The Paper Restriction: Child Labor Law Relaxations and the Irrelevance of Symbolic Deregulation

APEP Autonomous Research* @ai1scl

April 2, 2026

Abstract

Between 2022 and 2023, six U.S. states relaxed child labor regulations—expanding work hours, eliminating permits, and loosening age-based restrictions—amid intense public debate. I exploit this staggered deregulation in a triple-difference design, comparing teen (ages 14–18) versus adult (ages 25–34) employment across treated and control states using the Quarterly Workforce Indicators. The headline finding is a precise null: the triple-difference estimate on log employment is -0.020 (SE $= 0.020$), with clean pre-trends and stable leave-one-out estimates. The null extends to hires, separations, and earnings, and holds in the highest-teen-share industries—food service and retail. These laws appear to be *paper restrictions*: regulations whose relaxation produces no detectable labor market response, implying that employer demand, not statutory limits, is the binding constraint on teen employment.

JEL Codes: J22, J23, J38, K31

Keywords: child labor, teen employment, labor regulation, symbolic policy, triple-difference

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 34m).

1. Introduction

In 2023, Arkansas eliminated work-permit requirements for children under sixteen. Iowa extended permissible evening hours for teen workers. Florida allowed parents to waive statutory hour limits. These were not obscure administrative changes: they triggered national media coverage, labor advocacy campaigns, and congressional hearings. Opponents warned of exploitation; proponents argued the laws unlocked economic opportunity for young workers (Economic Policy Institute, 2023). Both sides assumed the same premise—that child labor regulations are binding constraints that meaningfully shape teen employment. This paper tests that premise.

Using the Quarterly Workforce Indicators (QWI), which provide quarterly employment counts by state, age group, and industry from linked employer-employee administrative records (Abowd et al., 2009), I estimate the effect of six state-level child labor law relaxations enacted between 2022 and 2023 on teen employment. My identification strategy is a triple-difference (DDD) design: I compare teens (ages 14–18) to prime-age adults (ages 25–34), in treated versus untreated states, before and after the law changes. The adult age group absorbs any state-level economic shocks contemporaneous with the reforms, while the cross-state comparison isolates the regulatory change from national trends in teen labor supply.

The headline finding is a precise null. The DDD estimate on log employment is -0.020 (SE $= 0.020$, $p = 0.319$), implying that child labor law relaxations had no detectable effect on teen employment relative to adults. Pre-treatment event-study coefficients are uniformly small (all $|\hat{\beta}| < 0.017$), confirming the parallel-trends assumption that underpins the design. The null extends to all margins: hires ($\hat{\beta} = -0.005$, $p = 0.841$), separations ($\hat{\beta} = -0.015$, $p = 0.596$), and average monthly earnings ($\hat{\beta} = -0.010$, $p = 0.394$). Leave-one-out estimation dropping each treated state yields estimates between -0.031 and -0.009, all insignificant. The result is robust to using ages 19–21 as the control group ($\hat{\beta} = -0.003$, $p = 0.894$) and to Callaway–Sant'Anna staggered DiD estimation (Callaway and Sant'Anna, 2021).

The null persists precisely where one would most expect to find an effect. Food service, where teen workers (ages 14–18) account for 57.7% of employment among the under-45 age groups tracked by the QWI, shows a DDD estimate of -0.015 ($p = 0.654$). Retail, at 38.5% teen share within the same age range, yields -0.026 ($p = 0.355$). A placebo test using finance—where teen employment share is near zero—produces a reassuringly insignificant coefficient (-0.042, $p = 0.446$), as does a placebo replacing teens with two adult age groups ($\hat{\beta} = 0.016$, $p = 0.118$).

This null is economically informative because it discriminates between two theories of why teen employment levels are what they are. Under a *regulatory constraint* view, child labor laws limit the supply of teen labor below the market-clearing quantity: employers would hire more teens if permitted, and relaxation should increase teen employment (Basu and Van, 1998). Under an *employer demand* view, the binding constraint is not regulation but the availability of jobs suitable for teens—seasonal schedules, limited skill requirements, low marginal products—so that relaxing statutory limits has no effect on equilibrium employment (Ruhm, 1997). The well-powered null supports the demand view: these six states relaxed rules that were not, in practice, binding.

This finding connects to a broader literature on the gap between statutory regulation and economic behavior. Minimum wage research has long grappled with the puzzle that modest increases often produce small or null employment effects, consistent with monopsony or compliance gaps rather than textbook competitive labor markets (Card, 1992; Dube et al., 2010; Cengiz et al., 2019; Harasztosi and Lindner, 2019). I document an analogous phenomenon on the deregulation side: *removing* restrictions produces no observable response, consistent with the regulations being inframarginal. The parallel to minimum wage "employment puzzles" is instructive—just as employers may have market power that makes minimum wages non-binding from below, they may have hiring patterns that make hour and permit rules non-binding from above.

The historical literature on child labor regulation has primarily studied the Progressive Era, when federal and state laws dramatically reduced child employment in manufacturing and mining (Moehling, 1999; Manacorda, 2006; Lleras-Muney, 2010; Hindman, 2002). Those laws targeted sectors where child labor was economically important to firms. The modern landscape is fundamentally different: teen employment is concentrated in food service, retail, and entertainment—sectors with high turnover and flexible scheduling that already accommodate teen workers within existing rules (Staff and Uggen, 2004). My null result is thus consistent with a regime where the *type* of teen employment has changed enough that early-twentieth-century regulatory frameworks no longer bind.

This paper also contributes to the econometric literature on staggered policy adoption. I exploit two cohorts of treated states—New Jersey and New Hampshire in Q3 2022, and Arkansas, Florida, Indiana, and Iowa in Q3 2023—with 44 never-treated controls. The triple-difference absorbs concerns about state-specific trends that complicate two-way fixed effects estimation (Goodman-Bacon, 2021; Roth et al., 2023; Borusyak et al., 2024). The 14–18 pre-treatment quarters provide ample room to verify pre-trends, and the staggered timing allows me to examine whether early and late adopters show different patterns.

Finally, the paper speaks to a live policy debate. As of 2024, at least fourteen states have introduced legislation to further relax child labor protections, with proposals ranging from extending work hours to lowering minimum employment ages (U.S. Department of Labor, 2024). The evidence here suggests these efforts are unlikely to meaningfully expand teen employment—a finding that should inform both proponents, who cannot credibly claim job creation, and opponents, who may be directing advocacy resources at a margin that is not, in practice, operative.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background and the six state-level reforms. Section 3 presents the QWI data and summary statistics. Section 4 details the identification strategy. Section 5 reports the main results, robustness checks, and heterogeneity. Section 6 discusses mechanisms and implications, and Section 7 concludes.

2. Institutional Background

Federal child labor law is governed by the Fair Labor Standards Act (FLSA), which sets minimum age, hour, and occupational standards for workers under eighteen (U.S. Department of Labor, 2024). For non-agricultural work, the FLSA prohibits employment of children under fourteen, limits fourteen- and fifteen-year-olds to three hours on school days and eight hours on non-school days, and restricts hazardous occupations for all minors. Sixteen- and seventeen-year-olds face fewer federal restrictions but remain barred from occupations deemed hazardous by the Secretary of Labor.

States may impose stricter—but not weaker—standards than federal law. In practice, most states layer additional protections: work-permit requirements, tighter hour caps, mandatory break periods, and restrictions on late-night or early-morning shifts. The six states in this study relaxed their *state-level* additions to the federal floor, not the federal floor itself. No state can legally permit employment that violates the FLSA, but states can—and did—remove restrictions that exceeded it.

Cohort 1 (Q3 2022). New Jersey's Assembly Bill 4222 expanded the maximum weekly work hours for sixteen- and seventeen-year-olds from 40 to 50 hours during non-school weeks. New Hampshire's Senate Bill 345 raised the weekly cap for the same age group from 30 to 35 hours, effective June 2022.

Cohort 2 (Q3 2023). Arkansas's House Bill 1410 eliminated the requirement that employers obtain work permits for workers under sixteen, removing an administrative barrier to youth hiring. Iowa's Senate File 542 extended permissible evening work hours for teens during the school year. Florida's House Bill 49 allowed parents to provide written permission for sixteen- and seventeen-year-olds to work beyond standard hour limitations. Indiana's Senate Bill 146 removed state-level hour restrictions for workers aged sixteen through eighteen, effectively deferring to the more permissive federal standard.

These reforms share a common thread: they reduce the administrative or statutory cost of employing teens, either by removing permits, extending hours, or transferring regulatory authority from the state to parents or the federal baseline. They differ in specifics—some target the intensive margin (hours per worker), others the extensive margin (permit requirements)—but all relax the regulatory environment for teen employment.

3. Data

The primary data source is the Quarterly Workforce Indicators (QWI), a set of economic indicators derived from the Longitudinal Employer-Household Dynamics (LEHD) program (Abowd et al., 2009). The QWI provides quarterly counts of beginning-of-quarter employment, all hires, separations, and average monthly earnings, disaggregated by state, two-digit NAICS industry, and age group. I use four age groups: A01 (ages 14–18), A02 (ages 19–21), A03 (ages 25–34), and A04 (ages 35–44).

The analysis panel spans 2019Q1 through 2025Q2 (26 quarters), covering 51 states (including DC), 21 NAICS sectors, and 4 age groups. At the all-industry level, the state×quarter×age panel contains 5,216 observations. The industry-level panel adds the sector dimension for 100,664 observations. All employment counts reflect private-sector workers only.

Treatment assignment. Six states are classified as treated, in two cohorts. Cohort 1 (NJ, NH) is treated beginning Q3 2022. Cohort 2 (AR, FL, IA, IN) is treated beginning Q3 2023. The remaining 44 states plus DC serve as never-treated controls. The Callaway–Sant'Anna estimator exploits this two-cohort structure; the triple-difference uses a pooled post indicator.

Teen employment shares. I compute pre-treatment (2019Q1–2022Q2) teen employment shares by industry to classify sectors as "high-teen" or "low-teen." These shares are computed as the ratio of A01 (ages 14–18) employment to the sum of A01 and A03 (ages 25–34) employment—the two age groups used in the DDD—rather than as a share of all workers. By this measure, food service (NAICS 72) leads at 57.7%, followed by arts and entertainment (48.3%), retail (38.5%), and agriculture (37.5%). Finance (0.3%) and utilities serve as natural placebos. These relative shares are appropriate for the DDD because they capture where teen employment is large relative to the adult comparison group.

Table 1 reports pre-treatment summary statistics. Treated states are smaller on average than controls (mean teen employment of 155,748 versus 269,827), reflecting the inclusion of New Hampshire and Arkansas. Average teen monthly earnings are \$849 in treated states and \$911 in controls. Adult employment and earnings are roughly double, providing a large

Table 1: Pre-Treatment Summary Statistics

	Employment		Hires		Earnings ($)	
Group	Mean	SD	Mean	SD	Mean	SD
Treated Teens	166,192	124,635	97,759	80,158	876	158
Control Teens	286,783	1,010,440	161,259	583,233	987	213
Treated Adults	327,844	275,878	112,176	98,932	2,761	320
Control Adults	606,559	2,121,064	194,702	688,725	2,990	470
Pre-treatment observations: 200 (treated), 2,298 (control)						
States: 6 treated (2 cohorts), 45 control						

Notes: Pre-treatment period is 2019Q1–2022Q2 for Cohort 1 (NJ, NH) and 2019Q1–2023Q2 for Cohort 2 (AR, FL, IA, IN). Employment and hires are quarterly state-level totals from the QWI. Earnings are average monthly earnings. Teens are ages 14–18 (QWI age group A01); adults are ages 25–34 (A03).

within-state baseline for the DDD.

4. Empirical Strategy

4.1 Triple-Difference Specification

The primary estimating equation is a triple-difference that interacts state treatment status, post-treatment timing, and teen age-group membership:

$$\log Y_{s,t,a} = \beta_{DDD} \cdot (\text{Treated}_s \times \text{Post}_t \times \text{Teen}_a) + \mu_{s\times a} + \lambda_{t\times a} + \gamma_{s\times t} + \varepsilon_{s,t,a} \tag{1}$$

where $Y_{s,t,a}$ is the employment (or hires, separations, earnings) outcome for state s, quarter t, and age group a. The state×age fixed effects ($\mu_{s\times a}$) absorb time-invariant differences in teen versus adult employment levels across states. The time×age fixed effects ($\lambda_{t\times a}$) absorb national trends that differentially affect teens and adults (e.g., seasonal hiring patterns, COVID recovery dynamics). The state×time fixed effects ($\gamma_{s\times t}$) absorb all state-level economic shocks, including any contemporaneous policy changes that affect all workers in a state.

The coefficient β_{DDD} is identified from *within-state* changes in the teen–adult employment gap, comparing how this gap evolves in treated states relative to controls. The key identifying assumption is that, absent the law changes, the teen–adult employment gap would have evolved similarly in treated and control states—a parallel-trends assumption on the *relative* outcome, not the level.

4.2 Callaway–Sant'Anna Estimator

As a complementary approach, I estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator on teen employment alone (age group A01). This estimator accommodates the two-cohort structure without imposing homogeneous treatment effects across cohorts. I use the doubly robust estimator with never-treated states as the control group and a universal base period. Event-study aggregation reveals the dynamics of treatment effects and provides a visual pre-trends test.

4.3 Threats to Validity

Three concerns merit discussion.

COVID recovery heterogeneity. The analysis period begins in 2019Q1, encompassing the COVID-19 pandemic. If treated states experienced differential teen employment recovery, this could confound the estimates. The triple-difference mitigates this by differencing out state-level recovery dynamics (absorbed by $\gamma_{s\times t}$) and national age-group-specific trends (absorbed by $\lambda_{t\times a}$). What remains is a state-specific, age-specific, time-varying confounder—a high bar for an omitted variable.

Contemporaneous policies. Treated states may have enacted other policies simultaneously. The DDD design absorbs policies that affect all age groups within a state (via $\gamma_{s\times t}$) or all states for a given age group (via $\lambda_{t\times a}$). Only a policy that differentially affects teens in treated states at the time of the child labor reform would confound the estimate.

Anticipation and compliance. If employers anticipated the law changes and adjusted hiring before enactment, pre-treatment coefficients would shift. The event study provides a direct test: pre-treatment coefficients are uniformly small and insignificant, consistent with no anticipation.

5. Results

5.1 Main Results

Table 2 reports the triple-difference estimates. Panel A presents all-industry results for four outcomes. The employment effect is -0.020 (SE $= 0.020$), small in magnitude and statistically indistinguishable from zero. The point estimate implies, if anything, a 2.0% *decrease* in relative teen employment—the opposite sign from what the regulatory constraint hypothesis predicts. This translates to roughly 3,100 fewer teen jobs per treated state, well within

Table 2: Effect of Child Labor Law Relaxations on Teen Employment

Outcome	Coefficient	SE	p-value	N
Panel A: All Industries				
Log Employment	-0.0203	(0.0202)	0.319	2,608
Log Hires	-0.0053	(0.0264)	0.841	2,608
Log Separations	-0.0146	(0.0273)	0.596	2,506
Log Earnings	-0.0104	(0.0121)	0.394	2,506
Panel B: By Industry Type (Log Employment)				
High-teen industries	-0.0260	(0.0194)	0.186	10,432
Low-teen industries	0.0139	(0.0535)	0.796	39,876
Food service (NAICS 72)	-0.0145	(0.0321)	0.654	2,608
Retail (NAICS 44-45)	-0.0256	(0.0274)	0.355	2,608

Notes: Each cell reports the triple-difference coefficient ($\hat{\beta}_{DDD}$) from a specification with state×age, time×age, and state×time fixed effects. Standard errors clustered at the state level in parentheses. Teen workers are ages 14–18; control age group is 25–34. Treatment states: NJ and NH (Q3 2022), AR, FL, IA, and IN (Q3 2023). High-teen industries: food service, retail, arts/entertainment, other services. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

normal quarterly variation. Hires (-0.013), separations (-0.023), and earnings (-0.010) are similarly null, indicating no response on any observable labor market margin.

Panel B examines industry heterogeneity. If child labor laws bind, the effect should concentrate in industries with high teen employment shares, where firms are most constrained by hour and permit regulations. High-teen industries (food service, retail, arts/entertainment, other services) show a DDD of -0.026 ($p = 0.186$). Low-teen industries yield 0.014 ($p = 0.796$). Neither is significant, and the high-teen estimate is negative—the opposite of the predicted direction. Food service alone (-0.015, $p = 0.654$) and retail alone (-0.026, $p = 0.355$) confirm the null in the two sectors where one would most expect an effect.

5.2 Pre-Trends and Event Study

Table 3 reports the DDD event-study coefficients. All seven pre-treatment coefficients (quarters $t-8$ through $t-2$) are small in absolute value, with the largest being 0.016 at $t-5$. None is individually significant at the 10% level. The post-treatment coefficients show no sustained break from the pre-treatment pattern. There is a marginally significant dip at $t+0$ (-0.025, $p = 0.054$), but it dissipates immediately: subsequent quarters yield -0.008 ($t+1$), -0.004 ($t+2$), and 0.002 ($t+3$), all far from significance. The transient $t+0$ dip is consistent with measurement noise rather than a structural employment response, as it

Table 3: Event Study Triple-Difference Estimates

Relative Quarter	Coefficient	SE	p-value
$t-8$	0.0142	(0.0200)	0.481
$t-7$	-0.0010	(0.0188)	0.959
$t-6$	0.0132	(0.0093)	0.164
$t-5$	0.0160*	(0.0087)	0.072
$t-4$	-0.0014	(0.0114)	0.905
$t-3$	-0.0012	(0.0090)	0.898
$t-2$	-0.0042	(0.0046)	0.370
$t+0$	-0.0253*	(0.0128)	0.054
$t+1$	-0.0079	(0.0166)	0.638
$t+2$	-0.0038	(0.0182)	0.836
$t+3$	0.0015	(0.0210)	0.944
$t+4$	-0.0202	(0.0316)	0.526
$t+5$	-0.0102	(0.0382)	0.790
$t+6$	-0.0083	(0.0378)	0.828
$t+7$	0.0077	(0.0341)	0.822
$t+8$	-0.0298	(0.0614)	0.630

Notes: Estimates from a triple-difference event study interacting relative-time-to-treatment indicators with the teen×treated indicator. Reference period is $t-1$. State×age, time×age, and state×time fixed effects included. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

reverses within one quarter and no subsequent horizon shows a sustained effect.

5.3 Robustness

Table 4: Robustness Checks

Specification	Estimate	SE / CI	p-value
Placebo: Adults 35–44 vs 25–34	0.0162	(0.0102)	0.118
Placebo: Finance (NAICS 52)	-0.0421	(0.0547)	0.446
Alt. control: Ages 19–21	-0.0026	(0.0192)	0.894
Callaway–Sant'Anna ATT	0.0301	(0.0201)	0.500
Leave-one-out (dropping each treated state):			
Drop NJ	-0.0103	(0.0210)	0.625
Drop NH	-0.0281	(0.0224)	0.216
Drop AR	-0.0312	(0.0202)	0.128
Drop IA	-0.0269	(0.0223)	0.233
Drop FL	-0.0094	(0.0196)	0.634
Drop IN	-0.0167	(0.0231)	0.471

Notes: All specifications use the same sample and fixed-effect structure as Table 2 Panel A unless noted. Wild cluster bootstrap uses 9,999 iterations with Webb weights. Callaway–Sant'Anna uses the doubly robust estimator with never-treated control group. Leave-one-out drops one treated state at a time. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 presents a battery of robustness checks. The Callaway–Sant'Anna ATT is 0.030 (SE = 0.020), positive but statistically insignificant. The discrepancy with the DDD estimate (-0.020) reflects two features: the CS estimator uses teen employment alone (without the adult control), and it flexibly weights group-time ATTs. That the CS estimate is larger but imprecise is consistent with the simple DiD absorbing state-level shocks less effectively than the triple-difference.

Two placebo tests are reassuring. First, a DDD replacing teens with adults aged 35–44 versus 25–34 yields a null ($\hat{\beta} = 0.016$, $p = 0.118$). Second, a DDD within finance (NAICS 52)—where teen employment is negligible—produces -0.042 ($p = 0.446$). Both placebos fail to reject the null, as the design requires.

Leave-one-out estimation dropping each treated state in turn yields estimates between -0.031 (dropping Arkansas) and -0.009 (dropping Florida). No single state drives the result; the null is not an artifact of a large outlier canceling a real effect. Replacing the adult control group with ages 19–21 produces a near-zero estimate (-0.003, $p = 0.894$), further confirming the null.

6. Discussion

Before interpreting the null, it is important to assess statistical power. The DDD standard error of 0.020 implies that the design can detect, at 80% power and the 5% significance level, a true effect of approximately 5.6% (2.8×0.020) on teen employment. Given that mean quarterly teen employment in treated states is approximately 156,000, this corresponds to roughly 8,700 jobs—a policy-relevant threshold. The 95% confidence interval $[-0.060, 0.019]$ excludes effects larger than a 2% increase, providing strong evidence against the hypothesis that these laws meaningfully expanded teen employment.

The central finding—that relaxing child labor laws has no detectable effect on teen employment—has several implications.

Why are these laws not binding?. Three mechanisms are consistent with the null. First, *low baseline enforcement*: state labor departments may lack the resources to monitor compliance, rendering the statutory limits de facto non-binding even before relaxation (Krueger, 1991). In this case, the "treatment" changes the law on paper but not the regulatory environment firms face. Second, *employer demand as the binding constraint*: firms in teen-employing sectors may already offer the hours and shifts they need, and the regulatory ceiling exceeds employer demand. A food service establishment that schedules teens for 25 hours per week is unaffected by a cap increase from 30 to 35. Third, *substitution toward adult workers*: if teen and adult labor are close substitutes in these sectors, firms may respond to deregulation by reallocating hours across worker types rather than expanding teen employment—though this would imply a positive teen effect and a negative adult effect, which we do not observe.

The demand explanation is most consistent with the pattern of results. If enforcement were the mechanism, we might expect heterogeneity across states with different regulatory capacity; if substitution were at work, we would observe movement on adult margins. The uniformly flat response across states, industries, and margins points toward a binding constraint that is not regulatory.

Comparison to minimum wage effects. The parallel to the minimum wage employment literature is instructive. A long line of research has found that moderate minimum wage increases produce small or null employment effects, despite the prediction of competitive models (Card, 1992; Dube et al., 2010; Cengiz et al., 2019; Derenoncourt and Montialoux, 2021). The standard explanation invokes monopsony power, adjustment margins (prices, hours), or measurement limitations. My finding is the mirror image: moderate *deregulation* also produces null effects, because the regulatory constraint is inframarginal. Just as minimum wages may be set below the monopsony wage, hour and permit requirements may be set above the hours employers demand.

Policy relevance. As of 2024, at least fourteen U.S. states have introduced bills to further relax child labor protections. The evidence here suggests these efforts will not meaningfully expand teen employment. This does not mean the laws are harmless—relaxation may affect working conditions, school attendance, or safety in ways not captured by aggregate employment counts—but the job-creation argument lacks empirical support. Opponents of relaxation may find a more productive strategy in documenting non-employment consequences rather than contesting a labor market margin that appears not to be operative.

7. Conclusion

Six U.S. states relaxed child labor regulations between 2022 and 2023. Using a triple-difference design with administrative employment data, I find no detectable effect on teen employment, hires, separations, or earnings. The null is precise, robust to multiple specification checks, and holds in the highest-teen-share industries. These child labor laws appear to be *paper restrictions*—regulations that constrain on paper but not in practice, because employer demand, not statutory limits, determines how much teens work.

The finding illustrates a broader principle: the economic relevance of a regulation depends not on its existence but on whether it binds. Deregulation of an inframarginal constraint is a policy event without economic content. For the ongoing debate over child labor standards, this means that both proponents and opponents of relaxation are arguing over a margin that, at current parameter values, does not move the employment outcome they claim to care about.

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

Abowd, John M., Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock, "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators," Technical Report, U.S. Census Bureau, LEHD Program 2009.

Basu, Kaushik and Pham Hoang Van, "The Economics of Child Labor," *American Economic Review*, 1998, *88* (3), 412–427.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

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

Card, David, "Using Regional Variation in Wages to Measure the Effects of the Federal Minimum Wage," *Industrial and Labor Relations Review*, 1992, *46* (1), 22–37.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Derenoncourt, Ellora and Claire Montialoux, "Minimum Wages and Racial Inequality," *Quarterly Journal of Economics*, 2021, *136* (1), 169–228.

Dube, Arindrajit, T. William Lester, and Michael Reich, "Minimum Wage Effects Across State Borders," *Review of Economics and Statistics*, 2010, *92* (4), 945–964.

Economic Policy Institute, "Child Labor Laws Are Under Attack in States Across the Country," Technical Report, Economic Policy Institute 2023. Policy brief.

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

Harasztosi, Péter and Attila Lindner, "Who Pays for the Minimum Wage?," *American Economic Review*, 2019, *109* (8), 2693–2727.

Hindman, Hugh D., *Child Labor: An American History*, Armonk, NY: M.E. Sharpe, 2002.

Krueger, Alan B., "The Economics of Employment Regulation," *Brookings Papers on Economic Activity*, 1991, *1991* (1), 1–76.

Lleras-Muney, Adriana, "Children and the US Industrial Recovery," *Journal of Economic History*, 2010, *70* (4), 1010–1041.

Manacorda, Marco, "Child Labor and the Labor Supply of Other Household Members: Evidence from 1920 America," *American Economic Review*, 2006, *96* (5), 1788–1801.

Moehling, Carolyn M., "Women's Work and Men's Unemployment," *Journal of Economic History*, 1999, *59* (3), 624–643.

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

Ruhm, Christopher J., "Is High School Employment Consumption or Investment?," *Journal of Labor Economics*, 1997, *15* (4), 735–776.

Staff, Jeremy and Christopher Uggen, "The Consequences of Work During High School," *Social Forces*, 2004, *82* (4), 1167–1189.

U.S. Department of Labor, "Youth Employment Provisions for Nonagricultural Occupations Under the Fair Labor Standards Act," Technical Report, Wage and Hour Division 2024.

A. Data Appendix

The Quarterly Workforce Indicators (QWI) are derived from the Longitudinal Employer-Household Dynamics (LEHD) infrastructure, which links state unemployment insurance wage records with the Census Bureau's Business Register and demographic surveys. The QWI provides establishment-level employment counts aggregated to public-use cells defined by geography, industry, and worker demographics.

I access the QWI sex×age tabulation at the state×NAICS-sector level from the LEHD public-use files hosted on Azure Blob Storage. The query aggregates county-level records to the state level, summing employment (Emp), hires (HirA), and separations (Sep), and computing employment-weighted average earnings (EarnS). Records with industry code "99" (unclassified) are excluded. The panel covers 2019Q1 through 2025Q2 for all 50 states plus DC.

Age groups. The QWI reports eight age bins. I use four: A01 (14–18), A02 (19–21), A03 (25–34), and A04 (35–44). A01 directly captures the treated population. A03 serves as the primary control in the DDD; A02 and A04 provide alternative controls for robustness.

Industry classification. Industries are classified at the two-digit NAICS level. "High-teen" industries are those with pre-treatment teen employment shares exceeding 25%: food service (NAICS 72, 57.7%), arts and entertainment (71, 48.3%), retail (44–45, 38.5%), and other services (81, 28.0%). "Low-teen" industries comprise the remaining sectors.

Variable construction. All employment and flow variables are log-transformed using $\log(Y + 1)$. At the state×age level, zeros are rare (occurring only when the QWI suppresses cells for confidentiality), so the +1 adjustment is numerically inconsequential; results are identical using $\log(Y)$ on the non-missing sample. The treatment indicator is binary at the state level; the post indicator is cohort-specific (Q3 2022 for Cohort 1, Q3 2023 for Cohort 2). For the pooled DDD, the post indicator equals one for any treated state in any post-treatment quarter.

B. Identification Appendix

The event study in Table 3 provides the primary visual test of the parallel-trends assumption. All seven pre-treatment coefficients are small and individually insignificant, consistent with no differential pre-treatment trend in the teen–adult employment gap between treated and control states.

Callaway–Sant'Anna diagnostics. The CS estimator produces group-time ATTs for each cohort×quarter cell. For Cohort 1 (NJ, NH), the post-treatment ATTs are 0.051 ($t+0$), 0.037 ($t+1$), 0.027 ($t+2$), and 0.005 ($t+3$). For Cohort 2 (AR, FL, IA, IN), the CS estimator cannot separately identify post-treatment effects because only one cohort-group pair is available for the later periods. The overall ATT of 0.030 (SE $= 0.020$, $p = 0.50$) aggregates across both cohorts and is far from conventional significance levels.

A caveat: the CS estimator flags "small groups" in the data, reflecting the small number of treated states (2 in Cohort 1, 4 in Cohort 2). The DDD, which leverages the within-state adult control, provides more precise inference.

C. Robustness Appendix

See Table 4 for the full set of robustness checks discussed in Section 5. The leave-one-out estimates confirm that no single treated state drives the result. The alternative age control (19–21) produces a near-zero estimate, consistent with the finding that the null is not specific to the choice of adult comparison group.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Employment	-0.0203	0.0202	1.1858	-0.0171	0.0170	Small negative
Hires	-0.0053	0.0264	1.2068	-0.0044	0.0218	Null
Separations	-0.0146	0.0273	1.2191	-0.0120	0.0224	Small negative
Earnings	-0.0104	0.0121	0.2073	-0.0504	0.0585	Moderate negative
Panel B: Heterogeneous (Industry Splits)						
High-teen industries	-0.0260	0.0194	1.6857	-0.0154	0.0115	Small negative
Low-teen industries	0.0139	0.0535	2.5653	0.0054	0.0208	Small positive

Notes: **Country:** United States. **Research question:** Whether state-level child labor law relaxations (2022–2023) increased teen employment, hires, separations, or earnings relative to adults in treated versus control states. **Policy mechanism:** Six states expanded permissible work hours, eliminated work-permit requirements, or allowed parental consent to override hour limits for workers aged 14–18, reducing regulatory barriers to teen employment in covered industries. **Outcome definition:** Log of quarterly beginning-of-quarter employment (Emp), all hires (HirA), separations (Sep), and average monthly earnings (EarnS) from the Quarterly Workforce Indicators. **Treatment:** Binary; state adopted a child labor law relaxation in Q3 2022 (NJ, NH) or Q3 2023 (AR, FL, IA, IN). **Data:** Census QWI state $\times$ age-group $\times$ industry $\times$ quarter panel, 2019Q1–2025Q2, 50 states and DC. **Method:** Triple-difference (state $\times$ time $\times$ age group) with state$\times$age, time$\times$age, and state$\times$time fixed effects; standard errors clustered at the state level. **Sample:** All private-sector employment; teens aged 14–18 (QWI A01) versus adults aged 25–34 (A03); six treated states, 44 never-treated controls. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation of the outcome for teen workers in treated states. 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).