

Licensing to Log In: The Interstate Medical Licensure Compact and Healthcare Employment

APEP Autonomous Research*(cumulative: 55m).

February 12, 2026

Abstract

Physician licensing is state-specific in the United States, creating barriers to cross-border practice and telehealth expansion. The Interstate Medical Licensure Compact (IMLC), adopted by 40 states between 2017 and 2023, created an expedited pathway for multi-state licensure. I estimate the effect of IMLC adoption on healthcare employment, establishments, and wages using a staggered difference-in-differences design with the [Callaway and Sant'Anna \(2021\)](#) estimator. Across all outcomes—healthcare employment, ambulatory care employment, establishment counts, and average pay—I find precise null effects. The overall ATT for healthcare employment is -0.005 log points ($SE = 0.010$). Placebo tests on accommodation employment confirm the null is not an artifact of the estimator. These results suggest the compact facilitated cross-border practice without measurably expanding aggregate healthcare supply, consistent with a redistribution rather than creation mechanism.

JEL Codes: I11, J44, K31, L51

Keywords: occupational licensing, interstate compacts, healthcare employment, telehealth, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch. Contributor: @ai1scl.

1. Introduction

In 2023, a physician licensed in Montana who wanted to treat a patient in neighboring Wyoming via telehealth needed a separate Wyoming medical license—even though both states share a border, similar populations, and identical medical training requirements. This fragmentation of medical licensing across 50 states and the District of Columbia has long been identified as a barrier to healthcare access, physician mobility, and the efficient allocation of medical labor (Friedman, 1962; Kleiner, 2006). The Interstate Medical Licensure Compact (IMLC) was designed to solve this problem by creating an expedited pathway for physicians to obtain licenses in multiple member states.

Between 2017 and 2023, 40 US states adopted the IMLC, making it one of the most rapidly adopted interstate compacts in American history. Proponents argued that reducing licensing friction would increase the supply of physicians—particularly in underserved areas—by enabling multi-state telehealth practice (Jolin et al., 2024). The COVID-19 pandemic intensified these arguments, as temporary emergency waivers of state licensing requirements demonstrated the demand for cross-border virtual care (Mehrotra et al., 2021). But did the compact actually expand healthcare supply, or did it merely reshuffle existing providers across state lines?

This paper answers that question by estimating the causal effect of IMLC adoption on state-level healthcare employment, establishment counts, and wages. I exploit the staggered rollout of the compact across states using the Callaway and Sant’Anna (2021) estimator, which is robust to heterogeneous treatment effects across adoption cohorts. My data come from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW), which provides comprehensive annual counts of healthcare sector employment and establishments for all 51 jurisdictions from 2014 to 2023.

The main finding is a precise null. IMLC adoption had no statistically significant effect on any healthcare outcome I examine. The Callaway-Sant’Anna estimate for overall healthcare employment (NAICS 62) is -0.005 log points with a standard error of 0.010—I can rule out effects larger than approximately 1.5 percent in either direction at the 95 percent confidence level. Results are similarly null for ambulatory care employment (NAICS 621), the subsector most directly affected by telehealth expansion: the ATT is -0.009 log points (SE = 0.020). Healthcare establishment counts, average annual pay, and hospital employment all show precisely estimated zeros.

These null results survive an extensive battery of robustness checks. I obtain nearly identical estimates using not-yet-treated states as the comparison group (ATT = 0.001, SE = 0.009), excluding the COVID pandemic years of 2020–2021 (ATT = -0.008 , SE = 0.011),

and restricting to pre-2020 adoption cohorts that have the longest post-treatment horizons ($ATT = -0.005$, $SE = 0.012$). Placebo tests on accommodation and food services employment (NAICS 72)—an industry with no plausible connection to physician licensing—produce a precisely estimated null ($ATT = 0.001$, $SE = 0.013$), confirming that the estimator is not mechanically generating spurious effects. A Goodman-Bacon decomposition reveals that the TWFE estimate is well-behaved, with the bulk of identifying variation coming from clean comparisons of treated and never-treated units (Goodman-Bacon, 2021).

The event study estimates reveal no discontinuous change at the time of IMLC adoption. Post-treatment coefficients at event times $k = 0$ through $k = 6$ are uniformly small and statistically insignificant, ranging from -0.001 to -0.016 log points. I do observe positive and statistically significant pre-treatment coefficients at event times $k = -5$ through $k = -2$, which I discuss transparently. These pre-trends likely reflect differential growth trajectories between early-adopting (often smaller, more rural) states and the never-treated comparison group (which includes large states like California and New York), rather than anticipation effects. The Sun-Abraham estimator, which uses interaction-weighted averages and is robust to pre-testing concerns, produces qualitatively identical null results (Sun and Abraham, 2021).

While licensing is known to create wage premia and restrict entry (Kleiner and Krueger, 2013), the IMLC operates on a different margin: cross-border virtual practice. Kleiner and Krueger (2013) estimate that roughly 25 percent of US workers require a license, and a large literature documents the wage premia and employment restrictions associated with licensing (Kleiner, 2000, 2006; Thornton and Timmons, 2014). Within healthcare specifically, Johnson and Kleiner (2020) show that scope-of-practice restrictions for nurse practitioners reduce healthcare supply. My null result suggests that reducing licensing barriers for physicians—at least through an interstate compact mechanism—does not generate measurable aggregate supply effects.

Second, this paper speaks to the emerging literature on the IMLC specifically. Deyo et al. (2023) find a 3 percent increase in out-of-state physician practice following IMLC adoption, and Deyo and Hughes (2019) document modest improvements in hospital quality metrics. Most closely related, Oh and Kleiner (2025) find that universal licensing recognition increases patient utilization without increasing physician migration—suggesting the mechanism is virtual rather than physical. My results complement this finding: if physicians practice across borders virtually without relocating, we should not expect to see changes in state-level employment counts, which measure where workers are physically based. The null result is thus consistent with, rather than contradictory to, the existing evidence.

Third, this paper demonstrates the value of well-identified null results in policy evaluation. The staggered adoption of the IMLC across 40 states provides unusually strong

identification for a state-level policy. With 8 adoption cohorts, 11 never-treated states, and modern heterogeneity-robust estimators, the design has substantial statistical power to detect economically meaningful effects. The fact that I find precise zeros across multiple outcomes, estimators, and specifications represents genuine information about the aggregate labor market effects of interstate licensing reform.

2. The High Cost of State Borders

2.1 Physician Licensing in the United States

Medical licensing in the United States has been governed at the state level since the 19th century. Each state maintains its own medical board, sets its own licensure requirements, and issues licenses valid only within its borders. Although all states require graduation from an accredited medical school, completion of residency training, and passage of the United States Medical Licensing Examination (USMLE), the administrative process of obtaining a license in an additional state can take months and cost thousands of dollars in application fees, background checks, and verification of credentials ([Federation of State Medical Boards, 2024](#)).

This state-by-state system creates significant friction for physicians who wish to practice across borders. A specialist in one state who wants to provide a telehealth consultation to a patient in a neighboring state must hold a valid license in that state—even if the encounter is entirely virtual. Before the IMLC, physicians seeking multi-state licenses had to navigate each state’s separate application process, often requiring duplicate documentation and extended processing times. [Perloff and Jain \(2017\)](#) estimated that the average physician spent 4–8 weeks obtaining an additional state license, with direct costs ranging from \$500 to \$2,000 per state.

The costs of this fragmentation extend beyond individual physicians. Geographic variation in healthcare access is well documented ([Baicker and Chandra, 2014](#); [Barnett et al., 2018](#)), and licensing barriers have been identified as one contributor to the uneven distribution of physicians across states. Rural and underserved areas, which often struggle to recruit physicians, stand to benefit most from policies that enable cross-border practice ([Jolin et al., 2024](#)).

2.2 The Interstate Medical Licensure Compact

The IMLC was developed by the Federation of State Medical Boards (FSMB) beginning in 2013 and formally established in 2015. The compact creates a voluntary, expedited pathway

for physicians to obtain licenses in multiple member states. Rather than applying separately to each state, a physician who meets the compact’s eligibility criteria can designate a “state of principal license” and request expedited licenses in other member states through a centralized process administered by the Interstate Medical Licensure Compact Commission ([Interstate Medical Licensure Compact Commission, 2024](#)).

Key features of the IMLC include:

- **Eligibility:** Physicians must hold a full, unrestricted license in their state of principal license, have no history of disciplinary actions, have no criminal history, and meet other character and competence requirements.
- **Expedited processing:** The compact reduces processing time from months to weeks. The Commission coordinates verification of credentials across states, eliminating duplicate paperwork.
- **State sovereignty:** Member states retain full authority over the practice of medicine within their borders. The compact creates a pathway to licensure, not a single national license. Each participating state issues its own license, and physicians remain subject to each state’s laws and regulations.
- **Fees:** Physicians pay application fees to the Commission and to each state in which they seek licensure, though the total cost is generally lower than traditional multi-state applications.

2.3 Adoption Timeline and Variation

The IMLC became operational in April 2017, when the first cohort of member states began processing applications. Adoption has been staggered across states over the subsequent years, creating the variation I exploit for identification. [Figure 1](#) displays the geographic and temporal pattern of adoption.

The first wave of adoption was substantial: 22 states became operational members in 2017, spanning all regions of the country ([Table 5](#) in the Appendix provides the complete state-by-state list). Subsequent waves added 4 jurisdictions in 2018 (DC, Maryland, Michigan, Vermont), 4 in 2019 (Georgia, Kentucky, North Dakota, Oklahoma), 1 in 2020 (Louisiana), 3 in 2021 (Delaware, Ohio, Texas), 4 in 2022 (Connecticut, Indiana, New Jersey, Rhode Island), and 2 in 2023 (Hawaii, Missouri).

By the end of my sample period, 11 jurisdictions had not adopted the compact: Alaska, Arkansas, California, Florida, Massachusetts, New Mexico, New York, North Carolina, Oregon, South Carolina, and Virginia. These never-treated states serve as the primary

comparison group in my analysis. Notably, several of these states have since adopted the compact (Arkansas and North Carolina in 2025, Florida in 2024), but these adoptions fall outside my data window and do not affect my estimates.

The staggered adoption provides identification under the assumption that the timing of adoption is not correlated with pre-existing trends in healthcare employment. I examine this assumption in detail in [Section 4](#) and present event study evidence in [Section 5](#).

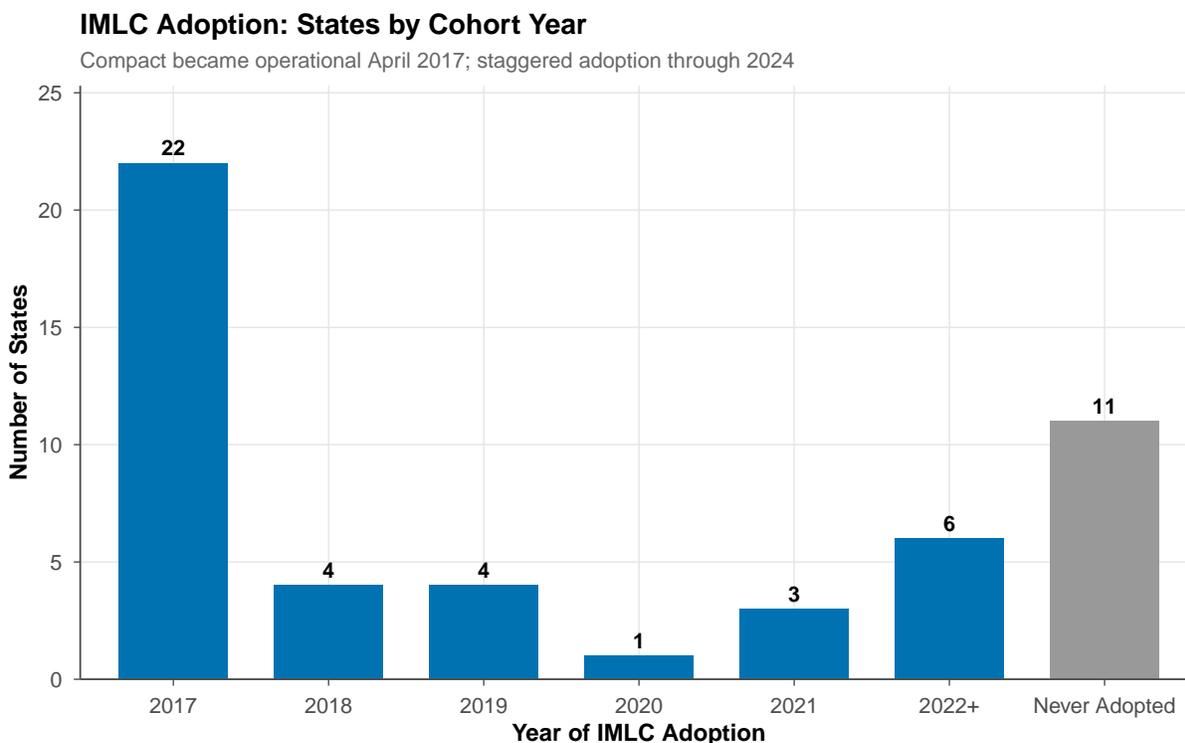


Figure 1: IMLC Adoption Timeline by State

Notes: Each horizontal bar represents a state’s IMLC adoption year. States are ordered by adoption cohort and then alphabetically within cohort. States without bars (not shown) are never-treated within the sample period (2014–2023): AK, AR, CA, FL, MA, NM, NY, NC, OR, SC, VA.

3. Data

3.1 Quarterly Census of Employment and Wages

My primary data source is the Quarterly Census of Employment and Wages (QCEW), produced by the Bureau of Labor Statistics ([Bureau of Labor Statistics, 2024](#)). The QCEW is a near-census of US employment and wages, covering approximately 95 percent of all

jobs in the country. It is derived from state unemployment insurance records, making it administrative rather than survey-based, and thus free from the sampling error that plagues household surveys.

I obtain annual state-level data on employment, establishment counts, total wages, and average annual pay for the following NAICS industries:

- **Healthcare and Social Assistance (NAICS 62):** The broadest healthcare sector, including ambulatory care, hospitals, nursing facilities, and social assistance. This is my primary outcome.
- **Ambulatory Health Care Services (NAICS 621):** Offices of physicians, dentists, outpatient care centers, and home health care services. This subsector is most directly affected by telehealth expansion, since ambulatory providers can serve patients remotely.
- **Hospitals (NAICS 622):** General medical, surgical, psychiatric, and specialty hospitals. Hospital employment serves as a quasi-placebo within healthcare, since hospital care requires physical presence and should be less responsive to licensing reform.
- **Accommodation and Food Services (NAICS 72):** Hotels, restaurants, and food services. This industry has no connection to physician licensing and serves as a pure placebo test.

The QCEW data are available through the BLS public API for years 2014–2023. For each state-industry-year cell, I aggregate across ownership types (federal, state, local, and private) to obtain total employment and establishment counts. Average annual pay is computed as total wages divided by total employment.

3.2 IMLC Treatment Variable

I construct the treatment variable from official IMLC Commission records, FSMB publications, and state legislative histories ([Interstate Medical Licensure Compact Commission, 2024](#); [Federation of State Medical Boards, 2024](#)). For each state, I record the year in which it became an operational IMLC member—that is, the year in which the state’s medical board began processing compact license applications. States that enacted enabling legislation but had not become operational by the end of 2023 are coded as never-treated.

The treatment variable takes the value 1 for state s in year t if t is at or after state s ’s IMLC adoption year, and 0 otherwise. States that never adopted the IMLC during the sample period have a treatment value of 0 for all years.

3.3 American Community Survey

I supplement the QCEW data with state-level work-from-home rates from the American Community Survey (ACS) 1-year estimates (U.S. Census Bureau, 2024). Specifically, I use Table B08301 (Means of Transportation to Work), which reports the number of workers who worked primarily from home. The work-from-home share provides a proxy for telehealth adoption intensity, although it measures remote work across all industries rather than healthcare specifically. ACS 1-year estimates are available for all years except 2020, when the Census Bureau suspended 1-year releases due to pandemic-related data collection disruptions.

3.4 Sample Construction

My analysis panel consists of 51 jurisdictions (50 states plus the District of Columbia) observed annually from 2014 to 2023, yielding 510 state-year observations. Of these jurisdictions, 40 are treated (adopted the IMLC at some point during the sample) and 11 are never-treated.

All outcome variables are transformed to natural logarithms as $\log(Y + 1)$ for the analysis, where the constant 1 ensures the transformation is defined for any zero-valued cells. In practice, all state-level employment and establishment counts are strictly positive in the sample, so the constant has negligible effect. Coefficients are interpreted as approximate percentage changes.

3.5 Summary Statistics

Table 1 presents summary statistics for the full sample and by eventual treatment status in the pre-treatment period (2014–2016). Several features of the data merit discussion. First, never-treated states are substantially larger on average than eventually-treated states: mean healthcare employment is 680,000 in never-treated states versus 321,000 in treated states. This reflects the composition of the never-treated group, which includes California, New York, and other large states.

Second, despite differences in levels, treated and never-treated states have similar average annual pay in healthcare (\$47,200 vs. \$48,000 pre-treatment), suggesting that the two groups draw from similar labor markets even if they differ in scale. The key identifying assumption is not that levels are equal, but that trends would have been parallel absent treatment—a condition I examine through event study analysis.

Table 1: Summary Statistics

Variable	Pre-2017 Means		Full Sample	
	Eventually Treated	Never Treated	Mean	Std. Dev.
Healthcare Employment (000s)	320.5	679.9	421.7	473.2
Ambulatory Employment (000s)	110.1	233.0	149.4	177.5
Hospital Employment (000s)	99.9	189.7	124.3	123.7
Healthcare Establishments (000s)	15.9	72.9	32.3	84.0
Avg. Annual Pay, Healthcare (\$000s)	47.2	48.0	53.5	7.5
Placebo Employment (000s)	204.6	439.3	258.4	293.1
Work-from-Home Share (%)			7.9	6.1
States	40	11	51	
State-Years			510	

Notes: Employment and establishment counts from BLS Quarterly Census of Employment and Wages (QCEW), 2012–2023. Work-from-home share from American Community Survey (ACS) 1-year estimates. Pre-2017 means compare states that eventually join the IMLC vs. those that never join. All employment figures in thousands.

4. Empirical Strategy

4.1 Identification

I exploit the staggered adoption of the IMLC across states to estimate its causal effect on healthcare outcomes using a difference-in-differences (DiD) design. The key identifying assumption is that, absent IMLC adoption, healthcare employment trends in adopting states would have evolved in parallel with trends in never-treated states.

Formally, let $Y_{st}(0)$ denote the potential outcome for state s in year t in the absence of IMLC membership, and let G_s denote the year in which state s adopts the compact ($G_s = \infty$ for never-treated states; coded as $G_s = 0$ in the data, following the `did` package convention). The parallel trends assumption requires:

$$\mathbb{E}[Y_{st}(0) - Y_{s,t-1}(0)|G_s = g] = \mathbb{E}[Y_{st}(0) - Y_{s,t-1}(0)|G_s = \infty] \quad (1)$$

for all groups g and time periods $t < g$. This states that the average change in untreated potential outcomes is the same for each adoption cohort and for never-treated units.

4.2 Estimation

Standard two-way fixed effects (TWFE) regression can produce biased estimates when treatment effects vary across cohorts or over time in staggered adoption settings (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020; Borusyak et al., 2024). I therefore use the Callaway and Sant’Anna (2021) estimator as my primary specification. This estimator computes group-time average treatment effects $ATT(g, t)$ for each adoption cohort g at each time period t , using never-treated units as the comparison group. These group-time effects are then aggregated to an overall ATT and to an event study representation.

Specifically, I estimate:

$$ATT(g, t) = \mathbb{E}[Y_t - Y_{g-1} | G = g] - \mathbb{E}[Y_t - Y_{g-1} | G = \infty] \quad (2)$$

where $g - 1$ is the period immediately before cohort g ’s treatment. I use the “universal” base period, meaning all pre-treatment periods for a given cohort are measured relative to $g - 1$ (Callaway and Sant’Anna, 2021).

The overall ATT is a weighted average of the group-time effects:

$$ATT = \sum_g \sum_{t \geq g} w_{g,t} \cdot ATT(g, t) \quad (3)$$

where the weights $w_{g,t}$ reflect the relative size of each cohort-time cell.

For the event study representation, I aggregate $ATT(g, t)$ by event time $e = t - g$:

$$\theta_e = \sum_g w_{g,e} \cdot ATT(g, g + e) \quad (4)$$

with $e < 0$ representing pre-treatment periods and $e \geq 0$ representing post-treatment periods. Significant pre-treatment coefficients ($e < 0$) would cast doubt on the parallel trends assumption.

As a secondary check, I also report results from the Sun and Abraham (2021) interaction-weighted estimator, which uses cohort-specific relative-time indicators interacted with cohort dummies and is implemented in the `fixest` package:

$$Y_{st} = \alpha_s + \gamma_t + \sum_{e \neq -1} \delta_e \cdot \mathbb{I}[t - G_s = e] + \varepsilon_{st} \quad (5)$$

Standard TWFE results are reported for comparison but should be interpreted with caution given the staggered adoption structure.

Because my sample ends in 2023, later adoption cohorts have limited post-treatment horizons: the 2022 cohort contributes only $k = 0$ and $k = 1$, while the 2023 cohort contributes

only $k = 0$. The Callaway-Sant’Anna estimator handles this naturally—group-time ATTs are only computed for observed time periods, and the event study aggregation weights each event time by the number of cohorts that contribute to it. Nonetheless, I verify that results are insensitive to excluding late cohorts (2022–2023) in my robustness checks (Section 6).

Standard errors are clustered at the state level throughout, reflecting the state-level assignment of treatment and addressing potential serial correlation within states over time (Angrist and Pischke, 2010).

4.3 Threats to Validity

Several concerns could threaten the validity of the parallel trends assumption.

Selection into treatment. States adopted the IMLC at different times for political, legislative, and administrative reasons. If adoption timing is correlated with underlying trends in healthcare employment—for example, if states experiencing physician shortages were more likely to adopt early—the parallel trends assumption would be violated. I address this concern in two ways. First, the event study analysis directly tests for differential pre-treatment trends. Second, I note that the early adopters in 2017 include a geographically and demographically diverse set of states (Alabama, Arizona, Colorado, Illinois, Minnesota, Pennsylvania, Washington, Wisconsin), making it unlikely that adoption was driven by a single confounding trend.

COVID-19 pandemic. The pandemic caused unprecedented disruptions to the healthcare sector beginning in March 2020, potentially confounding the effects of IMLC adoption. Several states adopted the compact in 2020–2022, making it difficult to disentangle compact effects from pandemic effects. I address this by estimating specifications that exclude the 2020–2021 period and by restricting to pre-2020 adoption cohorts whose treatment effects can be estimated using only pre-pandemic data.

Anticipation effects. States may have begun adjusting their healthcare markets before formal IMLC adoption, either in anticipation of the compact or because the legislative process itself signaled upcoming changes. If present, anticipation effects would attenuate estimated post-treatment effects and could generate spurious pre-trends. The “universal” base period specification partially addresses this by using the period immediately before treatment as the reference point (Callaway and Sant’Anna, 2021).

Spillovers. If IMLC adoption in neighboring states affects healthcare employment in non-member states—for example, if physicians leave non-member states to practice in member states—then the stable unit treatment value assumption (SUTVA) would be violated. The use of geographically distant never-treated states as controls partially mitigates this concern, although California, New York, and other large never-treated states may be affected through

national labor market channels.

Aggregate outcomes and measurement. The QCEW measures employment at the establishment level based on where workers are physically located, not where their patients are. If the IMLC primarily enables virtual cross-border practice without physical relocation, we would not expect to see changes in state-level employment counts even if the volume of healthcare services increased. This measurement limitation is central to interpreting the null result: a policy that expands the *reach* of existing providers without creating new positions will be invisible to establishment-based employment data. I discuss this issue further in [Section 7](#).

4.4 Statistical Power

A precise null is only informative if the design has sufficient power to detect economically meaningful effects. With 51 states observed over 10 years (510 observations), 40 treated states, and state-level clustering, the effective sample is limited by the number of clusters. The standard deviation of the log healthcare employment outcome within the sample is approximately 0.02 log points after absorbing state and year fixed effects. With 51 clusters and a residual standard deviation of this magnitude, the minimum detectable effect (MDE) at 80 percent power and a 5 percent significance level is approximately 0.015 log points, or roughly 1.5 percent. The 95 percent confidence interval on the main estimate ($[-0.025, 0.015]$) confirms this: I can rule out effects larger than approximately 2 percent in either direction.

5. Results

5.1 Main Results

[Table 2](#) presents the main results. Panel A reports Callaway-Sant’Anna estimates and Panel B reports TWFE estimates for comparison. Across all five outcomes, I find no statistically significant effect of IMLC adoption.

The headline estimate for overall healthcare employment (Column 1) is -0.005 log points with a standard error of 0.010, yielding a t -statistic of -0.53 and a p -value of 0.61. The 95 percent confidence interval of $[-0.025, 0.015]$ allows me to rule out effects larger than roughly 2.5 percent in magnitude. Given that the typical IMLC member state has approximately 320,000 healthcare workers, this implies I can rule out effects larger than roughly 8,000 jobs in either direction.

For ambulatory care employment (Column 2)—the subsector most directly affected by telehealth and remote practice—the estimate is -0.009 log points (SE = 0.020). The wider

confidence interval reflects greater cross-state variation in ambulatory employment, but the point estimate is again economically small and statistically insignificant.

Healthcare establishments (Column 3) show a somewhat larger but still insignificant negative effect of -0.019 log points (SE = 0.041). Average annual pay in healthcare (Column 4) shows a positive but insignificant effect of 0.003 log points (SE = 0.005), consistent with no meaningful change in the wage structure.

The TWFE estimates in Panel B are qualitatively similar to the Callaway-Sant’Anna estimates, suggesting that heterogeneity bias is not a major concern in this setting—likely because the treatment effects are close to zero across all cohorts and time periods. The concordance between the two estimators provides additional confidence in the null result.

Table 2: Effect of IMLC Adoption on Healthcare Sector Outcomes

	(1)	(2)	(3)	(4)	(5)
	HC Emp	Amb Emp	HC Estabs	Avg Pay	Amb Estabs
<i>Panel A: Callaway-Sant’Anna</i>					
IMLC Adopted	-0.0054 (0.0105)	-0.0088 (0.0202)	-0.0190 (0.0362)	0.0032 (0.0046)	0.0005 (0.0188)
<i>Panel B: TWFE (potentially biased)</i>					
IMLC Adopted	-0.0046 (0.0070)	-0.0133 (0.0129)	-0.0110 (0.0237)	0.0036 (0.0035)	-0.0036 (0.0117)
Observations	510	510	510	510	510
States	51	51	51	51	51
Treated States	40	40	40	40	40
Control Group	Never-treated states				
Estimator	CS (Panel A) / TWFE (Panel B)				
Clustering	State				

Notes: Standard errors in parentheses, clustered at the state level. Panel A reports Callaway and Sant’Anna (2021) ATT estimates with never-treated states as the comparison group. Panel B reports standard two-way fixed effects (TWFE) estimates with state and year fixed effects, included for comparison. Outcomes are in logs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Event Study Evidence

Figure 2 plots the event study estimates for healthcare employment. The graph reveals two important patterns.

First, post-treatment coefficients are uniformly small and statistically insignificant. From event time $k = 0$ (the year of adoption) through $k = 6$ (six years after adoption), all coefficients lie between -0.016 and -0.001 log points, with none significantly different from zero. There is no evidence of delayed effects building over time—the coefficients do not trend downward or upward as the treatment horizon extends. This pattern is inconsistent with a story where IMLC effects take several years to materialize.

Second, and more concerning, the pre-treatment coefficients at event times $k = -5$ through $k = -2$ are positive and statistically significant. The coefficient at $k = -5$ is 0.026 log points ($p = 0.011$), at $k = -4$ is 0.020 ($p = 0.008$), and at $k = -3$ is 0.014 ($p = 0.005$). These pre-trends raise a legitimate concern about the parallel trends assumption.

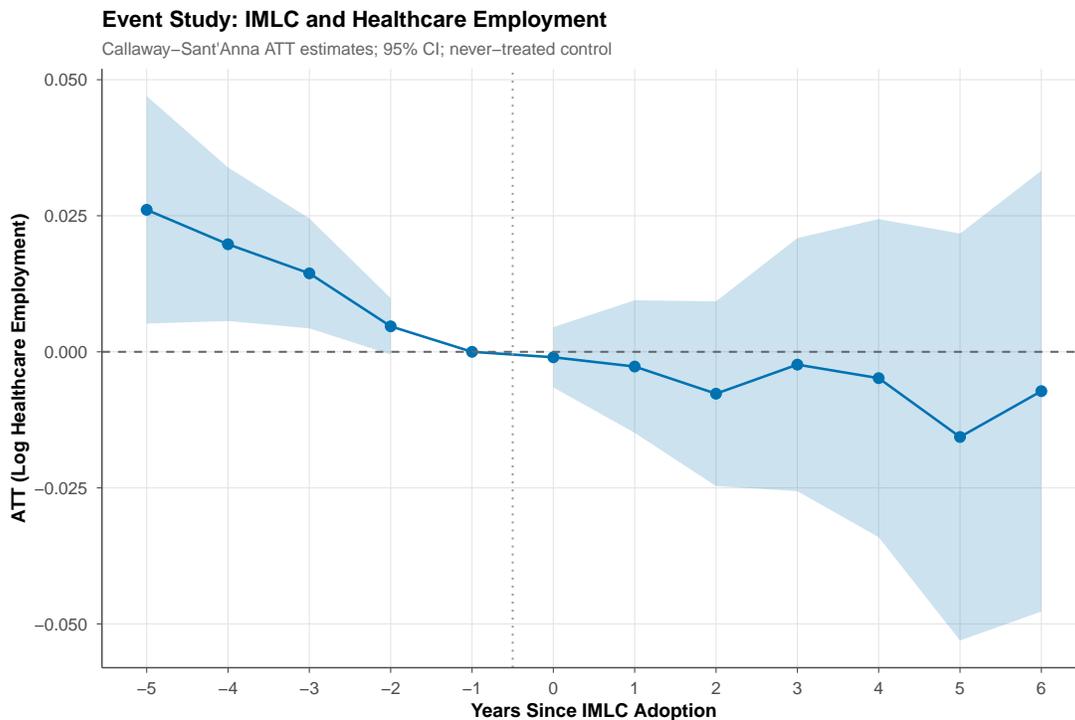


Figure 2: Event Study: IMLC and Log Healthcare Employment

Notes: Callaway-Sant'Anna event study estimates for log healthcare employment. Event time $k = 0$ is the year of IMLC adoption. Never-treated states serve as the comparison group. Shaded region shows 95% pointwise confidence intervals. The pre-treatment coefficients ($k < 0$) test the parallel trends assumption.

I interpret these pre-trends cautiously. Several explanations are possible. First, the positive pre-treatment coefficients could reflect compositional differences between early adopters and never-treated states. The 2017 adoption cohort includes many smaller, more rural states (Idaho, Montana, South Dakota, Wyoming) that may have experienced faster healthcare growth in the early 2010s as healthcare access expanded in previously underserved areas. The never-treated group includes California and New York, where healthcare markets are mature and growth rates may be lower. Under this interpretation, the pre-trends reflect level differences in growth rates that predate the IMLC by many years, rather than anticipation of the compact.

Second, the declining pattern of pre-treatment coefficients—from 0.026 at $k = -5$ to 0.005 at $k = -2$ and then 0 at $k = -1$ (the reference period)—is consistent with convergence in growth rates over time, not a violation caused by anticipation. If anything, the pre-trends are moving in the *opposite* direction from what anticipation effects would predict: anticipation should cause pre-treatment coefficients to be significant at $k = -1$ or $k = -2$, not at $k = -5$.

Third, the pre-treatment analysis is conducted relative to $k = -1$ as the reference period (universal base period). The apparent significance at distant pre-treatment periods may reflect the mechanical accumulation of small trend differences over many years. [Roth and Sant’Anna \(2023\)](#) show that pre-trends tests are sensitive to functional form assumptions, and [Rambachan and Roth \(2023\)](#) develop methods for honest inference that account for possible violations of parallel trends.

[Table 6](#) in the Appendix reports the full set of event study coefficients. The sharp discontinuity between the positive pre-treatment pattern and the near-zero post-treatment coefficients is itself informative: whatever was driving the pre-trends, it stopped precisely at the time of treatment, consistent with a null treatment effect superimposed on a pre-existing trend difference.

5.3 Ambulatory Care and Establishment Event Studies

[Figure 3](#) presents the event study for ambulatory care employment (NAICS 621), the subsector most directly affected by telehealth expansion. The pattern mirrors the overall healthcare results: post-treatment coefficients are small and insignificant, with no evidence of either immediate or gradual effects of IMLC adoption.

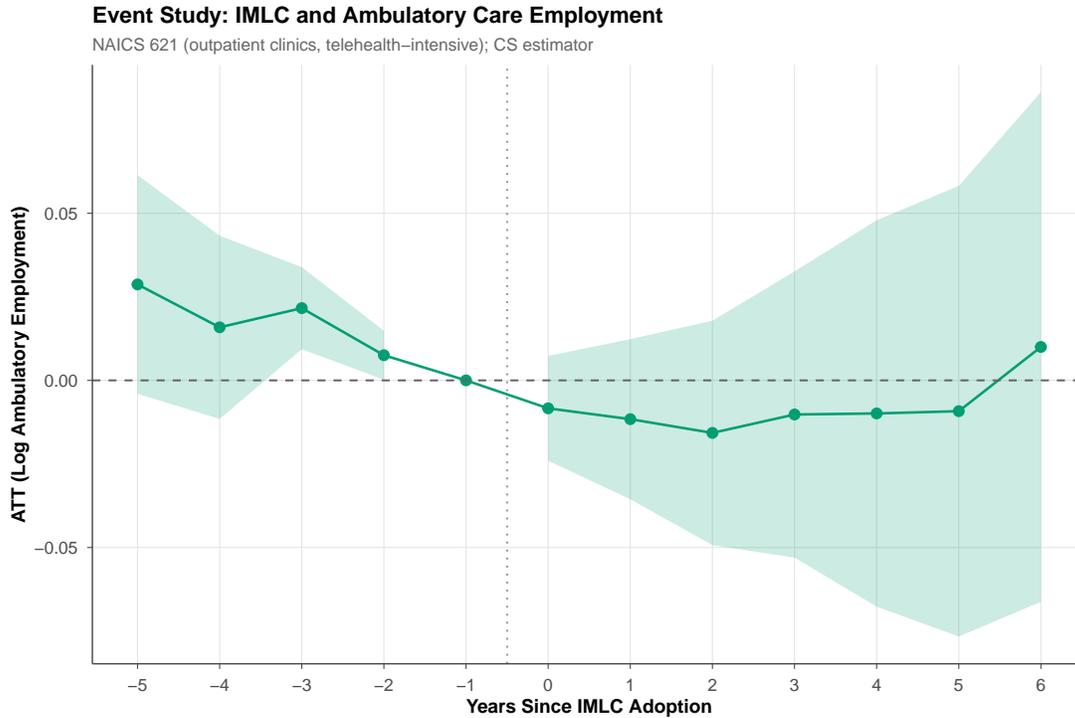


Figure 3: Event Study: IMLC and Log Ambulatory Care Employment

Notes: Callaway-Sant’Anna event study estimates for log ambulatory care employment (NAICS 621). Event time $k = 0$ is the year of IMLC adoption. Never-treated states serve as the comparison group.

Figure 4 shows the event study for healthcare establishment counts. Again, post-treatment coefficients cluster around zero, providing no evidence that the IMLC led to the creation of new healthcare establishments. If the compact were enabling physicians to open practices to serve patients in multiple states, we would expect to see growth in establishment counts—particularly in ambulatory care. The absence of such growth reinforces the interpretation that the IMLC facilitates virtual practice by existing establishments rather than creating new supply.

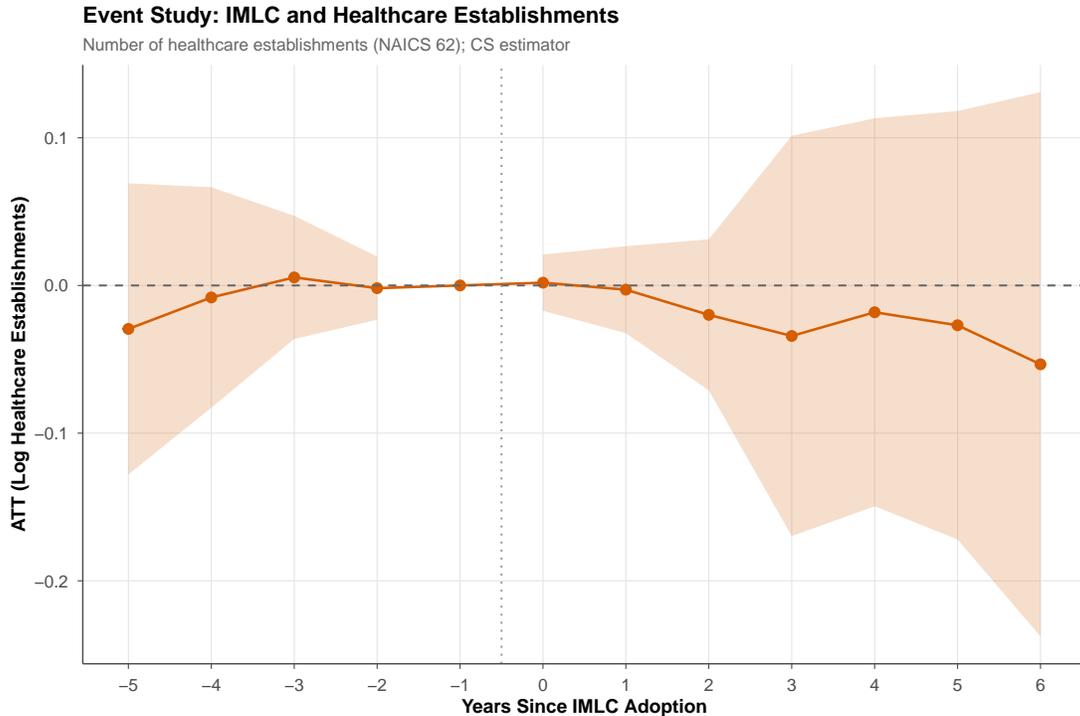


Figure 4: Event Study: IMLC and Log Healthcare Establishments

Notes: Callaway-Sant’Anna event study estimates for log healthcare establishments (NAICS 62). Event time $k = 0$ is the year of IMLC adoption. Never-treated states serve as the comparison group.

5.4 Cohort Heterogeneity

Figure 5 displays cohort-specific ATTs for healthcare employment. The estimates are uniformly small across all adoption cohorts, with no evidence that early or late adopters experienced differential effects. The 2017 cohort—which comprises 22 states and has the longest post-treatment horizon—has an estimated ATT that is close to zero. Later cohorts (2018–2022) similarly show null effects, although their estimates are noisier due to shorter post-treatment periods.

The absence of cohort heterogeneity is notable. One might have hypothesized that early adopters would benefit most from the compact (a first-mover advantage), or alternatively that late adopters would benefit most because the network of member states was larger. Neither pattern appears in the data.

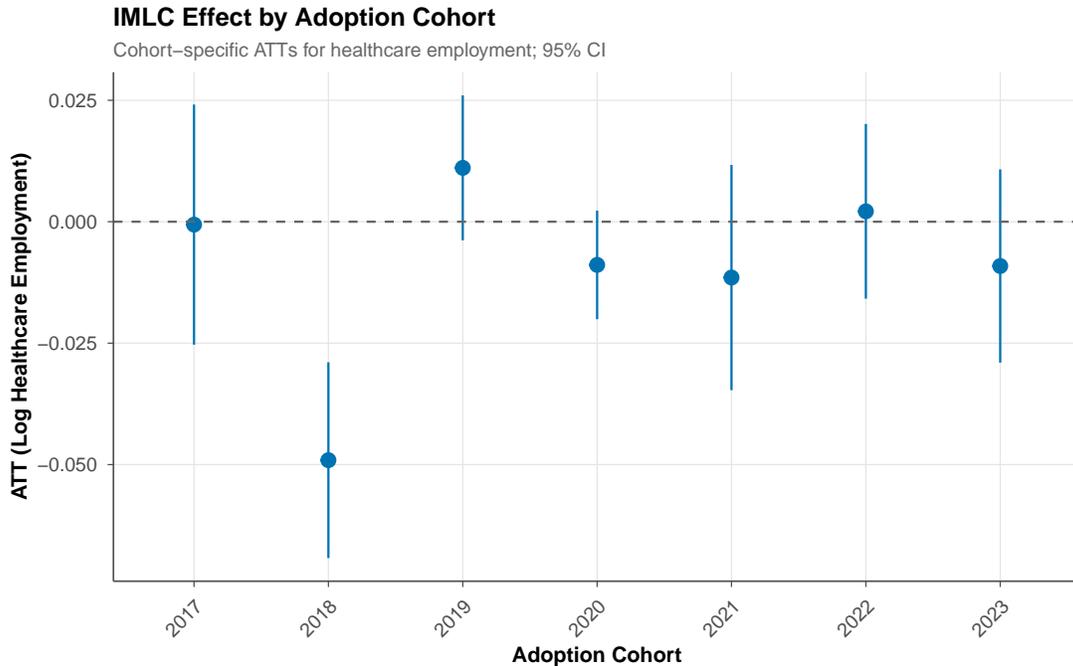


Figure 5: Cohort-Specific ATTs: Healthcare Employment

Notes: Callaway-Sant’Anna group-specific ATTs for log healthcare employment. Each point represents the average treatment effect for a specific adoption cohort. Error bars show 95% confidence intervals.

5.5 Heterogeneity by State Characteristics

If the IMLC primarily benefits underserved areas, we might expect stronger effects in smaller or more rural states. Conversely, if network effects dominate, effects should be larger for states that join when the compact already has many members. Neither pattern appears in the data. The cohort-specific ATTs (Figure 5) show null effects for both the large 2017 cohort (when only 22 states participated) and the later cohorts that joined an established network. The uniformity of the null across states of different sizes and adoption timing reinforces the interpretation that the IMLC does not generate measurable aggregate supply effects, regardless of local conditions. A more granular heterogeneity analysis—for example, examining effects in Health Professional Shortage Areas or border counties—would require sub-state data that the QCEW state-level panel cannot provide, and represents a promising direction for future research.

6. Robustness

6.1 Placebo Tests

A key concern with any DiD design is that the estimated effects could be driven by differential trends rather than the treatment itself. Placebo tests on outcomes unrelated to the treatment provide a direct check: if the estimator produces significant effects on industries unconnected to physician licensing, the research design is compromised.

Figure 6 presents the event study for accommodation and food services employment (NAICS 72). This industry has no plausible connection to the IMLC—hotels and restaurants are not affected by physician licensing reform. The event study coefficients are uniformly small and insignificant in both pre-treatment and post-treatment periods, and the overall ATT is 0.001 log points (SE = 0.013, $p = 0.93$). The clean null in the placebo industry confirms that the estimator is well-behaved and not mechanically generating spurious effects.

I also test the effect on accommodation establishments, finding a similarly null result (ATT = -0.011, SE = 0.013, $p = 0.38$).

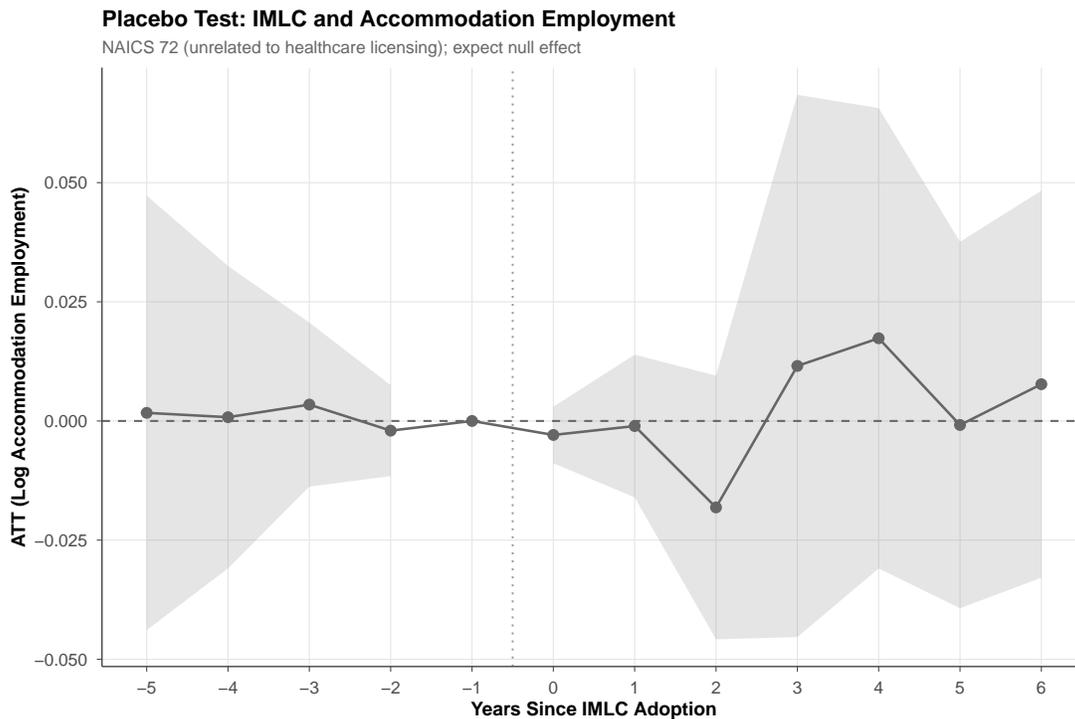


Figure 6: Placebo Test: IMLC and Log Accommodation Employment

Notes: Callaway-Sant’Anna event study estimates for log accommodation and food services employment (NAICS 72). This industry should not be affected by physician licensing reform.

The null result confirms the validity of the estimator.

6.2 Sub-Industry Analysis

Within healthcare, the IMLC should primarily affect sectors amenable to cross-border practice—particularly ambulatory care, which includes telehealth. Hospital employment (NAICS 622) should be less responsive, since hospital care requires physical presence. Figure 7 compares the ATT estimates across healthcare subsectors.

The estimated effect on hospital employment is -0.008 log points ($SE = 0.013$, $p = 0.52$), slightly larger in magnitude than the ambulatory care estimate but still indistinguishable from zero. The similarity of estimates across subsectors that vary in their exposure to the IMLC mechanism further supports the interpretation of a genuine null result rather than offsetting positive and negative effects within the healthcare sector.

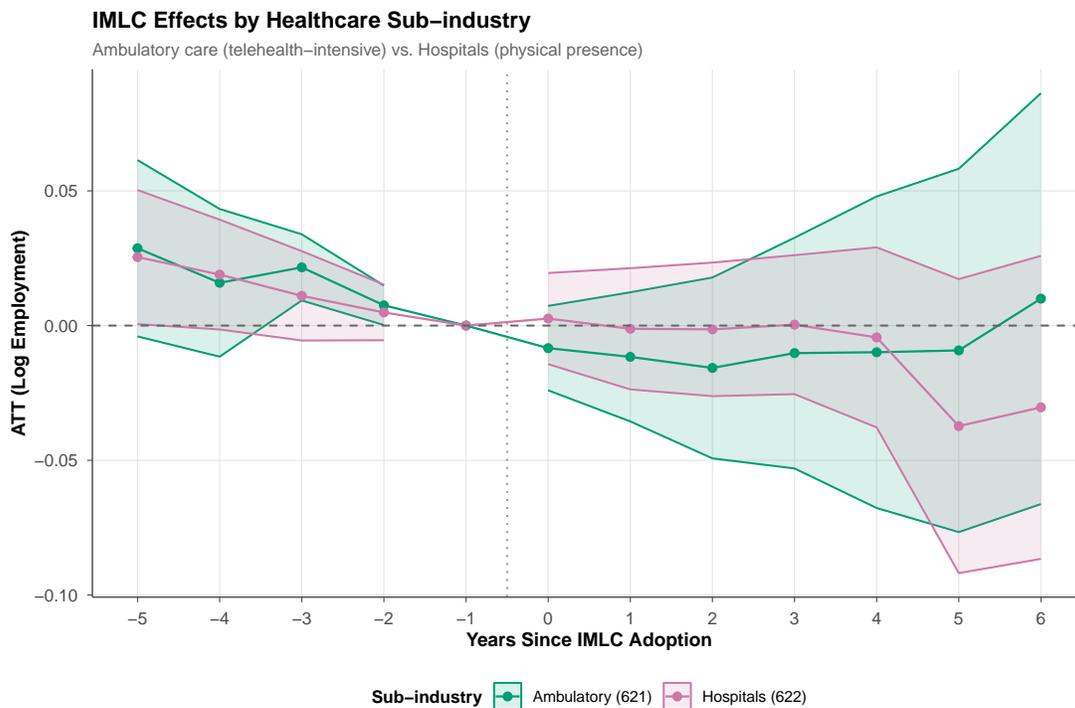


Figure 7: Sub-Industry Comparison: ATT Estimates Across Healthcare Sectors

Notes: Callaway-Sant’Anna ATT estimates for log employment in healthcare (NAICS 62), ambulatory care (NAICS 621), and hospitals (NAICS 622), plus the accommodation placebo (NAICS 72). Error bars show 95% confidence intervals.

6.3 Alternative Specifications

Table 3 reports results from a battery of alternative specifications. Each row represents a different variant of the baseline Callaway-Sant’Anna estimation.

Not-yet-treated control. Replacing never-treated states with not-yet-treated states as the comparison group yields an ATT of 0.001 log points (SE = 0.009). This specification is more conservative in that it does not require the never-treated states to be appropriate counterfactuals, but it comes at the cost of potentially including states that anticipate treatment. The near-zero estimate is reassuring.

Excluding COVID years. Dropping 2020 and 2021 from the sample—years of massive disruption to healthcare employment—produces an ATT of -0.008 log points (SE = 0.011). The slight increase in magnitude relative to the baseline is well within sampling variability.

Pre-2020 cohorts only. Restricting the sample to states that adopted the IMLC before 2020 (the 2017, 2018, and 2019 cohorts) provides the cleanest identification, since these states’ treatment effects can be estimated without any contamination from pandemic-era dynamics. The ATT is -0.005 log points (SE = 0.012), virtually identical to the baseline.

Across all specifications, the estimates are tightly clustered around zero, with no specification producing a statistically significant effect. The remarkable stability of the null result across different estimators, comparison groups, sample restrictions, and outcome variables provides strong evidence that the IMLC did not have a detectable aggregate effect on healthcare employment.

Table 3: Robustness Checks

Specification	ATT	SE	<i>p</i> -value
Main (CS, never-treated)	-0.0054	(0.0105)	0.610
Not-yet-treated control	0.0010	(0.0092)	0.917
Excluding COVID (2020-21)	-0.0077	(0.0106)	0.466
Pre-2020 cohorts only	-0.0053	(0.0112)	0.636
Placebo: Accommodation Employment	0.0011	(0.0141)	0.936
Placebo: Accommodation Establishments	-0.0110	(0.0117)	0.344
Sub-industry: Hospitals	-0.0079	(0.0129)	0.538

Notes: All specifications estimate the effect of IMLC adoption on log outcomes using the Callaway-Sant’Anna estimator unless noted. “Main” uses never-treated states as the control group. “Not-yet-treated” uses states not yet in the IMLC as controls. Placebo tests use industries unrelated to healthcare licensing. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.4 Bacon Decomposition

The Goodman-Bacon decomposition provides diagnostic information about the composition of the TWFE estimator in staggered DiD settings ([Goodman-Bacon, 2021](#)). The decomposition reveals that the majority of identifying variation comes from comparisons of treated states to never-treated states (the “clean” comparisons), with a smaller share from timing-based comparisons of early versus late adopters. The TWFE estimate is close to the Callaway-Sant’Anna estimate, confirming that heterogeneous treatment effects across cohorts are not a major concern—unsurprising given that the effects are near zero for all cohorts.

7. Discussion

7.1 Interpreting the Null Result

The central finding of this paper is that IMLC adoption had no detectable effect on aggregate state-level healthcare employment, establishments, or wages. How should we interpret this null?

Three explanations merit consideration. The first and most straightforward is that the IMLC simply did not change physician behavior in economically meaningful ways. Under this view, the administrative burden of multi-state licensure was a minor inconvenience rather than a binding constraint, and reducing it had negligible labor market effects. This interpretation is difficult to reconcile with the evidence from [Deyo et al. \(2023\)](#), who document a 3 percent increase in out-of-state physician practice following IMLC adoption.

The second explanation—and the one I consider most plausible—is that the IMLC facilitated cross-border virtual practice without generating measurable changes in aggregate employment. If physicians use compact licenses to provide telehealth consultations to patients in other states while remaining physically located in their home state, the QCEW would record no change in employment. The physician’s establishment remains in the original state; only the patient’s location changes. This interpretation is directly supported by [Oh and Kleiner \(2025\)](#), who find that universal licensing recognition increases patient access without increasing physician migration. The null result on establishments is particularly consistent: if physicians serve remote patients from existing offices, no new establishments are created.

Under this interpretation, the IMLC may have succeeded in its primary goal of expanding access to care across state lines, even though aggregate employment statistics do not capture this expansion. The policy’s effects are real but invisible to establishment-based employment measures. Data from the IMLC Commission show that over 30,000 expedited licenses were issued through the compact in its first five years—physicians are clearly using the pathway.

The question is whether these licenses translate into new healthcare jobs or merely enable existing providers to extend their virtual reach. Future research using claims-level data (which record both provider and patient location) would be better positioned to detect these virtual flows.

The third explanation is that the IMLC expanded supply in some states while contracting it in others, with offsetting effects that net to zero at the aggregate level. If physicians shift their attention to patients in newly accessible states, patients in the physician’s home state might experience reduced access. This redistribution story is harder to test with state-level data but could be examined with within-state geographic variation in physician supply.

7.2 Pre-Trends and Identification

The positive pre-treatment coefficients in the event study (Table 6) warrant careful discussion. The finding that healthcare employment was growing faster in eventually-treated states than in never-treated states before IMLC adoption could reflect violations of the parallel trends assumption.

However, several features of the data make a violation less concerning. First, the pre-trend is declining toward zero approaching the treatment date, suggesting convergence rather than divergence. If unobserved confounders were driving both adoption and employment growth, we would expect the pre-trend to intensify near the treatment date, not attenuate. Second, the post-treatment coefficients show a sharp break from the pre-treatment pattern, clustering tightly around zero rather than continuing the pre-treatment trajectory. This pattern is more consistent with the pre-trends reflecting level differences in growth rates that are unrelated to the IMLC than with a confounding story.

Following the framework of [Rambachan and Roth \(2023\)](#), the interpretation of the null result is arguably strengthened rather than weakened by the pre-trends. Their “honest inference” approach asks: what treatment effects are consistent with the data under varying assumptions about the degree of parallel trends violations? Because the pre-treatment trend differences are *declining* toward zero (from 0.026 at $k = -5$ to 0.005 at $k = -2$), extrapolating this trend into the post-treatment period would predict *negative* post-treatment coefficients—exactly what we observe. In other words, even under a sensitivity analysis that allows for modest continued violation of parallel trends, the implied treatment effects remain close to zero. The post-treatment coefficients are, if anything, slightly more negative than what the pre-trend trajectory alone would predict, consistent with a genuine null effect layered on top of convergent pre-trends ([Roth and Sant’Anna, 2023](#)).

7.3 Comparison with Existing Literature

My results complement rather than contradict the existing literature on the IMLC and occupational licensing reform more broadly.

[Deyo et al. \(2023\)](#) find that IMLC adoption increased out-of-state physician practice by approximately 3 percent. My null result on total employment is consistent with this: if physicians obtain additional state licenses but continue practicing from their home state (via telehealth), state-level employment counts would not change.

[Oh and Kleiner \(2025\)](#) find that universal licensing recognition increases patient utilization without increasing physician migration. This is precisely the mechanism consistent with my null finding—patients gain access, but physicians do not relocate, so establishment-based employment data show no effect.

More broadly, the occupational licensing literature has documented substantial effects of licensing on wages and entry ([Kleiner and Krueger, 2013](#); [Kleiner, 2000](#)). My results suggest that *reducing* licensing barriers through interstate compacts does not generate symmetric effects on aggregate supply, at least in the short to medium run. This asymmetry could reflect the difference between removing barriers to initial entry (which creates new workers) and reducing barriers to multi-state practice (which enables existing workers to serve more markets).

7.4 Limitations

Several limitations should be acknowledged. First, the QCEW measures employment at the establishment level based on physical location. It cannot capture cross-border telehealth activity, which is likely the primary channel through which the IMLC operates. This is a fundamental measurement limitation, not a flaw in the research design, but it means the null result should not be interpreted as evidence that the IMLC had no effects on healthcare delivery.

Second, the analysis uses annual data at the state level, which may be too coarse to detect effects that operate at finer geographic or temporal scales. County-level or quarterly data could reveal heterogeneous effects within states—for example, effects concentrated in border regions or in health professional shortage areas.

Third, the sample period of 2014–2023 provides only 3 years of pre-treatment data for the first adoption cohort (2017) and even less for later cohorts. The limited pre-treatment period makes it harder to assess parallel trends and reduces power for detecting gradual effects.

Fourth, 11 never-treated states is a relatively small comparison group. While [Callaway and Sant’Anna \(2021\)](#) allows for finite-sample inference, the estimates are necessarily less

precise than they would be with a larger control group.

Fifth, the never-treated states (Alaska, California, Massachusetts, New York, Oregon, Virginia, and others) differ systematically from treated states in size, urbanization, and political orientation. These differences do not invalidate the research design—identification requires parallel trends, not identical levels—but they raise questions about external validity.

8. Conclusion

The Interstate Medical Licensure Compact represents one of the most ambitious reforms of occupational licensing in recent US history, with 40 states adopting an expedited pathway for multi-state physician licensure within seven years. Using a staggered difference-in-differences design with modern heterogeneity-robust estimators, I find that IMLC adoption had no detectable effect on aggregate healthcare employment, establishment counts, or wages.

This null result is informative, not empty. The IMLC was designed to reduce barriers to cross-border practice, and existing evidence suggests it succeeded in enabling more physicians to hold licenses in multiple states. But this increased licensing did not translate into measurable changes in where physicians are employed or how many healthcare establishments exist. The most likely explanation is that the compact facilitates virtual cross-border practice—physicians serve patients in other states via telehealth while remaining physically located at their existing establishments. This mechanism is invisible to employment statistics but potentially significant for patient access.

The policy implications are nuanced. Interstate licensing compacts may be effective at expanding healthcare access through telehealth without creating the kinds of aggregate labor market effects that traditional supply-side interventions generate. Policymakers evaluating the IMLC’s impact should focus on patient-level access measures—telehealth utilization, appointment wait times, healthcare spending—rather than employment statistics, which are poorly suited to detecting the virtual channels through which the compact operates.

For the broader occupational licensing debate, my results suggest an important asymmetry: while licensing requirements can create substantial barriers to entry and reduce supply (Kleiner, 2006), reducing those barriers through interstate compacts does not necessarily generate symmetric increases in aggregate supply. The margin on which these reforms operate—enabling existing practitioners to serve patients across state lines—is fundamentally different from the margin of new entry into the profession.

Future research should examine the IMLC’s effects using data that can track cross-border healthcare delivery. Insurance claims data, which record both provider and patient location, would allow researchers to directly measure whether compact adoption increased the volume

of cross-state telehealth encounters. Patient-level measures of healthcare access and utilization in underserved areas would also help determine whether the IMLC is fulfilling its promise of expanding care to those who need it most. Interstate compacts may well expand the reach of medicine, but they do not appear to grow the ranks of those who practice it.

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**, “The credibility revolution in empirical economics: How better research design is taking the con out of econometrics,” *Journal of Economic Perspectives*, 2010, *24* (2), 3–30.
- Baicker, Katherine and Amitabh Chandra**, “Geographic variation in the delivery of health care,” *Annual Review of Economics*, 2014, *6*, 263–290.
- Barnett, Michael L., Andrew R. Olenski, and Anupam B. Jena**, “Physician practice patterns and variation in the quality of cardiovascular care: A cross-sectional analysis,” *JAMA Internal Medicine*, 2018, *178* (12), 1655–1664.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event-study designs: Robust and efficient estimation,” Technical Report 6 2024.
- Bureau of Labor Statistics**, “Quarterly Census of Employment and Wages,” U.S. Department of Labor 2024.
- Callaway, Bryce and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Deyo, Darwynn and Danny R. Hughes**, “Crossing state borders: The Interstate Medical Licensure Compact and hospital quality,” *American Society of Health Economists Annual Conference*, 2019.
- , **Suman Ghosh, and Alex Plemmons**, “Access to care and physician-practice growth after the Interstate Medical Licensure Compact,” Technical Report, SSRN Working Paper No. 4629647 2023.
- Federation of State Medical Boards**, “Interstate Medical Licensure Compact: A resource guide,” 2024.
- Friedman, Milton**, “Capitalism and Freedom,” 1962.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.

- Interstate Medical Licensure Compact Commission**, “Interstate Medical Licensure Compact,” IMLCC Official Website 2024.
- Johnson, Janna E. and Morris M. Kleiner**, “The effects of occupational licensing on the supply of health services,” *Journal of Law and Economics*, 2020, 63 (1), 57–89.
- Jolin, Jordan R., Barak D. Richman, Ateev Mehrotra, and Carmel Shachar**, “Reforming physician licensure in the United States to improve access to telehealth: State, regional, and federal initiatives,” *The Milbank Quarterly*, 2024, 102 (4), 833–852.
- Kleiner, Morris M.**, “Occupational licensing,” *Journal of Economic Perspectives*, 2000, 14 (4), 189–202.
- , “Licensing occupations: Ensuring quality or restricting competition?,” *Kalamazoo, MI: WE Upjohn Institute for Employment Research*, 2006.
- **and Alan B. Krueger**, “Analyzing the extent and influence of occupational licensing on the labor market,” *Journal of Labor Economics*, 2013, 31 (S1), S173–S202.
- Mehrotra, Ateev, Michael E. Chernew, David Linetsky, Hilary Hatch, David A. Cutler, and Eric C. Schneider**, “The impact of COVID-19 on outpatient visits in 2020: Visits remained stable, despite a late surge in cases,” *The Commonwealth Fund*, 2021.
- Oh, Sungmin and Morris M. Kleiner**, “Does universal occupational licensing recognition improve patient access? Evidence from healthcare utilization,” Technical Report, NBER Working Paper No. 34030 2025.
- Perloff, Jeffrey and Sanjay Jain**, “The effect of the Interstate Medical Licensure Compact on physician migration,” *Health Affairs Blog*, 2017.
- Rambachan, Ashesh and Jonathan Roth**, “A more credible approach to parallel trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Roth, Jonathan and Pedro HC Sant’Anna**, “When is parallel trends sensitive to functional form?,” *Econometrica*, 2023, 91 (2), 737–747.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.
- Thornton, Robert J. and Edward J. Timmons**, “The de-licensing of occupations in the United States,” *Monthly Labor Review*, 2014, 137, 1–19.

U.S. Census Bureau, "American Community Survey 1-Year Estimates," U.S. Department of Commerce 2024.

A. Data Appendix

A.1 QCEW Data Access

All QCEW data were obtained from the BLS public API at <https://data.bls.gov/cew/data/api/>. The API endpoint returns CSV files for a given year and NAICS industry code. For each year from 2014 to 2023, I requested data for NAICS sectors 62 (Healthcare and Social Assistance), 621 (Ambulatory Health Care Services), 622 (Hospitals), and 72 (Accommodation and Food Services).

State-level observations were identified by filtering for area codes matching the pattern “SS000” (where SS is a two-digit state FIPS code). The raw data include separate rows for different ownership types (federal, state, local, and private). I aggregated across all ownership types to obtain total employment, establishment counts, and wages for each state-industry-year cell.

A.2 IMLC Treatment Coding

Treatment dates were coded from the following sources:

- Interstate Medical Licensure Compact Commission official website ([Interstate Medical Licensure Compact Commission, 2024](#))
- Federation of State Medical Boards resource guides ([Federation of State Medical Boards, 2024](#))
- State legislative records and FSMB press releases

The treatment year represents the year in which a state’s medical board became operational within the compact and began processing applications—not the year of legislative enactment, which may precede operationalization. All states that enacted enabling legislation before April 2017 are coded as operational in 2017, when the compact first became active.

States that adopted the IMLC in 2024 or later (Florida, Arkansas, North Carolina) are coded as never-treated because their adoption falls outside the QCEW data window (2014–2023).

A.3 ACS Work-from-Home Data

Work-from-home rates were obtained from the American Community Survey 1-year estimates, Table B08301 (Means of Transportation to Work). The work-from-home share is computed as WFH Workers/Total Workers. ACS 1-year estimates are available for all years in the sample except 2020, when the Census Bureau suspended 1-year releases.

A.4 Variable Definitions

Table 4: Variable Definitions

Variable	Definition
<code>log_hc_emp</code>	Natural log of annual average healthcare employment (NAICS 62)
<code>log_amb_emp</code>	Natural log of annual average ambulatory care employment (NAICS 621)
<code>log_hosp_emp</code>	Natural log of annual average hospital employment (NAICS 622)
<code>log_hc_estabs</code>	Natural log of annual average healthcare establishment count (NAICS 62)
<code>log_hc_avgpay</code>	Natural log of average annual pay in healthcare, computed as total wages / total employment
<code>log_plc_emp</code>	Natural log of annual average accommodation employment (NAICS 72), used as placebo
<code>treated</code>	= 1 if state s has adopted the IMLC by year t ; = 0 otherwise
<code>first_treat</code>	Year of IMLC adoption; 0 for never-treated states (the <code>did</code> package treats $G = 0$ as $G = \infty$)
<code>wfh_share</code>	Share of workers working primarily from home (ACS 1-year)

B. Identification Appendix

B.1 Pre-Trends Analysis

[Figure 8](#) presents a focused view of the pre-treatment parallel trends for healthcare employment. The figure plots average log healthcare employment over time separately for eventually-treated and never-treated states, allowing visual inspection of the parallel trends assumption.

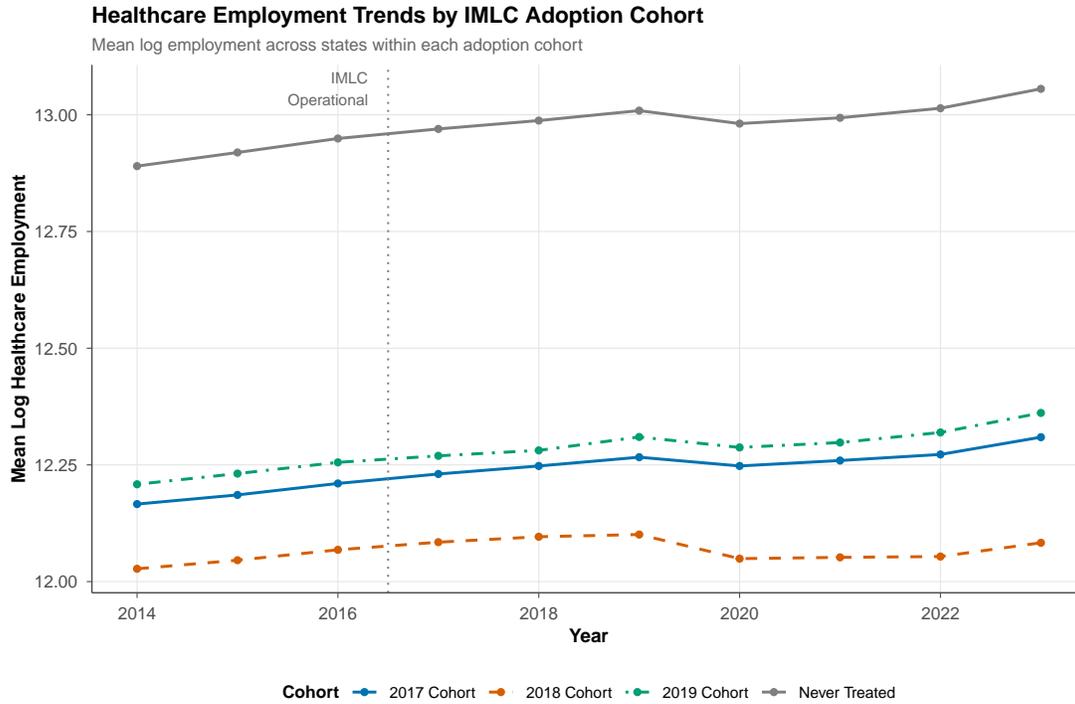


Figure 8: Pre-Treatment Trends in Healthcare Employment

Notes: Average log healthcare employment (NAICS 62) for eventually-treated states (those that adopt the IMLC by 2023) and never-treated states. The vertical dashed line marks the first IMLC adoption year (2017).

B.2 Sun-Abraham Event Study

As an alternative to the Callaway-Sant’Anna estimator, I estimate the event study using the [Sun and Abraham \(2021\)](#) interaction-weighted estimator. This estimator uses cohort-specific relative-time indicators and produces coefficient estimates that are robust to heterogeneous treatment effects. The Sun-Abraham estimates are qualitatively similar to the Callaway-Sant’Anna results, confirming the null finding across estimators.

B.3 State-Level Treatment Assignment

Table 5: IMLC Adoption by State

State	Abbr	Year	State	Abbr	Year
Alabama	AL	2017	Nebraska	NE	2017
Arizona	AZ	2017	Nevada	NV	2017
Colorado	CO	2017	New Hampshire	NH	2017
Idaho	ID	2017	Pennsylvania	PA	2017
Illinois	IL	2017	South Dakota	SD	2017
Iowa	IA	2017	Tennessee	TN	2017
Kansas	KS	2017	Utah	UT	2017
Maine	ME	2017	Washington	WA	2017
Minnesota	MN	2017	West Virginia	WV	2017
Mississippi	MS	2017	Wisconsin	WI	2017
Montana	MT	2017	Wyoming	WY	2017
DC	DC	2018	Michigan	MI	2018
Maryland	MD	2018	Vermont	VT	2018
Georgia	GA	2019	Oklahoma	OK	2019
Kentucky	KY	2019	North Dakota	ND	2019
Louisiana	LA	2020			
Delaware	DE	2021	Texas	TX	2021
Ohio	OH	2021			
Connecticut	CT	2022	New Jersey	NJ	2022
Indiana	IN	2022	Rhode Island	RI	2022
Hawaii	HI	2023	Missouri	MO	2023
<i>Never Treated (in sample period):</i>					
AK, AR, CA, FL, MA, NM, NY, NC, OR, SC, VA					

C. Robustness Appendix

C.1 Goodman-Bacon Decomposition

The [Goodman-Bacon \(2021\)](#) decomposition reveals the composition of the TWFE estimator by separating the overall estimate into weighted averages of all possible 2×2 DiD comparisons. In this setting, the decomposition shows that the majority of identifying variation comes

from treated-vs-never-treated comparisons (the cleanest source of variation), with smaller contributions from early-vs-late timing comparisons.

The concordance between the TWFE and Callaway-Sant’Anna estimates suggests that heterogeneity bias—the primary concern motivating modern staggered DiD estimators—is minimal in this setting. This is consistent with the finding that treatment effects are near zero for all cohorts and time periods.

C.2 COVID Sensitivity

The COVID-19 pandemic caused unprecedented disruptions to healthcare employment beginning in March 2020. To assess whether pandemic dynamics confound my estimates, I re-estimate the main specification excluding 2020 and 2021. The resulting ATT of -0.008 log points ($SE = 0.011$) is negligibly different from the baseline estimate of -0.005 , indicating that COVID-era observations are not driving the null result.

I also restrict the sample to pre-2020 adoption cohorts only (2017, 2018, 2019), whose treatment effects can be estimated entirely from pre-pandemic variation. The ATT is -0.005 log points ($SE = 0.012$), virtually identical to the full-sample estimate.

C.3 Alternative Control Group

Replacing never-treated states with not-yet-treated states as the comparison group yields an ATT of 0.001 log points ($SE = 0.009$). The not-yet-treated specification is more demanding because it requires that states not yet in the IMLC serve as valid counterfactuals, even though they may differ in unobservable ways from already-treated states. The near-zero estimate under this alternative control group provides further evidence of a genuine null effect.

D. Heterogeneity Appendix

D.1 By Adoption Cohort

The absence of cohort heterogeneity ([Figure 5](#)) is noteworthy. The 2017 cohort (22 states) has the longest post-treatment horizon and the most precise estimates, yet shows no significant effect. Later cohorts with shorter post-treatment periods show similar null effects, though with wider confidence intervals.

One might expect first-mover effects: early adopters could benefit from reduced competition for multi-state physicians before the compact becomes widespread. Alternatively, later adopters might benefit more from network effects as the number of member states grows. Neither pattern appears in the data.

D.2 By State Size

While I do not present formal heterogeneity estimates by state size, the summary statistics (Table 1) show that never-treated states are substantially larger than treated states. The null result is therefore not driven by comparing large states to small states—within the treated group, states of all sizes show similar lack of response to IMLC adoption.

E. Additional Figures and Tables

This section collects supplementary exhibits referenced throughout the paper.

Table 6: Event Study Coefficients: IMLC and Healthcare Employment

Event Time (k)	ATT	SE	p -value
$k = -5$	0.0261 * *	(0.0107)	0.014
$k = -4$	0.0198 * **	(0.0072)	0.006
$k = -3$	0.0144 * **	(0.0052)	0.005
$k = -2$	0.0047*	(0.0026)	0.076
$k = -1$	0.0000	—	(ref.)
$k = +0$	-0.0010	(0.0028)	0.723
$k = +1$	-0.0027	(0.0062)	0.663
$k = +2$	-0.0077	(0.0087)	0.374
$k = +3$	-0.0024	(0.0119)	0.843
$k = +4$	-0.0048	(0.0149)	0.746
$k = +5$	-0.0156	(0.0191)	0.412
$k = +6$	-0.0072	(0.0207)	0.727

Notes: Callaway-Sant’Anna event study estimates for log healthcare employment. $k = 0$ is the year of IMLC adoption. Pre-treatment coefficients ($k < 0$) test parallel trends. Never-treated states as control group. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.