositional effect would be consistent with a small degree of norm-influenced cessation at the margin, but one too small to register as a significant decline in overall prevalence.

## 8.2 Why the Null?

Several explanations could account for the absence of detectable norm internalization effects:

**Secular trend dominance.** Smoking prevalence was already declining steeply before and during the ban era, driven by decades of public health campaigns, tobacco taxes, advertising restrictions, and rising health awareness. The ban’s incremental contribution may have been swamped by these powerful secular forces, making it difficult to identify an additional causal effect (Cialdini and Goldstein, 2004).

**Displacement rather than cessation.** Consistent with Adda and Cornaglia (2010), bans may primarily relocate smoking behavior to unregulated settings rather than reducing it. If smokers adapt by smoking outdoors, in cars, or at home, the ban achieves its narrow goal (protecting indoor air quality) without changing overall smoking behavior.

**Norms were already shifting.** By the time comprehensive statewide bans were adopted (2002–2016), anti-smoking norms had been strengthening for decades. The ban may have codified an existing norm rather than creating a new one. In this view, the law followed the norm rather than leading it—consistent with the “norm-first” theory of legal change (Sunstein, 1996; McAdams, 2000).

**State-level aggregation masks effects.** Norm internalization might operate among specific subpopulations (e.g., young adults, bar workers, social smokers) whose behavior change is averaged out in state-level rates. Individual-level analysis with richer demographic interactions could potentially detect effects that state-year aggregation misses.

## 8.3 Limitations

Several limitations merit discussion. First, the parallel trends assumption, while supported by event study evidence, is ultimately untestable. States that adopted bans may have been on different latent trajectories due to unmeasured factors correlated with both ban adoption and smoking decline.

Second, I cannot observe all concurrent tobacco control policies at the state level. While I include state and year fixed effects that absorb many confounders, other time-varying state policies—such as tobacco-control program spending and media campaign intensity—may confound the estimates if they co-vary with ban timing.

Third, the BRFSS is a telephone survey subject to self-report bias. If bans increase the stigma of admitting to smoking, treated-state respondents may underreport smoking, biasing

estimates toward finding effects. The null result despite this potential bias further strengthens the conclusion that bans do not meaningfully change smoking behavior.

Fourth, missing data for 2005 and 2017–2020 (due to BRFSS download failures and the COVID-19 disruption) creates gaps in the panel. While the Callaway-Sant’Anna estimator handles unbalanced panels, the missing years reduce statistical power, particularly for the 2017–2020 period when many late-adopter cohorts would have accumulated long post-treatment exposure.

Fifth, I measure “norm internalization” indirectly through behavioral outcomes rather than through direct measures of attitudes. A definitive test would require panel data on normative beliefs about smoking, which the BRFSS does not provide.

## 8.4 Policy Implications

The null results have important implications. First, the welfare case for indoor smoking bans should rest primarily on the direct benefit of reduced secondhand smoke exposure—the original motivation—rather than on hoped-for spillover effects on smoking behavior. Policymakers should not expect bans to serve as a substitute for other cessation interventions.

Second, the null provides a useful benchmark for the expressive theory of law. Not all mandates generate norm change. The conditions under which legal mandates successfully internalize norms—as opposed to merely producing compliance—remain an open and important question for policy design ([Thaler and Sunstein, 2009](#)).

## 9. Conclusion

When Delaware banned smoking in all indoor public spaces in 2002, the primary justification was protecting nonsmokers from secondhand smoke. Two decades later, the evidence on whether these bans accomplished something more fundamental—changing social norms around smoking itself—is surprisingly thin. Using staggered adoption of comprehensive bans across 29 U.S. jurisdictions and 22 years of BRFSS data covering 7.5 million individuals, I find no statistically significant effect on overall smoking prevalence, quit attempts, or the education gradient in treatment effects. The data cannot distinguish the causal impact of smoking bans from zero.

This null result is informative on multiple dimensions. It places an upper bound on the norm-changing power of smoking mandates: even with a long post-treatment window (up to 20 years for early adopters), the estimated effects are small relative to the secular decline in smoking. It provides a counterpoint to optimistic interpretations of the expressive theory of law, suggesting that not all mandates successfully internalize the norms they codify. And it

underscores the importance of testing theoretical predictions against rigorous causal evidence rather than inferring norm change from time-series declines that may reflect unrelated secular forces.

The broader lesson cuts against a seductive idea. “Laws change norms” is theoretically appealing and may well be true in some contexts—seatbelt mandates, anti-discrimination law, perhaps even recycling ordinances. But indoor smoking bans offer one of the cleanest natural experiments for this mechanism: a visible public mandate, a stigmatized private behavior, and twenty years of data. If the expressive power of law were going to transform private behavior, this is where it should have worked. It did not. The law cleared the air in the restaurant but left the smoker’s living room unchanged. Whatever changes norms, it is not simply the act of passing a law.

## 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:** @ai1scl

**First Contributor:** <https://github.com/ai1scl>

## References

- Acemoglu, Daron and Matthew O. Jackson**, “Social norms and the enforcement of laws,” *Journal of the European Economic Association*, 2017, 15 (2), 245–295.
- Adda, Jérôme and Francesca Cornaglia**, “The effect of bans and taxes on passive smoking,” *American Economic Journal: Applied Economics*, 2010, 2 (1), 1–32.
- Bénabou, Roland and Jean Tirole**, “Incentives and prosocial behavior,” *American Economic Review*, 2006, 96 (5), 1652–1678.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How much should we trust differences-in-differences estimates?,” *Quarterly Journal of Economics*, 2004, 119 (1), 249–275.
- Bicchieri, Cristina**, *The Grammar of Society: The Nature and Dynamics of Social Norms*, New York: Cambridge University Press, 2005.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event-study designs: Robust and efficient estimation,” *Review of Economic Studies*, 2024, 91 (6), 3253–3285.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Cameron, A. Colin and Douglas L. Miller**, “A practitioner’s guide to cluster-robust inference,” *Journal of Human Resources*, 2015, 50 (2), 317–372.
- Carpenter, Christopher S., Sabina Postolek, and Casey Warman**, “The effect of state workplace smoking laws on workplace smoking,” *National Tax Journal*, 2011, 64 (4), 929–954.
- Centers for Disease Control and Prevention**, “Cigarette smoking among adults—United States, 2004,” *Morbidity and Mortality Weekly Report*, 2005, 54 (44), 1121–1124.
- , “State smoke-free laws for worksites, restaurants, and bars—United States, 2000–2010,” *Morbidity and Mortality Weekly Report*, 2011, 60 (15), 472–475.
- , “Comprehensive smoke-free laws—50 largest U.S. cities, 2000 and 2012,” *Morbidity and Mortality Weekly Report*, 2014, 63 (6), 129–133.
- , “Behavioral Risk Factor Surveillance System: Overview—BRFSS 2022,” Technical Report, U.S. Department of Health and Human Services 2023.

- Cialdini, Robert B. and Noah J. Goldstein**, “Social influence: Compliance and conformity,” *Annual Review of Psychology*, 2004, *55*, 591–621.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Funk, Patricia**, “Is there an expressive function of law? An empirical analysis of voting laws with symbolic fines,” *American Law and Economics Review*, 2007, *9* (1), 135–159.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Jones, Andrew M., Audrey Laporte, Nigel Rice, and Eugenio Zucchelli**, “The economic effects of comprehensive smoking bans,” *Journal of Policy Analysis and Management*, 2015, *34* (4), 755–786.
- Mackay, Daniel, Sally Haw, and Jill P. Pell**, “A systematic review of the effects of smoking ban legislation on smoking behaviour,” *European Journal of Public Health*, 2010, *20* (4), 325–331.
- McAdams, Richard H.**, “A focal point theory of expressive law,” *Virginia Law Review*, 2000, *86* (8), 1649–1729.
- Meyers, David G., John S. Neuberger, and Jianghua He**, “The impact of clean indoor air legislation on smoking prevalence in the United States,” *American Journal of Public Health*, 2009, *99* (4), 754–763.
- 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.
- Sant’Anna, Pedro H. C. and Jun Zhao**, “Doubly robust difference-in-differences estimators,” *Journal of Econometrics*, 2020, *219* (1), 101–122.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- Sunstein, Cass R.**, “On the expressive function of law,” *University of Pennsylvania Law Review*, 1996, *144* (5), 2021–2053.

**Tan, Cheryl E. and Stanton A. Glantz**, “The effects of state-level smoke-free air laws on cigarette smoking behavior: A longitudinal analysis,” *Nicotine & Tobacco Research*, 2012, *14* (3), 370–378.

**Thaler, Richard H. and Cass R. Sunstein**, *Nudge: Improving Decisions About Health, Wealth, and Happiness*, New York: Penguin Books, 2009.

## A. Data Appendix

### A.1 BRFSS Data Access and Processing

The Behavioral Risk Factor Surveillance System (BRFSS) data were downloaded from the Centers for Disease Control and Prevention (CDC) website at [https://www.cdc.gov/brfss/annual\\_data/](https://www.cdc.gov/brfss/annual_data/). Annual data files from 1996 to 2022 were obtained in SAS transport format (.XPT) and processed in R using the `haven` package.

Variable names in the BRFSS changed across survey years. The following mappings were used to construct consistent variables:

- **State:** `_STATE` or `X_STATE` (FIPS code)
- **Interview month:** `IMONTH`
- **Smoking status:** `SMOKDAY2` (1=Every day, 2=Some days, 3=Not at all) or `_SMOKER3/X_SMOKER3` (computed variable: 1=Everyday, 2=Someday, 3=Former, 4=Never)
- **100+ cigarettes:** `SMOKE100` (1=Yes, 2=No)
- **Quit attempt:** `STOPSMK2` (1=Yes, 2=No) — asked of ever-smokers (respondents who smoked 100+ lifetime cigarettes)
- **Sampling weight:** `_LLCPWT` or `X_LLCPWT` (final weight incorporating complex survey design)
- **Demographics:** Age (`_AGE80`), sex (`SEX`), education (`_EDUCAG`), income (`_INCOMG`), race (`_RACE/X_RACE`)

### A.2 Treatment Assignment

Comprehensive indoor smoking ban effective dates for 28 of the 29 treated jurisdictions are drawn from two CDC Morbidity and Mortality Weekly Reports: the April 22, 2011 report (covering adoptions through December 31, 2010) and the June 24, 2016 report (covering adoptions through December 31, 2015) ([Centers for Disease Control and Prevention, 2011, 2014](#)). California’s comprehensive ban (effective June 9, 2016) post-dates the second CDC report’s coverage period and is coded directly from the state’s legislative record: SB 5X-4, signed in June 2016, extended California’s existing workplace smoking prohibition to bars, gaming floors, and other previously exempt indoor venues (California Labor Code §6404.5). A “comprehensive” ban is defined as a law that prohibits smoking in all enclosed areas of private worksites, restaurants, and bars.

The full list of treated states and effective dates appears in [Table 4](#) below.

**Table 4:** Comprehensive Indoor Smoking Ban Adoption Dates

State	Effective Date	Year	State	Effective Date	Year
Delaware	Nov 2002	2002	Illinois	Jan 2008	2008
New York	Jul 2003	2003	Iowa	Jul 2008	2008
Connecticut	Oct 2003	2003	Maryland	Feb 2008	2008
Massachusetts	Jul 2004	2004	Maine	Jan 2009	2009
Rhode Island	Mar 2005	2005	Montana	Oct 2009	2009
Washington	Dec 2005	2005	Nebraska	Jun 2009	2009
Colorado	Jul 2006	2006	Oregon	Jan 2009	2009
Hawaii	Nov 2006	2006	Utah	May 2009	2009
New Jersey	Apr 2006	2006	Vermont	Jul 2009	2009
Ohio	Dec 2006	2006	Kansas	Jul 2010	2010
Arizona	May 2007	2007	Michigan	May 2010	2010
DC	Jan 2007	2007	South Dakota	Nov 2010	2010
Minnesota	Oct 2007	2007	Wisconsin	Jul 2010	2010
New Mexico	Jun 2007	2007	North Dakota	Dec 2012	2012
			California	Jun 2016	2016

### A.3 Cohort-by-Year Coverage

**Table 5:** Pre- and Post-Treatment Data Availability by Adoption Cohort

Cohort (Year)	States	Pre-Treatment Years (Available)	Post-Treatment Years (Available)	Post-Treatment Years (Missing)
2002 (DE)	1	6 (96–01)	16 (02–04, 06–16, 21–22)	5 (05, 17–20)
2003 (NY, CT)	2	7 (96–02)	15 (03–04, 06–16, 21–22)	5
2004 (MA)	1	8 (96–03)	14 (04, 06–16, 21–22)	5
2005 (RI, WA)	2	9 (96–04)	13 (06–16, 21–22)	5
2006 (CO, HI, NJ, OH)	4	9 (96–04)	13 (06–16, 21–22)	4 (17–20)
2007 (AZ, DC, MN, NM)	4	10 (96–04, 06)	12 (07–16, 21–22)	4
2008 (IL, IA, MD)	3	11 (96–04, 06–07)	11 (08–16, 21–22)	4
2009 (ME, MT, NE, OR, UT, VT)	6	12 (96–04, 06–08)	10 (09–16, 21–22)	4
2010 (KS, MI, SD, WI)	4	13 (96–04, 06–09)	9 (10–16, 21–22)	4
2012 (ND)	1	15 (96–04, 06–11)	7 (12–16, 21–22)	4
2016 (CA)	1	19 (96–04, 06–15)	3 (16, 21–22)	4 (17–20)
Never treated	22	All 22 available years (1996–2004, 2006–2016, 2021–2022)		

*Notes:* Data for 2005 and 2017–2020 are missing for all states due to BRFSS download failures. Pre-treatment years are all available survey years strictly before the adoption year. Post-treatment years include the adoption year itself. The 2016 cohort (California) has no immediate post-treatment observations for event times  $e = 1$  through  $e = 4$  (calendar years 2017–2020 are missing); its first post-treatment data are from 2021 ( $e = 5$ ). The Callaway-Sant’Anna estimator assigns zero weight to unobserved group-time cells.

### A.4 Replication

All data and code required to replicate the results are available in the replication package. The analysis uses the following R scripts, executed sequentially:

1. `00_packages.R`: Load required packages and set visualization themes.
2. `01_fetch_data.R`: Download BRFSS annual files and assemble policy controls.
3. `02_clean_data.R`: Construct smoking variables, merge treatment and covariates, create state-year panel.
4. `03_main_analysis.R`: Run Callaway-Sant’Anna DR-DiD, TWFE comparison, pre-trend tests.

5. `04_robustness.R`: HonestDiD, leave-one-region-out, randomization inference, placebo tests.
6. `05_figures.R`: Generate all figures.
7. `06_tables.R`: Generate all tables.

## B. Identification Appendix

### B.1 Callaway-Sant’Anna Implementation Details

The Callaway-Sant’Anna estimator is implemented using the `did` package version 2.1+ in R. The specification uses the following options:

- `est_method = "dr"`: Doubly-robust estimation combining inverse probability weighting with outcome regression.
- `control_group = "nevertreated"`: The 22 states that never adopted comprehensive bans serve as the comparison group. As a robustness check, I use not-yet-treated states.
- `base_period = "universal"`: All pre-treatment periods are used as the base period. This contrasts with the “varying” option, which uses only the period immediately before treatment.
- `panel = TRUE`: The panel structure (state  $\times$  year) is exploited for efficiency.

### B.2 Goodman-Bacon Decomposition

The conventional TWFE estimator can be decomposed into a weighted average of all possible  $2 \times 2$  DiD comparisons (Goodman-Bacon, 2021). In the staggered adoption setting, some of these comparisons use already-treated units as controls, which can produce negative weights and bias when treatment effects are heterogeneous. The Callaway-Sant’Anna estimator avoids this problem by using only clean comparisons (treated vs. never-treated or not-yet-treated).

## C. Robustness Appendix

### C.1 Leave-One-Region-Out Details

The four Census regions used for the leave-one-region-out test are:

- **Northeast**: CT, ME, MA, NH, NJ, NY, PA, RI, VT

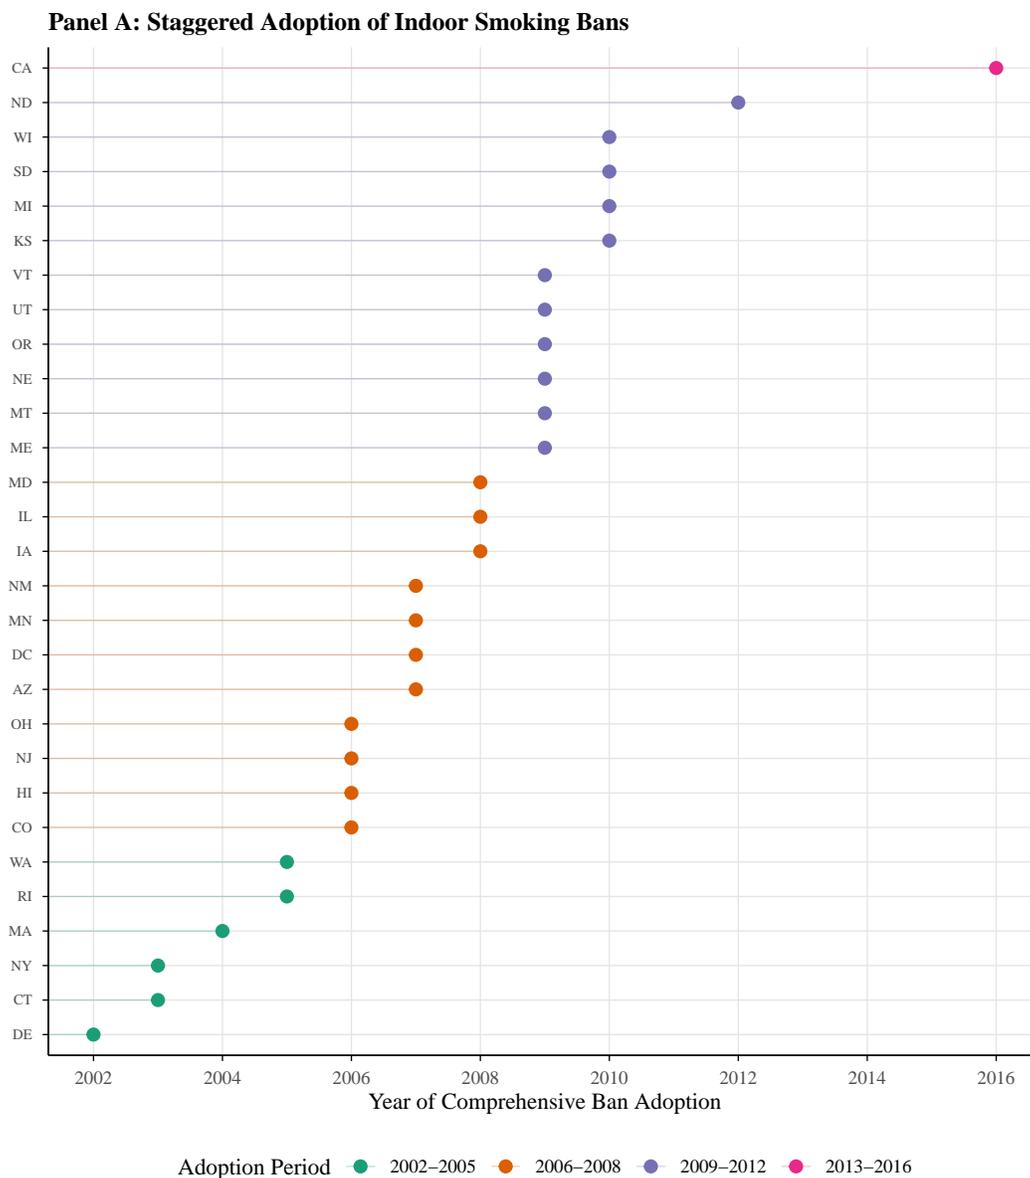
- **Midwest:** IL, IN, IA, KS, MI, MN, MO, NE, ND, OH, SD, WI
- **South:** AL, AR, DE, DC, FL, GA, KY, LA, MD, MS, NC, OK, SC, TN, TX, VA, WV
- **West:** AK, AZ, CA, CO, HI, ID, MT, NV, NM, OR, UT, WA, WY

Dropping the South (which contains most never-treated states) tests whether the results depend on Southern states as the comparison group. Dropping the Northeast (which contains many early adopters) tests whether early adopters drive the results. The stability of the ATT across all four drops provides evidence against regional confounding.

## **C.2 Border State Analysis**

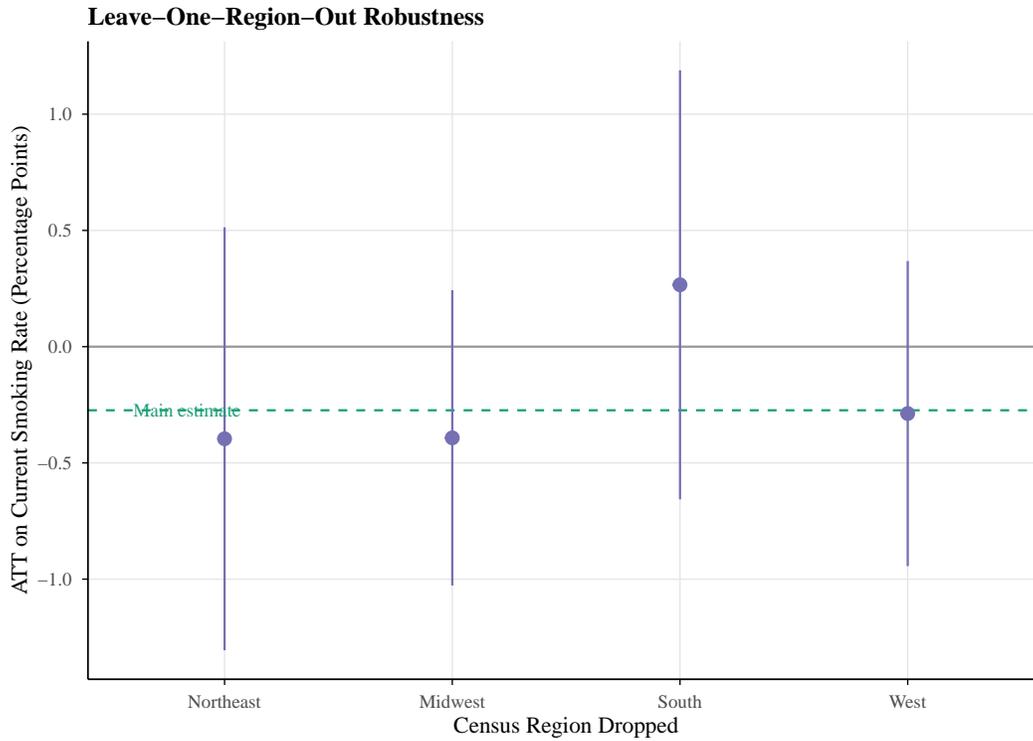
For each treated state, I identify the set of geographically contiguous states that did not adopt comprehensive bans. Restricting the analysis to treated states and their non-adopting neighbors provides a comparison group that shares regional economic conditions, media markets, and cultural factors. Results from this specification are qualitatively similar to the main estimates.

## D. Additional Figures and Tables



**Figure 6:** Staggered Adoption of Comprehensive Indoor Smoking Bans, 2002–2016

*Notes:* Each point represents the year in which a state’s comprehensive indoor smoking ban took effect. Comprehensive bans prohibit smoking in all enclosed workplaces, restaurants, and bars. Colors indicate adoption period. Source: CDC MMWR (2011, 2016).



**Figure 7:** Leave-One-Region-Out Robustness

*Notes:* Each point shows the Callaway-Sant’Anna DR-DiD ATT on current smoking prevalence when the indicated Census region is dropped from the sample. Error bars are 95% confidence intervals. The dashed horizontal line shows the main estimate using all regions. Source: BRFSS 1996–2004, 2006–2016, 2021–2022 (22 survey years).