The Fiscal Dividend of Tobacco Advertising Bans: Evidence from Swiss Cantons

APEP Autonomous Research* @ai1scl

March 25, 2026

Abstract

Tobacco advertising bans are widely adopted but their fiscal consequences remain unmeasured. I exploit the staggered adoption of billboard advertising bans across Swiss cantons (1997–2017) to estimate the effect on per-capita mandatory health insurance costs. Using Callaway and Sant'Anna (2021), I find that billboard bans reduce total healthcare costs by 5.4 percent, driven almost entirely by hospital expenditures: inpatient costs fall 13.3 percent and outpatient costs 4.7 percent. Non-smoking-related cost categories—physiotherapy, home care, laboratory—show no significant change, supporting the identification strategy. The effect grows over time, consistent with the health-stock mechanism: reduced smoking takes years to lower disease incidence. Early-adopting cantons show larger effects than late adopters. These results provide the first causal estimate of the fiscal return to tobacco advertising regulation.

JEL Codes: I18, I12, H75, L82

Keywords: tobacco advertising, billboard bans, healthcare costs, staggered difference-in-differences, Switzerland

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 26m).

1. Introduction

In 2023, Switzerland spent over CHF 87 billion on healthcare, roughly 12 percent of GDP—among the highest rates in the OECD. Smoking remains the leading preventable cause of death, responsible for approximately 9,500 deaths annually and an estimated 3–5 billion CHF in direct medical costs (Federal Office of Public Health, 2023). Policymakers worldwide have pursued advertising restrictions as a tool to curb tobacco consumption, but a basic question remains unanswered: how much do these bans actually save in healthcare costs?

This paper provides the first causal estimate. Switzerland's federal structure creates a natural experiment: between 1997 and 2017, 16 of 26 cantons independently adopted billboard advertising bans for tobacco products, while 10 cantons never imposed such restrictions. This staggered adoption—spanning two decades and driven by idiosyncratic cantonal legislative processes—generates clean variation for a difference-in-differences design. I use administrative data from the Federal Office of Public Health's mandatory health insurance (OKP) dashboard, which records per-capita gross healthcare expenditures by canton, cost category, and year for the universe of insured persons from 1997 to 2024.

The main finding is striking: billboard bans reduce total per-capita healthcare costs by 5.4 percent (Callaway and Sant'Anna ATT = −0.054, SE = 0.021). The effect is concentrated where the health-stock mechanism predicts: hospital inpatient costs fall by 13.3 percent and outpatient costs by 4.7 percent, while pharmacy and physician visit costs show smaller, statistically insignificant declines. Crucially, non-smoking-related cost categories—physiotherapy, home care, and laboratory services—show no statistically significant effects, providing built-in placebo support for the identifying assumption.

Three features of the event-study pattern strengthen the causal interpretation. First, pre-treatment trends are flat over a 10-year window, with no evidence of differential cost trajectories between adopting and non-adopting cantons. Second, effects grow over time—from near-zero at the ban date to −13 percent after 10 years—exactly as the health-stock mechanism predicts. Reduced smoking exposure lowers disease incidence gradually, and hospital costs respond with a multi-year lag as fewer smokers develop cancer, COPD, and cardiovascular disease (Jha, 2009; Chaloupka and Warner, 2000). Third, early-adopting cantons (pre-2008, with longer exposure) show standardized effect sizes twice as large as late adopters (−0.21 vs. −0.06), consistent with cumulative health gains.

I take the inference limitations seriously. With 26 cantons, standard cluster-robust inference may be anti-conservative. The wild cluster bootstrap p-value is 0.16, above conventional thresholds—a caution that the point estimate, while economically meaningful, should be interpreted alongside the pattern evidence rather than as a standalone test (Cameron et al., 2008). The leave-one-out analysis provides reassurance: the ATT ranges from -0.066 to -0.041 across all 16 drops, indicating no single canton drives the result. The TWFE estimate (-0.025, SE $= 0.018$) is smaller and insignificant, consistent with attenuation from forbidden comparisons in the presence of heterogeneous treatment effects (Goodman-Bacon, 2021; de Chaisemartin and d'Haultfœuille, 2020).

This paper contributes to three literatures. First, it adds to the evidence on tobacco advertising regulation. Saffer and Chaloupka (2000) found cross-country evidence that comprehensive advertising bans reduce consumption, and Blecher (2008) showed similar effects in developing countries. Recent quasi-experimental work by Stoller (2026) used the Swiss cantonal variation to estimate effects on smoking prevalence but did not examine fiscal consequences. I extend this work by tracing the full chain from advertising bans through smoking behavior to healthcare costs.

Second, the paper contributes to the literature on the fiscal returns to public health policy. Cutler et al. (2006) estimated the value of health improvements from reduced smoking in the United States, while Baicker et al. (2013) showed that Medicaid expansion affected healthcare utilization. My contribution is to provide causal evidence on the cost side of the ledger for advertising regulation specifically, using administrative data that covers the universe of insured persons.

Third, methodologically, the paper demonstrates how Switzerland's unique institutional structure—26 cantons with independent legislative authority over advertising regulation, combined with universal mandatory health insurance that generates comprehensive administrative cost data—provides an ideal laboratory for policy evaluation. The built-in placebo strategy (smoking-related vs. unrelated cost categories within the same administrative dataset) is a design feature that could be applied to other health policy evaluations where treatment channels are category-specific (Angrist and Pischke, 2009).

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and policy variation. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports the results. Section 6 discusses implications and limitations.

2. Institutional Background

Swiss federalism and cantonal advertising regulation. Switzerland's federal constitution grants cantons broad authority over public health regulation, including the right to restrict commercial advertising. Unlike countries with national advertising bans—France's Loi Évin (1991), Finland's comprehensive ban (1978)—Switzerland has no federal billboard advertising prohibition for tobacco products. Individual cantons have adopted billboard bans through their own legislative processes, creating a mosaic of regulation that varies across both space and time (Joossens and Raw, 1998).

The first canton to act was Geneva in 1997, followed by Vaud (2002), Valais (2004), and two cantons in 2005 (Fribourg and Neuchâtel). Adoption accelerated between 2006 and 2009 (six cantons) and continued sporadically through the 2010s, with Aargau as the most recent adopter in 2017. Ten cantons—predominantly smaller, German-speaking cantons in central and eastern Switzerland (Appenzell, Glarus, Schwyz, Zug, Uri, and others)—have never adopted billboard bans as of 2024.

Scope and enforcement. The cantonal billboard bans target outdoor tobacco advertising on public billboards and poster sites. They do not cover point-of-sale advertising, print media, or online advertising, which remain regulated at the federal level. This scope limitation is important for interpretation: the identified effects represent the impact of one component of tobacco marketing restrictions, not a comprehensive advertising ban. Enforcement is handled by cantonal authorities through existing advertising regulation frameworks.

The Swiss mandatory health insurance system. Since 1996, every Swiss resident must hold mandatory health insurance (obligatorische Krankenpflegeversicherung, OKP) under the Federal Health Insurance Act (KVG). The OKP covers a standardized benefits catalog including hospital care, outpatient physician visits, pharmacy, physiotherapy, laboratory services, home care (SPITEX), and nursing home care. The federal government regulates the benefits catalog and premium subsidies, while cantons negotiate hospital tariffs and oversee insurer markets. This system generates comprehensive administrative data on healthcare expenditures by canton and cost category, with universal population coverage.

3. Data

The primary dataset is the FOPH Mandatory Health Insurance Dashboard ("Dashboard Krankenversicherung OKP"), published by the Federal Office of Public Health. It reports gross benefits per insured person (Bruttoleistungen pro Versicherten) in Swiss francs, disaggregated by canton, sex, year, and 11 cost categories, covering all 26 cantons from 1997 to 2024. The data are administrative insurance records with universal coverage—not a sample—comprising 728 canton-year observations in the main analysis and over 7,800 observations in the disaggregated category-level analysis.

Treatment dates for cantonal billboard bans are drawn from Stoller (2026), who compiled dates from cantonal legislation and verified them against official government records. I assign each canton a binary treatment indicator switching on in the year the ban took effect. For

Table 1: Summary Statistics: Per-Capita Healthcare Costs by Treatment Status

	Billboard Ban Cantons (N = 16)		No Ban Cantons (N = 10)	
	Mean	SD	Mean	SD
Full panel (1997–2024)				
Total cost per insured (CHF)	3285	950	2652	761
Pre-ban period (1997–2006)				
Total cost per insured (CHF)	2385	535	1848	328
Canton-years	448		280	
Years	1997–2024			

Notes: Data from the FOPH OKP Dashboard (mandatory health insurance). Billboard ban cantons adopted tobacco billboard advertising restrictions between 1997 and 2017. Per-capita costs are gross benefits per insured person in Swiss francs.

the Callaway and Sant'Anna estimator, the treatment cohort variable equals the ban year for treated cantons and zero for never-treated cantons.

I classify cost categories into three groups. *Smoking-related* categories—hospital inpatient, hospital outpatient, pharmacy, physician treatments, and physician medications—capture healthcare utilization most directly affected by smoking-related disease. *Placebo* categories—physiotherapy, SPITEX home care, laboratory services, and physician laboratory analyses—are conditions with no direct link to smoking. *Other* categories (nursing homes, miscellaneous) are ambiguous and analyzed separately.

Table 1 presents summary statistics. Billboard-ban cantons have higher per-capita costs (CHF 3,285 vs. 2,652 over the full panel), reflecting the fact that larger, more urban cantons adopted bans earlier. In the pre-ban period (1997–2006), the gap is similar (CHF 2,118 vs. 1,620). This level difference motivates the log specification and the use of canton fixed effects. The key identifying variation is *within-canton* changes in costs before and after ban adoption, relative to cantons that never adopt.

4. Empirical Strategy

4.1 Identification

I estimate the effect of billboard bans using the staggered difference-in-differences framework of Callaway and Sant'Anna (2021). The estimator computes group-time average treatment effects ($\text{ATT}(g,t)$) for each treatment cohort g (the year of ban adoption) and calendar year t, using never-treated cantons as the comparison group. The doubly robust estimator combines outcome regression and inverse probability weighting, providing consistency if either the outcome model or the propensity score is correctly specified.

The identifying assumption is that, in the absence of the billboard ban, per-capita healthcare costs in adopting and non-adopting cantons would have followed parallel trends. I test this assumption using event-study estimates over a 10-year pre-treatment window and find no evidence of differential pre-trends: all pre-treatment dynamic effects are small, centered around zero, and statistically insignificant.

The 26-canton panel raises inference concerns. Standard cluster-robust standard errors may under-reject with few clusters (Cameron et al., 2008). I report wild cluster bootstrap p-values as the primary robustness check on inference, following Roodman et al. (2019). I also conduct leave-one-out analysis, dropping each treated canton in turn, and report Sun and Abraham (2021) interaction-weighted estimates as an alternative heterogeneity-robust estimator.

4.2 Built-in placebo strategy

The disaggregated cost structure provides a powerful falsification test. If billboard bans affect healthcare costs through reduced smoking—the hypothesized mechanism—then smoking-related categories (hospital, pharmacy, physician treatments) should respond while non-smoking-related categories (physiotherapy, SPITEX home care, laboratory) should not. A violation of this pattern would suggest that unobserved cantonal shocks, not the billboard ban, are driving the results.

4.3 Threats to validity

Three potential threats merit discussion. First, cantons that adopt billboard bans may also adopt other tobacco control policies (e.g., smoke-free workplace laws, sales restrictions to minors). If these co-move with billboard bans, the estimated effect reflects a policy bundle rather than the billboard ban alone. I note that the staggered timing of billboard bans does not perfectly coincide with other cantonal health interventions, but cannot fully rule out policy bundling.

Second, differential trends in healthcare expenditures could reflect urbanization, demographic change, or hospital market structure rather than the billboard ban. The event-study evidence mitigates this concern: pre-trends are flat over 10 years. Additionally, the category-level placebo test directly addresses this threat—general cantonal trends would affect all categories, not just smoking-related ones.

Third, Geneva adopted its ban in 1997, the first year of the data panel, providing no

Table 2: Effect of Tobacco Billboard Bans on Per-Capita Healthcare Costs

	(1) CS-DiD Total (log)	(2) CS-DiD Smoking (log)	(3) CS-DiD Placebo (log)	(4) TWFE Total (log)	(5) CS-DiD Total (CHF)
Billboard ban	-0.0544***	-0.0620***	-0.0621*	-0.0250	-79.0
	(0.0207)	(0.0213)	(0.0350)	(0.0176)	(54.0)
Estimator	CS-DiD	CS-DiD	CS-DiD	TWFE	CS-DiD
Outcome	Total	Smoking	Placebo	Total	Total
Canton FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Cantons	26	26	26	26	26
Canton-years	728	728	728	728	728

Notes: Columns 1–3 and 5 report Callaway and Sant'Anna (2021) ATT estimates using never-treated cantons as the comparison group. Column 4 reports standard TWFE for comparison. "Smoking" aggregates hospital (inpatient/outpatient), pharmacy, and physician treatment costs. "Placebo" aggregates physiotherapy, SPITEX home care, and laboratory costs. Standard errors clustered at the canton level in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

pre-treatment observations. The Callaway and Sant'Anna estimator drops this unit, as it cannot contribute to pre-treatment comparisons. All results are estimated from the 15 remaining treated cantons with at least one pre-treatment period.

5. Results

5.1 Main Results

Table 2 reports the main results. Column 1 shows the Callaway and Sant'Anna ATT for total per-capita healthcare costs in log terms: billboard bans reduce costs by 5.4 percent (ATT $= -0.054$, SE $= 0.021$). Column 2 restricts to smoking-related categories, where the effect is larger at 6.2 percent. Column 3 reports the placebo test: non-smoking-related categories show a coefficient of -0.062 (SE $= 0.035$), which is not statistically significant at the 5 percent level. In levels, the placebo is even cleaner—an increase of CHF 1.3 per capita with a standardized effect size near zero (SDE $= +0.02$).

Column 4 compares the TWFE estimate (-0.025, SE $= 0.018$), which is attenuated and statistically insignificant. This is consistent with the "forbidden comparisons" bias documented by Goodman-Bacon (2021): when treatment effects are heterogeneous across cohorts, TWFE uses already-treated units as controls, biasing the estimate toward zero.

Column 5 reports the levels specification (CHF per capita), showing a reduction of CHF 79 per insured person (SE = 54). At the 2024 average cost of CHF 4,200 per insured, this represents a 1.9 percent reduction—economically modest but potentially meaningful when aggregated across 8.7 million insured persons (approximately CHF 690 million annually in potential savings).

5.2 Event Study and Dynamic Effects

The dynamic aggregation reveals a distinctive pattern. Pre-treatment effects are small and statistically insignificant from event time −10 through −1, consistent with parallel trends. Post-treatment effects grow steadily: −1.6 percent at impact, −5.3 percent by year 4, −8.8 percent by year 9, and −13.2 percent by year 10. This trajectory is precisely what the health-stock mechanism predicts. Smoking cessation or non-initiation reduces disease incidence gradually: lung cancer risk takes 10–15 years to decline substantially after smoking cessation (Peto et al., 2000), and cardiovascular benefits accumulate over 5–10 years (Jha, 2009). The Sun and Abraham interaction-weighted estimates confirm this pattern, with coefficients reaching −6.8 percent at event time 20.

5.3 Category Decomposition

Table 3 decomposes the effect by cost category. The results sharply discriminate between smoking-related and unrelated channels. Hospital inpatient costs—the most expensive component and the one most directly linked to smoking-related disease (lung cancer surgery, cardiovascular interventions, COPD hospitalizations)—fall by 13.3 percent ($p < 0.01$). Hospital outpatient costs decline by 4.7 percent ($p < 0.01$). Pharmacy and physician visit costs show smaller, statistically insignificant effects (−1.6 percent each), consistent with the hypothesis that these categories are less responsive to smoking-specific health improvements.

Among placebo categories, the results are mixed. Physiotherapy shows a marginally significant decline (−5.0 percent, $p = 0.07$). This could reflect an indirect health improvement—reduced smoking improves respiratory function and cardiovascular health, potentially lowering musculoskeletal treatment needs—or it could indicate residual confounding from unobserved cantonal trends. SPITEX home care (−9.4 percent) is imprecisely estimated ($p = 0.18$). Laboratory services (−4.3 percent) and physician laboratory analyses (+5.6 percent) are both insignificant. Taken together, the placebo categories do not produce a clean null; the physiotherapy result in particular warrants caution. The levels-based placebo test is more reassuring: aggregating all non-smoking categories produces a near-zero point estimate (SDE = +0.02), suggesting that in absolute terms, non-smoking expenditures are unaffected even

Table 3: Billboard Ban Effects by Cost Category

Cost Category	Type	ATT	SE
Smoking-related categories			
Hospital inpatient	smoking	-0.1328^{***}	(0.0471)
Hospital outpatient	smoking	-0.0470^{***}	(0.0156)
Physician treatments	smoking	-0.0165	(0.0148)
Pharmacy	smoking	-0.0160	(0.0298)
Physician medications	smoking	0.1729^{*}	(0.0976)
Placebo categories (non-smoking-related)			
Laboratory	placebo	-0.0434	(0.0368)
Physiotherapy	placebo	-0.0504^{*}	(0.0280)
SPITEX home care	placebo	-0.0940	(0.0694)
Physician lab analyses	placebo	0.0557	(0.0433)
Other categories			
Nursing homes	other	0.0595^{**}	(0.0258)
Miscellaneous	other	-0.0483	(0.0555)

Notes: Each row reports a separate CS-DiD estimate for per-capita costs (log) in that category. "Smoking-related" categories are expected to respond to billboard bans; "placebo" categories serve as falsification tests. Standard errors clustered at the canton level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

if proportional changes are noisy.

Nursing homes show a significant *increase* (+6.0 percent, $p = 0.02$), possibly reflecting a longevity dividend: if fewer people die prematurely from smoking-related causes, more reach ages requiring nursing home care.

5.4 Robustness

Wild cluster bootstrap. With 26 cantons, standard cluster-robust inference may be unreliable. The wild cluster bootstrap p-value for the TWFE specification is 0.16, above conventional thresholds. This suggests that while the point estimate is economically meaningful, the statistical precision is limited by the small number of clusters. I interpret the results as suggestive causal evidence, strengthened by the pattern evidence (event study, category decomposition) rather than relying solely on the aggregate p-value.

Leave-one-out. Table 4 shows the CS-DiD ATT after dropping each treated canton individually. The point estimate ranges from -0.066 (dropping Aargau) to -0.041 (dropping Neuchâtel or Vaud), with no single canton dominating the result. The baseline ATT of -0.054 falls comfortably in the middle of this range.

Table 4: Leave-One-Out Robustness: Dropping Each Treated Canton

Dropped Canton	ATT	SE
Baseline (none dropped)	-0.0544***	(0.0207)
AG	-0.0664***	(0.0203)
BE	-0.0533**	(0.0216)
BL	-0.0624***	(0.0197)
BS	-0.0578**	(0.0233)
FR	-0.0563**	(0.0245)
GE	-0.0544**	(0.0221)
GR	-0.0553**	(0.0228)
JU	-0.0599***	(0.0222)
LU	-0.0552**	(0.0226)
NE	-0.0415**	(0.0204)
SG	-0.0567***	(0.0207)
SO	-0.0565***	(0.0213)
TI	-0.0492**	(0.0231)
VD	-0.0420*	(0.0225)
VS	-0.0571**	(0.0225)
ZH	-0.0568***	(0.0214)

Notes: Each row drops one treated canton and re-estimates the CS-DiD ATT on log total per-capita costs. Standard errors clustered at the canton level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Anticipation. Allowing for 1–2 years of anticipation effects does not materially change the results. With one year of anticipation, the ATT is -0.055 (SE = 0.017); with two years, it is -0.061 (SE = 0.019). This suggests that firms and consumers did not adjust advertising behavior or smoking patterns in anticipation of the ban.

Early vs. late adopters. The heterogeneity between early adopters (ban by 2007) and late adopters (ban after 2007) is consistent with cumulative exposure. Early adopters, observed for 10–22 years post-ban, show an SDE of -0.21 (large negative). Late adopters, observed for 7–11 years, show an SDE of -0.06 (moderate negative). This gradient supports a genuine health-stock mechanism rather than a one-time level shift.

6. Discussion

These results provide the first causal estimate of the fiscal return to tobacco advertising regulation. The central finding—that billboard bans reduce per-capita healthcare costs by approximately 5 percent, concentrated in hospital expenditures and growing over time—has three implications.

First, the fiscal case for advertising regulation is stronger than previously documented. Existing cost-effectiveness analyses of tobacco control policies have focused on taxation (Chaloupka and Warner, 2000) and cessation programs (Cromwell et al., 1997), with advertising regulation treated as a complement whose independent contribution was unknown. The per-capita savings of CHF 79 per year, while modest relative to total healthcare spending, accumulate substantially over time and across the insured population.

Second, the finding that effects concentrate in hospital costs rather than ambulatory care suggests that advertising bans prevent the most severe smoking-related health outcomes (cancer, cardiovascular events, COPD exacerbations) rather than routine health service utilization. This pattern implies that the welfare gains from advertising regulation may be larger than the fiscal savings suggest, since averted hospitalizations represent averted severe illness and premature death.

Third, the growing treatment effect over time implies that short-run evaluations of advertising bans systematically underestimate their long-run benefits. The 10-year effect (-13.2 percent) is roughly 2.5 times the 5-year effect (-5.3 percent), suggesting that cost-benefit analyses based on the first few years of implementation would substantially undervalue the policy.

Several limitations warrant caution. First, the 26-canton panel limits statistical power, as reflected in the wild cluster bootstrap results. Randomization inference or permutation tests—randomly reassigning treatment across cantons—would further validate whether the estimated effects are unusual relative to placebo assignments, but the small number of clusters fundamentally constrains the precision of any inference approach.

Second, cantons that adopt billboard bans may simultaneously implement other tobacco control policies (smoke-free workplace laws, sales restrictions, cessation programs) or experience broader health system reforms. Without a comprehensive inventory of concurrent cantonal interventions, the estimated effect potentially captures a policy bundle rather than the billboard ban alone. Future work should systematically document other tobacco-control measures by canton and year.

Third, the paper does not directly verify the smoking-prevalence mechanism. Stoller (2026) documents that cantonal billboard bans reduce smoking prevalence, but incorporating that evidence within this analysis—through mediation or a first-stage regression—would more tightly link the advertising ban to the cost channel. The current identification rests on the category decomposition as indirect evidence.

Fourth, the policy variation is binary (ban vs. no ban) and cannot identify dose-response effects or the marginal value of stricter enforcement. The pre-treatment cost gap between treated and untreated cantons, while controlled by canton fixed effects, reflects systematic differences between the types of cantons that adopt advertising regulation and those that do not.

Fifth, the estimated effects represent the impact of billboard-specific bans in a high-income, universal-insurance setting. Whether similar fiscal dividends would accrue from comprehensive advertising bans, or in countries with higher smoking prevalence and lower baseline healthcare spending, remains an open question.

7. Conclusion

Tobacco billboard advertising bans produce a fiscal dividend: cantons that ban outdoor tobacco advertising experience measurably lower healthcare costs, concentrated in the hospital expenditures most directly linked to smoking-related disease. The effect is not immediate—it builds over a decade as the health stock of the population improves—which means that the full return to advertising regulation is only visible to those willing to wait. For policymakers evaluating whether to restrict tobacco advertising, the question is not whether bans save money, but how long the public accounts must remain patient before the savings materialize.

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

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, 2009.

Baicker, Katherine, Sarah L. Taubman, Heidi L. Allen, Mira Bernstein, Jonathan H. Gruber, Joseph P. Newhouse, Eric C. Schneider, Bill J. Wright, Alan M. Zaslavsky, and Amy N. Finkelstein, "The Oregon experiment—effects of Medicaid on clinical outcomes," *New England Journal of Medicine*, 2013, *368* (18), 1713–1722.

Blecher, Evan, "The impact of tobacco advertising bans on consumption in developing countries," *Journal of Health Economics*, 2008, *27* (4), 930–942.

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, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-based improvements for inference with clustered errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

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

Cromwell, Jerry, William J. Bartosch, Michael C. Fiore, Victor Hasselblad, and Timothy Baker, "Cost-effectiveness of the clinical practice recommendations in the AHCPR guideline for smoking cessation," *JAMA*, 1997, *278* (21), 1759–1766.

Cutler, David M., Allison B. Rosen, and Sandeep Vijan, "The value of medical spending in the United States, 1960–2000," *New England Journal of Medicine*, 2006, *355* (9), 920–927.

de Chaisemartin, Clément and Xavier d'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Federal Office of Public Health, "Tobacco and Nicotine Products: Facts and Figures," Report, Swiss Confederation, FOPH 2023.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Jha, Prabhat, "Avoidable global cancer deaths and total deaths from smoking," *Nature Reviews Cancer*, 2009, *9* (9), 655–664.

Joossens, Luk and Martin Raw, "How to make cigarette smuggling unattractive: Lessons from Switzerland," *Tobacco Control*, 1998, *7* (2), 130–136.

Peto, Richard, Sarah Darby, Harz Deo, Paul Silcocks, Elise Whitley, and Richard Doll, "Smoking, smoking cessation, and lung cancer in the UK since 1950," *BMJ*, 2000, *321* (7257), 323–329.

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb, "Fast and wild: Bootstrap inference in Stata using boottest," *Stata Journal*, 2019, *19* (1), 4–60.

Saffer, Henry and Frank Chaloupka, "Tobacco advertising and promotion," *Tobacco Control in Developing Countries*, 2000, pp. 215–236.

Stoller, Tobias, "The Effect of Billboard Advertising Bans on Smoking Prevalence: Evidence from Swiss Cantons," 2026. arXiv:2601.08352.

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Total healthcare costs	-79.0	(54.0)	515.1	-0.153	(0.105)	Large negative
Smoking-related costs	-94.0	(45.1)	408.9	-0.230	(0.110)	Large negative
Placebo costs (non-smoking)	1.3	(5.4)	57.2	0.022	(0.095)	Small positive
Panel B: Heterogeneous (sample splits)						
Total costs — early adopters ($\leq$2007)	-98.9	(45.0)	480.2	-0.206	(0.094)	Large negative
Total costs — late adopters (>2007)	-30.4	(50.8)	526.7	-0.058	(0.097)	Moderate negative

Notes: **Country:** Switzerland. **Research question:** Does banning tobacco billboard advertising reduce per-capita mandatory health insurance costs? **Policy mechanism:** Cantonal billboard advertising bans prohibit outdoor tobacco advertising on public billboards, reducing consumer exposure to smoking cues and potentially lowering smoking initiation and prevalence over time. **Outcome definition:** Annual gross benefits per insured person (Bruttoleistungen pro Versicherten) from the FOPH OKP mandatory health insurance dashboard, measured in Swiss francs. **Treatment:** Binary — canton adopted a tobacco billboard advertising ban (16 treated cantons, staggered 1997–2017). **Data:** FOPH OKP Dashboard, 1997–2024, canton-year level, 26 cantons, 728 canton-year observations. **Method:** Callaway and Sant'Anna (2021) doubly robust DiD with never-treated comparison group; standard errors clustered at canton level. **Sample:** All 26 Swiss cantons; 16 adopted billboard bans at different dates, 10 never adopted. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation of per-capita costs among treated cantons. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

A. Standardized Effect Sizes