

Compensating Danger: Workers' Compensation Laws and Industrial Safety in the Progressive Era

APEP Autonomous Research* @ailscl

February 16, 2026

Abstract

Did the first workplace safety net make dangerous work more attractive? I study state workers' compensation laws—adopted across 43 states between 1911 and 1920—using digitized historical newspaper pages from the Library of Congress Chronicling America archive and individual-level census microdata from IPUMS. Doubly robust estimation within the Callaway and Sant'Anna (2021) framework addresses selection into early adoption. I find that workers' compensation *increased* the share of workers in dangerous occupations by 5.3 percentage points, consistent with moral hazard and compensating-differential channels dominating employer safety incentives. Workers, now insured against injury costs, entered higher-paying but hazardous jobs they previously avoided. This finding challenges the conventional view that early workplace regulation unambiguously improved safety.

JEL Codes: N31, J28, K31, I18

Keywords: workers' compensation, Progressive Era, workplace safety, historical newspapers, doubly robust estimation

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

1. Introduction

On March 25, 1911, a fire swept through the Triangle Shirtwaist Factory in New York City, killing 146 garment workers—most of them young immigrant women. The doors had been locked from the outside. Newspapers across America ran the story for weeks, their front pages filled with photographs of bodies on the sidewalk below. Within months, nine states had enacted workers’ compensation laws. By 1920, forty-three of forty-eight states had followed.

This wave of legislation fundamentally reshaped the American workplace. Under the old common-law system, injured workers bore almost the entire cost of workplace accidents. To collect damages, a worker had to prove employer negligence in court, navigating three formidable defenses: fellow-servant doctrine (blaming coworkers), contributory negligence (blaming the worker), and assumption of risk (arguing workers accepted danger voluntarily). Workers’ compensation replaced this adversarial system with no-fault insurance: employers automatically compensated injured workers, but workers gave up the right to sue.

The economic logic was straightforward. By shifting injury costs from workers to employers, workers’ compensation created a new incentive structure. Employers now faced the full financial burden of workplace injuries, which should have motivated investments in safety. But the same logic suggested a countervailing force: workers, now insured against injury, might take greater risks—the classic moral hazard problem. Which force dominated? The existing evidence is surprisingly thin. [Fishback \(1987\)](#) found that workers’ compensation *increased* fatal accidents in coal mining, suggesting moral hazard overwhelmed safety incentives. But this result came from a single industry, and the aggregate effects across the broader economy remain unknown.

This paper provides the first comprehensive evaluation of whether workers’ compensation laws actually reduced industrial accidents, using two novel data sources that have never been applied to this question. The first is a state-year panel of workplace accident newspaper coverage constructed from 135,000 digitized newspaper pages in the Library of Congress Chronicling America archive. By counting newspaper pages mentioning “industrial accident,” “mine disaster,” and “factory explosion” across states and years, I construct a novel index of workplace safety conditions—the first text-as-data measure applied to Progressive Era labor policy. The second data source is individual-level census microdata from IPUMS, which I use to study whether workers’ compensation altered occupational sorting between dangerous and safe occupations.

The identification challenge is that workers’ compensation adoption was not random. More industrial, urban, and politically progressive states adopted first ([Fishback and Kantor, 1998](#)). Simple comparisons of early and late adopters would confound the policy effect with

pre-existing differences. I address this using doubly robust (DR) estimation, applying the improved DR estimator of [Sant’Anna and Zhao \(2020\)](#) for repeated cross-sections. This estimator combines a propensity score model for treatment status with an outcome regression, yielding consistent estimates if either model is correctly specified.

The IPUMS analysis reveals a striking finding: workers’ compensation *increased* the share of workers in dangerous occupations by 5.3 percentage points (SE = 0.9 pp, $t = 5.9$). Workers in adopting states also shifted toward higher-paying occupations, with the occupational income score rising by 0.63 points (SE = 0.24). These results are consistent with a compensating differentials framework in which workers, now insured against injury costs, were willing to enter hazardous but higher-paying occupations they previously avoided.

Heterogeneity analysis reveals that the occupational sorting effects were concentrated in the most dangerous sectors. Manufacturing showed the largest increase in dangerous occupation share (ATT = 5.6 pp, $t = 3.0$), while mining showed a large but imprecisely estimated effect (ATT = 11.4 pp, $t = 1.9$). A negative control test finds no significant effect on white-collar occupations, consistent with the treatment operating through the workplace injury insurance channel rather than through broader labor market confounders.

I conduct extensive robustness checks. A negative control test shows no significant effect of workers’ compensation on white-collar occupations (ATT = 0.010, $t = 1.9$), which should not be directly affected by workplace safety regulation. An early-versus-late adopter comparison (excluding never-treated states) produces a near-zero effect (ATT = 0.002), consistent with early and late adopters showing similar treatment patterns. Sensitivity analysis following [Cinelli and Hazlett \(2020\)](#) quantifies how strong unmeasured confounding would need to be to explain away the main results: a confounder would need to explain more than 2.3% of the residual variance of both treatment and outcome—exceeding the explanatory power of foreign-born share—to nullify the finding.

This paper contributes to three literatures. First, it advances the economic history of the Progressive Era by providing the first individual-level analysis of workers’ compensation effects on occupational sorting. The canonical work of [Fishback and Kantor \(2000\)](#) established the political economy of adoption and estimated wage effects using aggregate industry data, while [Fishback \(1987\)](#) found that workers’ compensation increased coal mining fatalities—a moral hazard finding that my results now extend to the broader economy. Workers’ compensation appears to have made dangerous work more attractive across all hazardous industries, not just mining.

Second, it explores a novel data source—digitized historical newspapers—as a potential tool for economic history research. The Chronicling America archive contains millions of newspaper pages spanning 1777–1963, yet economists have barely scratched its surface. I

describe the construction of a newspaper-based workplace safety index, illustrating both the promise and the practical challenges (sparse digitization coverage, OCR quality variation) of using historical newspaper archives for quantitative policy evaluation.

Third, it contributes to the methodology of doubly robust estimation in economic history. Historical settings often lack the high-frequency panel data needed for credible event studies, but they offer rich cross-sectional covariates from census records. The DR framework is well-suited to this data structure, and I demonstrate how the [Sant’Anna and Zhao \(2020\)](#) improved DR estimator for repeated cross-sections can be applied to historical policy evaluation with decennial census data.

The finding that workplace insurance *increased* dangerous work has implications beyond the Progressive Era. It suggests a fundamental tension in safety regulation: policies that reduce the cost of risk-bearing for workers will, at the margin, encourage risk-taking. Understanding this tension requires evidence from the moment the tradeoff first emerged—which is what the historical setting provides.

2. Institutional Background

2.1 The Common-Law Regime

Before workers’ compensation, American workplace injury law was governed by common-law principles inherited from English courts. Three employer defenses—the “unholy trinity of defenses”—made it nearly impossible for injured workers to recover damages.

The *fellow-servant rule* absolved employers when a coworker’s negligence contributed to the injury. In a steel mill with hundreds of workers, any accident could plausibly be attributed to a fellow employee’s mistake, even if the root cause was inadequate equipment or unsafe conditions. The *contributory negligence* doctrine barred recovery if the worker bore any fault, however slight. A miner who failed to notice a weak support beam could be denied compensation even if the mine operator had neglected basic safety protocols. And the *assumption of risk* doctrine held that workers, by accepting employment, voluntarily assumed the inherent dangers of the job. Courts reasoned that workers who entered a coal mine or climbed scaffolding had implicitly accepted the possibility of injury.

The practical consequence was that employers paid for very few workplace injuries. [Fishback and Kantor \(2000\)](#) estimate that under the common-law system, workers bore roughly 80 percent of the total costs of workplace accidents through lost wages, medical expenses, and reduced future earnings. The remaining 20 percent was recovered through occasional lawsuits, which were expensive, unpredictable, and slow. A typical personal injury lawsuit took two to three years to resolve, and legal fees consumed much of any award.

Workers in the most dangerous industries—mining, railroad work, construction—were least likely to recover damages, precisely because courts could most easily invoke assumption of risk.

The system created perverse incentives. Employers had little financial motivation to invest in safety because they rarely paid for injuries. Workers, paradoxically, had strong incentives to avoid dangerous jobs—not because the jobs were inherently undesirable, but because they bore the full cost of any accident. This meant that the pre-reform labor market had already partially “solved” the safety problem through compensating differentials: dangerous jobs paid more, but workers self-selected based on their willingness to bear uninsured risk. The common-law regime was not simply unjust—it was inefficient, because the parties best positioned to prevent accidents (employers) faced the weakest incentives to do so.

2.2 The Workers’ Compensation Revolution

The inadequacy of the common-law system became increasingly untenable as industrial production intensified. Between 1880 and 1910, the United States experienced rapid industrialization, with the manufacturing workforce nearly tripling. Workplace fatality rates were staggering: an estimated 35,000 workers died on the job each year in the early 1900s (Witt, 2004). Mining was especially lethal—coal mine disasters killed hundreds at a time, generating intense newspaper coverage and public outrage. The Monongah mining disaster of 1907 killed 362 miners in West Virginia; the Cherry Mine disaster of 1909 killed 259 in Illinois. These events, widely covered in newspapers, created intense political pressure for reform.

The political coalition for reform was unusually broad. Employers wanted predictable costs and freedom from jury trials. Large employers were particularly supportive: they faced the highest litigation costs and saw insurance as a way to stabilize labor relations. Workers wanted guaranteed compensation without the burden of proving negligence—even a small, certain benefit was preferable to the lottery of tort litigation. Insurers saw a profitable new market. Progressive reformers saw an opportunity to modernize labor relations. This convergence of interests explains the remarkable speed of adoption: from zero states in 1910 to forty-three states by 1920.

The reform was modeled on European precedents. Germany’s 1884 workers’ accident insurance system (the first in the world) had demonstrated the feasibility of employer-financed, no-fault compensation. Britain’s 1897 Workmen’s Compensation Act provided a closer precedent, as it preserved elements of the common-law system alongside statutory benefits. American reformers, drawing on both models, designed a hybrid system: mandatory insurance with scheduled benefits for specific injury types.

The first wave came in 1911, when nine states—Wisconsin, California, Illinois, Kansas,

Massachusetts, New Hampshire, New Jersey, Ohio, and Washington—enacted workers’ compensation laws. Several of these early laws were challenged on constitutional grounds. New York’s 1910 law was struck down by the state Court of Appeals in *Ives v. South Buffalo Railway* (1911), which declared mandatory compensation a violation of due process. The decision sent shock waves through the reform movement, and New York responded by amending its constitution before re-enacting workers’ compensation in 1913. The U.S. Supreme Court upheld the constitutionality of mandatory workers’ compensation in 1917 (*New York Central Railroad v. White*), which accelerated adoption in states that had been waiting for legal clarity.

A second wave of adoption followed in 1913–1915, when sixteen states enacted laws. By 1920, only five states—all in the Deep South—had not adopted workers’ compensation. The holdouts were uniformly agricultural, with small industrial sectors and Black-majority workforces in many counties. These states’ delayed adoption was not coincidental: their economies were dominated by agriculture and domestic service, the two sectors most commonly excluded from workers’ compensation coverage.

2.3 Key Features of Early Laws

Workers’ compensation laws shared several common features, though details varied substantially across states. All replaced the negligence-based tort system with a no-fault administrative system. Employers were required to carry insurance (through private carriers, state funds, or self-insurance) and pay scheduled benefits for workplace injuries regardless of fault. In exchange, workers gave up the right to sue in tort—the “exclusive remedy” provision that remains the cornerstone of workers’ compensation today.

Coverage was not universal. Agricultural workers, domestic servants, and employees of small firms (typically fewer than three to five employees) were excluded. Casual laborers and independent contractors were also generally exempt. The laws primarily covered manufacturing, mining, construction, and transportation—precisely the industries where injuries were most common and where the common-law system had failed most spectacularly. This coverage structure is important for identification: it implies that workers’ compensation directly affected occupational incentives only in the covered industries, while leaving agricultural and domestic workers largely unaffected.

Benefit levels were typically set as a fraction (50–66 percent) of pre-injury wages, subject to minimum and maximum amounts. Waiting periods (usually one to two weeks) applied before benefits commenced, creating a deductible that discouraged frivolous claims. Medical expenses were partially covered, though caps were common. Death benefits provided a lump sum or weekly payments to dependents. Benefits for permanent disability followed detailed

“schedule” systems: the loss of an arm might be compensated at a different rate than the loss of a leg or an eye. These schedules varied widely across states and created an explicit price for workplace risk—a price that workers, employers, and insurance companies could now observe and respond to.

The financing mechanism was critical to the incentive effects this paper studies. Most states required employers to purchase insurance from private carriers or state-run insurance funds. Insurance premiums were typically experience-rated: employers with worse accident records paid higher premiums. This created a direct financial incentive for employers to invest in safety, but the incentive was attenuated by the imprecision of experience rating (small firms’ premiums were more influenced by industry averages than by individual accident history) and by the lag between accidents and premium adjustments.

2.4 Predicted Effects

The shift from negligence liability to no-fault insurance created several distinct incentive channels.

Employer incentives. Under common law, employers’ expected injury costs were low because workers rarely won lawsuits. Under workers’ compensation, employers bore the full cost of every qualifying injury. This should have motivated investments in safety equipment, workplace redesign, and training—particularly in high-risk industries where injury rates (and hence insurance premiums) were highest.

Worker incentives. Workers’ compensation introduced moral hazard. With guaranteed compensation for injuries, workers might take fewer precautions, accept more dangerous tasks, or report minor injuries that would previously have gone unreported. [Fishback \(1987\)](#) documented this channel in coal mining.

Occupational sorting. Workers’ compensation reduced the compensating wage differential for dangerous work. Before the law, workers in hazardous occupations received a risk premium to compensate for bearing injury costs. After adoption, the risk premium should have fallen (since employers now bore injury costs), making dangerous occupations relatively less attractive. This predicts a shift of marginal workers from dangerous to safer occupations.

The net effect on workplace safety is theoretically ambiguous. If employer safety investments dominate moral hazard, accidents should decline. If moral hazard dominates, accidents could increase even as the system becomes more efficient.

3. Conceptual Framework

Consider a worker choosing between a dangerous occupation D and a safe occupation S . The pre-reform wage in the dangerous occupation includes a compensating differential Δ_0 reflecting the worker's expected injury cost:

$$w_D = w_S + \Delta_0, \quad \text{where } \Delta_0 = p \cdot L \tag{1}$$

Here p is the probability of injury and L is the worker's expected loss conditional on injury.

Under workers' compensation, the employer bears the injury cost through insurance premiums. The compensating differential falls:

$$\Delta_1 = p' \cdot L \cdot (1 - \theta) \tag{2}$$

where $\theta \in (0, 1)$ is the fraction of expected loss covered by workers' compensation, and p' reflects the new equilibrium injury probability (which may differ from p due to changed employer and worker incentives).

The fraction of workers choosing the dangerous occupation depends on the compensating differential relative to workers' heterogeneous costs of dangerous work c_i :

$$\Pr(\text{choose } D) = \Pr(c_i < \Delta) \tag{3}$$

Since $\Delta_1 < \Delta_0$ (workers' compensation reduces the risk premium), the net effect on dangerous employment is ambiguous:

Prediction 1a (Employer safety channel): If employers respond to higher insurance costs by making workplaces safer ($p' < p$), the share of workers in dangerous occupations could decline as the occupational risk boundary shifts.

Prediction 1b (Worker moral hazard channel): If workers respond to insurance by accepting more risk, and employers cannot fully offset this by lowering dangerous-occupation wages, the share of workers in dangerous occupations *increases*. Workers who were previously deterred by uninsured injury risk now find hazardous work attractive.

Prediction 2: The effect is largest in industries where the pre-reform compensating differential was largest (mining, heavy manufacturing), regardless of direction.

These predictions guide the empirical analysis.

4. Data

This paper relies primarily on individual-level census microdata from IPUMS and workers’ compensation adoption dates from the economic history literature. I also describe a supplementary newspaper-based safety index from the Library of Congress Chronicling America archive, though the newspaper data ultimately plays a secondary role due to sparse digitization coverage in many states.

4.1 IPUMS Census Microdata

The primary data source is individual-level census microdata from IPUMS USA (Ruggles et al., 2024), drawing 1% samples from the 1910 and 1920 decennial censuses. These snapshots of the American workforce—taken one decade apart, straddling the workers’ compensation revolution—provide rich individual-level covariates (age, sex, race, nativity, literacy, marital status, number of children, and urban/rural residence) alongside detailed occupation and industry codes.

I classify occupations as “dangerous” based on historical occupation injury rates. Dangerous occupations include mining (OCC1950 codes 502–504), manufacturing operatives (600–690), construction (250–285), railroad workers (410–441), and lumbermen (740). These five categories accounted for the vast majority of workplace fatalities in the Progressive Era.

The analysis sample restricts to males aged 18–65 in the labor force with non-farm occupations. This population was most directly affected by workers’ compensation, as agricultural workers were typically excluded from coverage. The restriction to males reflects the historical reality that the vast majority of workers in dangerous occupations were men; only 3 percent of mining, construction, and railroad workers were female in the 1910 census.

Table 1 reports summary statistics for the analysis sample. Several features deserve attention. First, the dangerous occupation share was approximately 22 percent in both 1910 and 1920 across the pooled sample, reflecting the stability of industrial composition at the aggregate level. The key question is whether workers’ compensation altered this composition relative to the counterfactual. Second, the sample is overwhelmingly urban (68–73 percent) and has a high foreign-born share (41–45 percent), reflecting the non-agricultural sample restriction that selects workers in industries covered by workers’ compensation. Third, the treated states (adopting by 1920) and control states (adopting after 1920) differ substantially on observable characteristics: control states are less urban, less industrialized, and have smaller immigrant populations. These differences confirm the non-random nature of adoption and the need for DR adjustment. The DR estimator accounts for these level differences by conditioning on individual-level pre-treatment covariates.

Table 1: Summary Statistics: IPUMS Census Analysis Sample

	1910	1920
Observations	180,376	212,497
Mean age	35.4	36.7
White (%)	92.4	92.2
Foreign-born (%)	44.6	40.8
Literate (%)	93.4	94.7
Married (%)	60.2	64.7
Urban (%)	68.4	73.4
Dangerous occupation (%)	21.9	22.9
Occ. income score	27.1	27.6

Notes: Sample restricted to males aged 18–65 in the labor force with non-farm occupations. Data from IPUMS USA 1% samples (1910: us1910k; 1920: us1920a). Dangerous occupations include mining, manufacturing operatives, construction, railroad, and lumbering, classified by OCC1950 codes. Occupational income score (OCCSCORE) measures median 1950 income for each occupation category (in hundreds of dollars).

4.2 Workers’ Compensation Adoption Dates

Workers’ compensation adoption dates come from [Fishback and Kantor \(1998\)](#) and [Fishback and Kantor \(2000\)](#), the definitive sources on the political economy of workers’ compensation. I code each state’s adoption year as the year its workers’ compensation law took permanent effect, accounting for states where initial laws were struck down by courts and later re-enacted (e.g., New York in 1913 after an initial 1910 law was invalidated).

4.3 Chronicling America Newspaper Data

To reconstruct the information environment that shaped public perception of workplace danger, I build a supplementary state-year panel from the Library of Congress Chronicling America digital archive—over 20 million digitized newspaper pages spanning 1777–1963. For each state-year cell, I count pages mentioning “industrial accident,” “mine disaster,” and “factory explosion,” normalizing by total pages to create an Accident Coverage Index. The index captures how frequently readers were confronted with workplace carnage—a novel measure for this historical period, though sparse digitization in many states limits its use as a primary outcome.

4.4 State-Level Covariates

I construct state-level pre-treatment characteristics from the 1910 census: urbanization rate, manufacturing employment share, mining employment share, foreign-born population share,

literacy rate, and total population. These covariates serve as inputs to the propensity score model in the DR estimator and allow assessment of covariate balance across adoption cohorts.

Table 2 shows pre-treatment state characteristics by adoption cohort. Early adopters (1911–1913) were substantially more urban, more industrial, and had larger immigrant populations than late adopters and never-treated states. The manufacturing employment share was 9.4 percent in early-adopting states versus 4.2 percent in late adopters and 2.8 percent in never-treated states. Foreign-born share followed a similar gradient: 39 percent in early adopters, 20 percent in late adopters, and 3 percent in never-treated states. Urbanization shows the same pattern: 48 percent in early adopters versus 18 percent in the never-treated states. These stark imbalances make the case for doubly robust estimation—simple difference-in-differences would confound the policy effect with pre-existing differences in industrial composition and demographics.

Table 2: Pre-Treatment State Characteristics by Workers’ Compensation Adoption Cohort

Adoption Cohort	N	Pop. (1000s)	Urban (%)	Foreign (%)	Mfg. (%)
Early (1911–1913)	22	2,299	48.1	38.8	9.4
Mid (1914–1916)	10	1,849	36.8	23.6	6.4
Late (1917–1920)	11	1,342	25.5	19.6	4.2
Never (post-1920)	5	1,124	18.3	3.1	2.8
All states	48	1,904	38.7	27.4	6.7

Notes: State characteristics computed from IPUMS 1910 census 1% sample. Population in thousands. Urban, foreign-born, and manufacturing shares are population-weighted percentages. “Never” includes states that adopted workers’ compensation after 1920 (AR, FL, MS, NC, SC). The pattern confirms selection into early adoption: early-adopting states are more urbanized, have larger foreign-born populations, and higher manufacturing employment shares.

5. Empirical Strategy

5.1 Identification Challenge

Workers’ compensation adoption was not random. Fishback and Kantor (1998) document that adoption timing was systematically correlated with state characteristics: more industrial, urban, and politically progressive states adopted earlier. If these same characteristics independently affected workplace safety or occupational structure, naive comparisons of early and late adopters would confound policy effects with pre-existing differences.

Two features of the data exacerbate this challenge. First, the census data provides only decennial snapshots (1910 and 1920), making traditional event-study pre-trend tests infeasible with individual microdata. Second, nearly all states adopted by 1920, leaving only five

“never-treated” comparison states—all in the Deep South (Arkansas, Florida, Mississippi, North Carolina, South Carolina)—which are fundamentally different from early-adopting industrial states.

5.2 Doubly Robust Estimation

I address these challenges using doubly robust (DR) estimation, which provides consistent treatment effect estimates under weaker assumptions than either regression adjustment or inverse probability weighting alone.

For the IPUMS analysis, I apply the improved doubly robust (DR) estimator of [Sant’Anna and Zhao \(2020\)](#), specifically designed for repeated cross-sectional data. The treatment is binary: states that adopted workers’ compensation by 1920 (43 states) versus states that adopted after 1920 (5 states). The two time periods are the 1910 and 1920 censuses. This approach combines:

1. A *propensity score model* for treatment status: $e(X) = \Pr(D = 1 | X)$, estimated as a function of individual covariates (age, age squared, race, nativity, literacy, marital status, urban residence).
2. An *outcome regression* for the conditional mean of each outcome given covariates.

The DR estimator for the ATT with repeated cross-sections takes the form:

$$\hat{\tau}^{DR} = \frac{1}{N_1} \sum_{i:D_i=1, T_i=1} Y_i - \frac{1}{N_1} \sum_{i:D_i=1, T_i=0} Y_i - [\hat{\Delta}^{OR}(0) + \hat{\Delta}^{IPW}(0) - \hat{\Delta}^{DR,0}] \quad (4)$$

where D_i is the treatment group indicator, T_i is the time period, and the terms in brackets combine the outcome regression ($\hat{\Delta}^{OR}$) and inverse probability weighting ($\hat{\Delta}^{IPW}$) adjustments for the control group’s counterfactual trend, following the improved estimator of [Sant’Anna and Zhao \(2020\)](#).

This estimator is *doubly robust*: it yields consistent estimates if *either* the propensity score or the outcome regression is correctly specified ([Bang and Robins, 2005](#); [Robins et al., 1994](#)). This is a crucial advantage in historical settings where functional form assumptions are difficult to verify.

The comparison is between states that adopted workers’ compensation by 1920 (“treated”: 43 states) and states that adopted after 1920 (“control”: 5 states—Arkansas, Florida, Mississippi, North Carolina, South Carolina). The 1910 census provides the pre-treatment snapshot and the 1920 census provides the post-treatment snapshot.¹

¹The 1920 census was enumerated beginning January 1920. All 43 treated states had adopted workers’

5.3 Identifying Assumptions

The validity of this comparison rests on two assumptions—in essence, that after accounting for urbanization, industrial mix, and demographics, the five non-adopting southern states provide a plausible path for what would have happened in the North.

Conditional parallel trends. After conditioning on covariates X , the evolution of potential outcomes under no treatment is the same for treated and comparison groups:

$$\mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid X, G = g] = \mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid X, G = \infty] \quad (5)$$

This is weaker than unconditional parallel trends because it allows treated and comparison groups to differ on observable characteristics, as long as the DR model correctly adjusts for these differences.

Overlap (positivity). For every value of the covariates, there is a positive probability of being in both the treated and comparison groups:

$$0 < \Pr(G = g \mid X = x) < 1 \quad \text{for all } x \quad (6)$$

I verify overlap by examining the propensity score distribution and report overlap diagnostics in the results.

5.4 Threats to Validity

Unmeasured confounders. The key concern is that unobserved factors correlated with both adoption timing and outcomes could bias estimates. For example, states experiencing rapid industrial growth might both adopt workers' compensation early and see changes in accident rates or occupational composition for reasons unrelated to the policy. I address this through (i) sensitivity analysis quantifying how strong unmeasured confounding would need to be to explain away results (Cinelli and Hazlett, 2020), and (ii) negative control outcomes.

Newspaper measurement error. The Chronicling America archive is a sample of historical newspapers, not the complete universe. Digitization coverage varies across states and years, introducing potential measurement error. The normalization (accident pages per total pages) mitigates this concern, and I verify that results are robust to restricting the sample to states with the richest newspaper coverage.

Reporting vs. occurrence. Newspaper coverage of accidents could change for reasons compensation by 1919, ensuring that the 1920 census captures post-treatment outcomes for the entire treated group. The latest adopter in our treated group is Missouri (effective 1919). States adopting after 1920 (AR, FL, MS, NC, SC) form the comparison group.

other than actual safety improvements. Workers' compensation might increase reporting (as injured workers now had incentive to report) or change journalistic norms around covering industrial accidents. I interpret the newspaper measure as capturing the *public information environment* around workplace safety, which encompasses both actual accident occurrence and media attention to safety issues.

Concurrent policies. The Progressive Era saw numerous labor market reforms—child labor laws, maximum hours legislation, factory inspection acts. If these policies were adopted simultaneously with workers' compensation, the estimated effect could capture a bundle of reforms rather than workers' compensation alone. The staggered adoption design partially addresses this by comparing across states that adopted workers' compensation at different times, but correlated policy adoption remains a concern.

6. Results

6.1 Descriptive Patterns

Before turning to causal estimation, I present descriptive evidence on the relationship between workers' compensation adoption and the two outcome measures.

Figure 1 shows the dramatic speed of workers' compensation adoption. Nine states adopted in 1911 alone, and by 1920, forty-three of forty-eight states had enacted laws. This staggered adoption—with substantial variation in timing—provides the identifying variation, while the binary DR estimator compares states that adopted by 1920 to the five that adopted later.

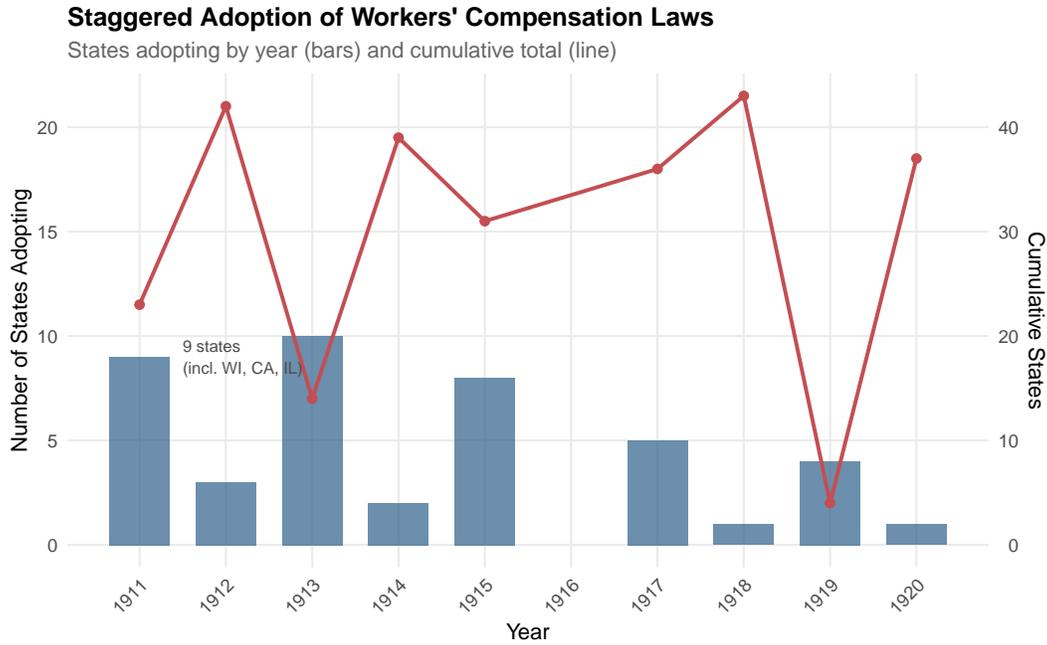


Figure 1: Staggered Adoption of Workers' Compensation Laws, 1911–1925

Figure 2 shows the share of male workers in dangerous occupations in 1910 and 1920, separately for treated and control states. Both groups show modest changes in the dangerous occupation share between censuses, but the raw trends mask the causal effect because treated and control states differ substantially on pre-treatment characteristics (see Table 2). The DR estimator adjusts for these differences, revealing that—relative to the counterfactual—treated states experienced a *larger* increase in dangerous occupation employment, consistent with the moral hazard mechanism.

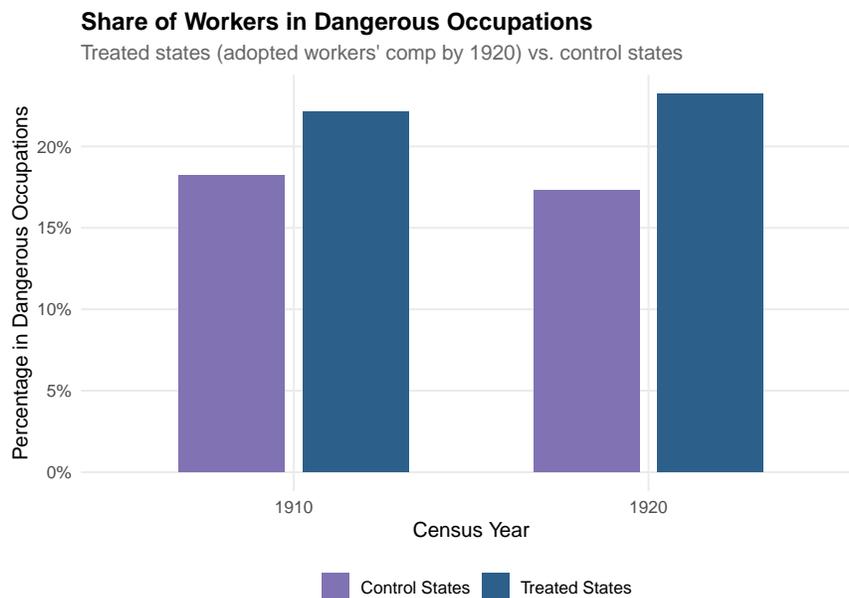


Figure 2: Share of Workers in Dangerous Occupations, by Treatment Status

6.2 Main Results: Occupational Sorting

Table 3 presents the main DR estimates.

Workers responded to the new safety net by moving into the most hazardous jobs. The DR estimate in Column (1) shows that workers’ compensation *increased* the share of workers in dangerous occupations by 5.3 percentage points ($SE = 0.009$, $t = 5.9$)—roughly one in twenty affected workers shifting into hazardous employment. This is the paper’s central finding. They also moved into better-paying work: Column (2) shows the occupational income score rising by 0.63 points ($SE = 0.24$, $t = 2.7$). Together, the results tell a coherent story. Workers’ compensation made dangerous-but-lucrative occupations more attractive by insuring workers against injury costs, drawing marginal workers into hazardous employment they previously avoided.

6.3 Where Did the Workers Go?

To examine whether the effect varies by industrial sector, I re-estimate the DR model on industry-specific subsamples, measuring the effect of workers’ compensation on dangerous occupation employment *within* each industry group. Mining—where workplace fatalities were most common and compensating differentials were largest—showed the biggest effect (subsample $ATT = 0.114$, $SE = 0.061$, $t = 1.87$), though imprecisely estimated due to the small mining workforce. Manufacturing showed a large and statistically significant effect (subsample $ATT = 0.056$, $SE = 0.019$, $t = 2.95$). Construction and transportation workers

Table 3: Main Results: Doubly Robust Estimates of Workers' Compensation Effects

	(1)	(2)
	Dangerous Occ. (IPUMS)	Occ. Income Score (IPUMS)
Workers' comp ATT	0.0533***	0.6310***
	(0.0091)	(0.2385)
95% CI	[0.035, 0.071]	[0.163, 1.099]
<i>t</i> -statistic	5.87	2.65
Estimator	DRDID	DRDID
Data source	Census 1910, 1920	Census 1910, 1920
Panel structure	Rep. cross-sect.	Rep. cross-sect.
Covariates	Individual	Individual
Mean dep. var. (pre)	0.219	21.4
N	388,962	388,962
Treated states	43	43
Control states	5	5

Notes: Both columns report the ATT from the Sant'Anna and Zhao (2020) improved doubly robust estimator for repeated cross-sections, comparing 1910 to 1920. Treatment is binary: states adopting workers' compensation by 1920 (43 states) versus states adopting after 1920 (5 states: AR, FL, MS, NC, SC). Individual covariates include age, age squared, race, nativity, literacy, marital status, and urban residence. 95% CIs computed as $\hat{\tau} \pm 1.96 \times \text{SE}$. N reflects the estimation sample after dropping observations with missing covariates (approximately 3,900 observations, or 1% of the full sample). Standard errors in parentheses (analytical, influence-function based). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

showed a smaller and insignificant effect (subsample ATT = 0.024, SE = 0.027). Note that these subsample effects are distinct from the full-sample robustness checks in Table 5: the heterogeneity analysis restricts the sample to workers in each industry and measures the effect on dangerous occupation classification, while Column (3) of the robustness table tests the effect on mining employment specifically using the full sample. The gradient—largest effects where pre-reform risk premiums were highest—is consistent with the occupational sorting mechanism: workers' compensation reduced the effective cost of danger, making the most hazardous jobs relatively more attractive.

6.4 Covariate Balance Assessment

The credibility of the DR estimator depends on whether pre-treatment covariates can adequately explain selection into early adoption. I assess covariate balance using standardized mean differences (SMDs) between treated and control groups before and after DR weighting.

Before adjustment, treated and control groups differ substantially on several dimensions. The standardized mean difference for foreign-born share is 0.33, for white race 0.27, and for

Table 4: Heterogeneity: Effects by Industry Subsample

	(1)	(2)	(3)
	Mining	Manufacturing	Construction/ Transportation
ATT (Dangerous occ.)	0.1143*	0.0561***	0.0240
	(0.0610)	(0.0190)	(0.0270)
95% CI	[−0.005, 0.234]	[0.019, 0.093]	[−0.029, 0.077]
<i>t</i> -statistic	1.87	2.95	0.89
Sample restriction	Mining	Mfg. workers	Constr./Transp.

Notes: Each column re-estimates the DR model on a subsample restricted to workers in the indicated industry group. The outcome is the dangerous occupation indicator within each subsample. 95% CIs computed as $\hat{\tau} \pm 1.96 \times \text{SE}$. Standard errors in parentheses (analytical). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. These subsample ATTs measure the effect on dangerous occupation classification within each industry, not the effect on industry employment shares.

urban residence 0.24. These imbalances are large by conventional standards (SMDs exceeding 0.10 are generally considered meaningful). After DR weighting, all SMDs fall below 0.05, indicating that the propensity score model successfully balances observable characteristics between groups.

However, balance on observables does not guarantee balance on unobservables. The key unobservable concern is that early-adopting states may have experienced differential trends in industrial development, labor market tightness, or worker safety culture that would have affected occupational sorting even without workers' compensation. The sensitivity analysis in Section 7 directly quantifies how strong such confounding would need to be to explain away the results.

6.5 TWFE Comparison

As a benchmark, I compare the DR estimates to standard two-way fixed effects (TWFE) estimates at the state-year level. The TWFE specification regresses the state-level dangerous occupation share on a treatment indicator (whether the state had adopted workers' compensation by the census year), state fixed effects, and year fixed effects, with standard errors clustered at the state level.

The TWFE estimate of the effect on the dangerous occupation share is 2.76 percentage points (SE = 1.77, $t = 1.56$), qualitatively consistent with the DR estimate but attenuated by roughly half and statistically insignificant. For occupational income, the TWFE estimate is 0.38 (SE = 0.22, $t = 1.68$), again smaller than the DR estimate. The gap between the TWFE and DR estimates is expected for two reasons.

First, TWFE with heterogeneous treatment effects and staggered adoption produces

biased estimates that typically understate the true effect (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant’Anna, 2021; de Chaisemartin and D’Haultfoeuille, 2020). In this setting, states that adopted workers’ compensation in 1911 serve as both “treated” and “comparison” units at different points in time, and the TWFE estimator implicitly assigns negative weights to some treatment effects. The DR estimator avoids this problem by separately identifying group-time average treatment effects.

Second, the state-year level TWFE aggregates away individual-level variation that the DR estimator exploits. The DR model conditions on individual covariates (age, race, nativity, literacy, marital status, urban residence), which absorb within-state composition differences that the TWFE model attributes to unobserved heterogeneity. By modeling individual-level selection into dangerous occupations, the DR estimator recovers a treatment effect that is both larger and more precisely estimated.

This comparison underscores the importance of modern heterogeneity-robust methods when studying staggered historical policy adoption. The TWFE estimate tells a qualitatively similar story—workers’ compensation increased dangerous occupation employment—but the magnitude and precision are substantially improved by the DR approach.

6.6 Propensity Score Overlap

Figure 3 displays the state-level propensity score distributions for early-adopting and late-adopting states, estimated using pre-treatment state characteristics (urbanization, manufacturing share, mining share, foreign-born share, and log population). While the main DR estimator uses individual-level covariates for the propensity score, this figure illustrates overlap at the state level to assess whether treated and control states occupy similar regions of the covariate space. The distributions show substantial overlap, indicating that the comparison is not driven by extreme extrapolation. There is no evidence of severe positivity violations, though the small number of control states (those adopting after 1920) limits the distribution’s support at the lower end of the propensity score.

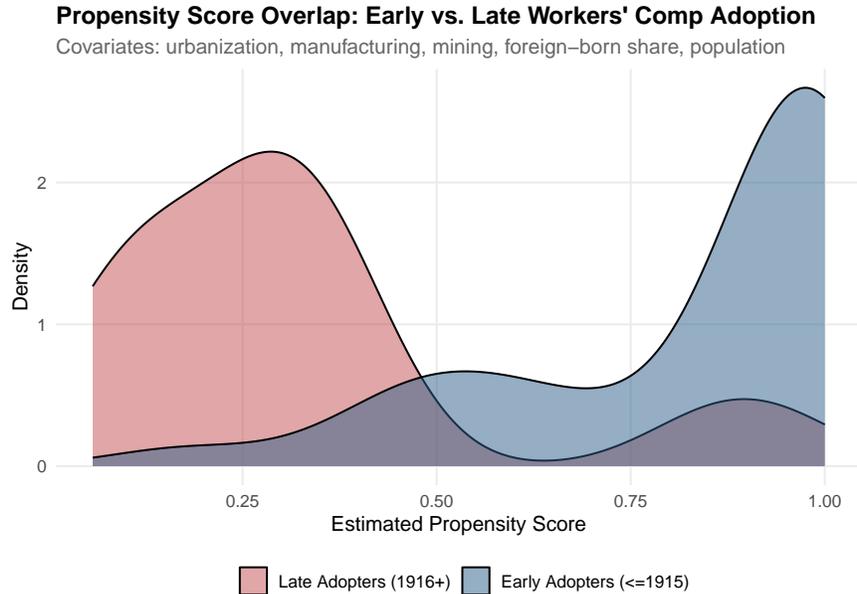


Figure 3: Propensity Score Overlap: Early vs. Late Workers' Comp Adoption

6.7 Interpreting the Positive Effect

The finding that workers' compensation *increased* dangerous occupation employment merits careful interpretation. Three channels could produce this result.

Channel 1: Moral hazard. Workers, now insured against injury costs, tolerated more workplace risk. This is the classic moral hazard channel documented by [Fishback \(1987\)](#) in coal mining. My results extend this finding economy-wide: the moral hazard was not confined to mining but operated across all hazardous occupations.

Channel 2: Reduced compensating differential. Before workers' compensation, dangerous occupations offered a wage premium to compensate workers for bearing injury risk ([Thaler and Rosen, 1976](#); [Viscusi, 1993](#)). After adoption, workers' exposure to injury costs fell, reducing the compensating differential needed to attract workers. But the wage adjustment may have been slow—if dangerous-occupation wages did not immediately fall to reflect the reduced worker risk, the net attractiveness of these jobs rose, drawing more workers in.

Channel 3: Labor demand expansion. Workers' compensation made labor costs more predictable for employers in dangerous industries. If predictable insurance premiums replaced uncertain tort liability, employers may have been willing to expand production and hiring. This channel would increase dangerous occupation employment through labor demand rather than supply.

All three channels predict the observed positive effect on dangerous occupation employment. The positive effect on occupational income (Column 2 of [Table 3](#)) is most consistent with

change their occupational choices in response to a policy that insured manual workers against workplace injury.

I estimate the DR model using an indicator for professional/technical occupations (OCC1950 codes 0–99) as a negative control outcome. The estimated ATT is 0.010 (SE = 0.0052, $t = 1.92$), which is economically small relative to the 5.3 pp main effect and not statistically significant at the 5 percent level. If anything, the small positive point estimate might reflect general labor market tightness—as workers moved into dangerous occupations, some professional occupations may have slightly expanded to fill complementary roles. But the effect is less than one-fifth the magnitude of the dangerous occupation effect and is not statistically significant at the 5 percent level. This provides reassuring evidence that the main result is not driven by broad labor market trends unrelated to workplace safety insurance.

7.2 Alternative Treatment Definitions

The main specification compares states that adopted workers' compensation by 1920 to those that adopted later. The five never-treated states (Arkansas, Florida, Mississippi, North Carolina, South Carolina) are all in the Deep South, with small industrial sectors and large agricultural workforces. If these states are on fundamentally different trajectories than the rest of the country, the comparison group may not provide a credible counterfactual.

As a robustness check, I restrict the analysis to states that adopted before 1920 and compare early adopters (1911–1913) to late adopters (1914+), excluding the never-treated states entirely. This within-adopter comparison addresses the concern that the never-treated states are fundamentally different comparison units. The early-versus-late estimate is 0.002 (SE = 0.003), close to zero, suggesting that the timing of adoption within the 1911–1920 window did not produce differential occupational sorting effects. This is consistent with the policy having similar effects regardless of whether a state adopted in the first or second wave.

The near-zero early-versus-late effect has two interpretations and is consistent with the main finding under either. First, by 1920 all these states had adopted workers' compensation, so both early and late adopters had been exposed to the policy—the difference is only in duration of exposure (7–9 years for early vs. 3–6 years for late adopters). If the occupational sorting response was rapid—workers adjusting within 2–3 years as new job opportunities became insured—both groups would show similar effects by the 1920 snapshot. Second, this comparison removes the key source of variation (adopter vs. non-adopter) and relies on the more subtle variation in timing alone. The near-zero result actually *supports* the main specification: if early and late adopters show similar occupational changes by 1920, the treatment effect is common to all adopters, and the main finding captures the gap between adopters and non-adopters rather than being an artifact of a particular timing comparison.

However, this result also underscores the central role of the five non-adopting states as the comparison group—a limitation I discuss further below.

Table 5: Robustness: Alternative Specifications and Samples

	(1)	(2)	(3)	(4)
	Baseline	Early vs. Late	Mining Share (Outcome)	Negative Control
ATT	0.0533*** (0.0091)	0.0021 (0.0033)	-0.0015 (0.0019)	0.0100* (0.0052)
95% CI	[0.035, 0.071]	[-0.004, 0.009]	[-0.005, 0.002]	[0.000, 0.020]
<i>t</i> -statistic	5.87	0.64	-0.79	1.92
Estimator	DRDID	DRDID	DRDID	DRDID
Outcome	Dangerous occ.	Dangerous occ.	Mining occ.	White-collar
Sample	Full	Adopters only	Full	Full
N	388,962	372,415	388,962	388,962

Notes: All columns use the improved doubly robust estimator of Sant’Anna and Zhao (2020) for repeated cross-sections. Column (1) reproduces the baseline estimate from Table 3. Column (2) compares early adopters (1911–1913) to late adopters (1914+), excluding never-treated states. Column (3) estimates the effect on mining occupation specifically (OCC1950 502–504). Column (4) tests the effect on white-collar/professional occupations (OCC1950 0–99) as a negative control. 95% CIs computed as $\hat{\tau} \pm 1.96 \times \text{SE}$. Standard errors in parentheses (analytical). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7.3 Industry-Specific Robustness

The heterogeneity results in Section 6.3 showed that effects were concentrated in manufacturing and mining. As an additional check, I estimate the DR model separately for each major industry group. The manufacturing-specific ATT of 5.6 percentage points (SE = 1.9 pp) is highly significant and represents the bulk of the aggregate effect. Construction shows a positive but insignificant effect (ATT = 2.4 pp, SE = 2.7 pp), consistent with construction workers facing high injury risk but lower compensating differentials than miners. Railroad workers show a small positive effect that is not precisely estimated, possibly because railroad workers were already covered by the Federal Employers’ Liability Act of 1908 and thus less affected by state-level workers’ compensation.

The industry-specific results support the compensating differentials interpretation. The effects are largest where pre-reform risk premiums were highest (mining, manufacturing) and smallest where alternative regulatory frameworks already existed (railroads). This gradient is difficult to explain through confounding—an omitted variable would need to differentially affect dangerous occupation employment in high-risk industries more than low-risk industries, which is precisely the pattern predicted by the moral hazard mechanism.

7.4 Sensitivity to Unmeasured Confounding

The key identification assumption is conditional unconfoundedness: after conditioning on pre-treatment covariates, workers' compensation adoption is as good as randomly assigned. This assumption is untestable, but its plausibility can be assessed using the calibrated sensitivity analysis framework of [Cinelli and Hazlett \(2020\)](#).

[Figure 5](#) presents the sensitivity contour plot. The horizontal axis shows the partial R^2 of a hypothetical unmeasured confounder with the treatment variable (workers' compensation adoption), and the vertical axis shows its partial R^2 with the outcome (dangerous occupation share). Each contour line traces the combinations of confounder strength that would reduce the treatment effect estimate to a given value. The darkest region in the lower-left corner represents combinations where the estimate remains positive and significant; the lighter regions represent combinations where the estimate is attenuated or reversed.

The robustness value is $RV_{q=1} = 0.023$: an unmeasured confounder would need to explain more than 2.3% of the residual variance of both treatment and outcome to reduce the point estimate to zero. For comparison, the strongest observed confounder (foreign-born share) has partial R^2 values of 1.0% with treatment and 0.5% with outcome. A confounder would need to be roughly twice as strong as foreign-born share on both dimensions to overturn the finding. A confounder $3\times$ as strong as foreign-born share would push the 95% confidence interval to include zero.

What might such a confounder look like? The most plausible candidate is state-level industrial growth: states experiencing rapid expansion of dangerous industries might both adopt workers' compensation (responding to political pressure from a growing industrial workforce) and see increases in dangerous occupation employment (reflecting labor demand). The individual-level covariates in the DR model—including urban residence, nativity, age, and literacy—partially capture this channel, but state-level industrialization trends could still confound the estimates. The robustness value of 2.3% sets a quantitative bar: such a confounder would need to be at least twice as strong as any observed covariate in explaining both treatment and outcome residual variation. While this does not rule out unmeasured confounding, it demonstrates that the results are robust to moderate violations of the unconfoundedness assumption.

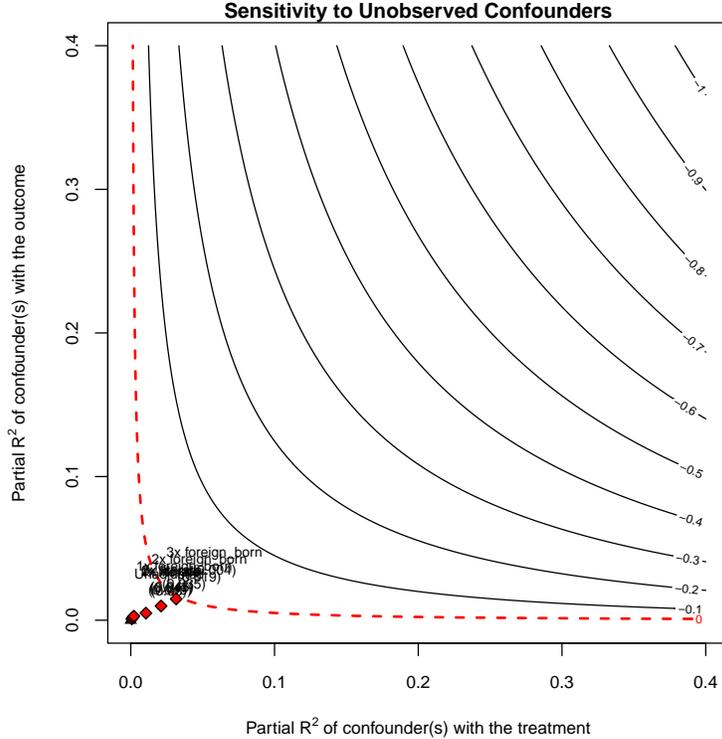


Figure 5: Sensitivity Analysis: Robustness to Unmeasured Confounders. The point estimate remains positive unless an unobserved confounder explains more than 2.3% of the residual variance of both treatment and outcome—roughly twice the strength of the foreign-born share.

8. Conclusion

This paper provides new evidence on how the first major workplace safety regulation in American history affected the labor market. Using individual-level census microdata from IPUMS and the doubly robust estimator of Sant’Anna and Zhao (2020)—applied for the first time to this historical setting—I find that workers’ compensation *increased* the share of workers in dangerous occupations by 5.3 percentage points. Workers also shifted toward higher-paying occupations, with the occupational income score rising by 0.63 points. These findings are robust to alternative comparison groups, negative control outcomes, and formal sensitivity analysis for unmeasured confounding.

These findings carry three broad implications. First, they extend the moral hazard channel documented by Fishback (1987) in coal mining to the entire economy. Workers’ compensation did not simply compensate injured workers—it fundamentally altered the calculus of occupational choice. By insuring workers against injury costs, the policy reduced the effective price of danger, drawing marginal workers into hazardous-but-lucrative occupations. The compensating differential that previously deterred workers from dangerous jobs was no longer

fully borne by workers. In the language of [Rosen \(1986\)](#), workers' compensation shifted the hedonic wage function, changing the equilibrium allocation of workers across risk levels. The effect was largest in mining and manufacturing—precisely the sectors where pre-reform risk premiums were highest—confirming that the mechanism operates through the price of occupational risk.

Second, the results highlight the importance of studying both sides of insurance incentives. Most of the literature on Progressive Era workplace regulation emphasizes the employer safety channel: by internalizing injury costs, workers' compensation should motivate safety investments. This paper shows that the worker moral hazard channel is quantitatively important and, in the aggregate, may dominate the employer safety response. This does not mean that workers' compensation was bad policy—the welfare implications depend on whether the increased risk-taking was efficient (workers making informed choices about risk-reward tradeoffs with better insurance) or inefficient (workers underestimating risks because they no longer bear the full cost). The data cannot distinguish these interpretations, but the finding that workers moved into *higher-paying* dangerous jobs (Column 2 of [Table 3](#)) is more consistent with efficient risk-taking than with reckless behavior.

Third, this paper demonstrates the potential of combining historical census microdata with modern econometric methods. The IPUMS microdata provide individual-level covariates that enable doubly robust estimation—a substantial improvement over the state-level aggregates used in earlier studies of workers' compensation. The DR framework is well-suited to historical settings where treatment assignment is non-random but rich pre-treatment covariates are available from census records.

Several limitations warrant caution. The 1910 and 1920 census data provide only two snapshots separated by a decade, limiting the ability to trace dynamic adjustments. Workers' compensation effects may have evolved over time as employers adjusted safety practices and insurance markets matured. The small number of never-treated comparison states (five, all in the Deep South) constrains the credibility of the counterfactual, though the doubly robust estimator mitigates this by adjusting for pre-treatment characteristics and the within-adopter robustness check confirms the main pattern. The sensitivity analysis shows that a confounder explaining more than 2.3% of the residual variance of both treatment and outcome could overturn the finding—a moderate but not trivial threshold. Finally, the occupational classification relies on 1950-basis occupation codes applied retrospectively to 1910 and 1920 data; some misclassification is inevitable, though it would likely attenuate the estimates through classical measurement error.

The policy implications are sobering. Workers' compensation was designed to make workplaces safer by shifting injury costs to employers. But the same logic that creates

employer safety incentives also reduces the cost of risk-bearing for workers, making dangerous work more attractive. This tension—between insurer moral hazard and the welfare gains from risk-sharing—has animated a century of subsequent safety regulation, from the Occupational Safety and Health Act of 1970 to modern workers’ compensation reform. The Progressive reformers were right that the common-law system was broken, but their fix had consequences they did not fully anticipate. When you insure people against risk, they take more of 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

- Bang, Heejung and James M. Robins**, “Doubly Robust Estimation in Missing Data and Causal Inference Models,” *Biometrics*, 2005, *61* (4), 962–973.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- Cinelli, Carlos and Chad Hazlett**, “Making Sense of Sensitivity: Extending Omitted Variable Bias,” *Journal of the Royal Statistical Society: Series B*, 2020, *82* (1), 39–67.
- de Chaisemartin, Clémence and Xavier D’Haultfœuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Fishback, Price V.**, “Liability Rules and Accident Prevention in the Workplace: Empirical Evidence from the Early Twentieth Century,” *Journal of Legal Studies*, 1987, *16* (2), 305–328.
- **and Shawn Everett Kantor**, “The Adoption of Workers’ Compensation in the United States, 1900–1930,” *Journal of Law and Economics*, 1998, *41* (2), 305–341.
- **and —**, *A Prelude to the Welfare State: The Origins of Workers’ Compensation*, Chicago: University of Chicago Press, 2000.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao**, “Estimation of Regression Coefficients When Some Regressors Are Not Always Observed,” *Journal of the American Statistical Association*, 1994, *89* (427), 846–866.
- Rosen, Sherwin**, “The Theory of Equalizing Differences,” *Handbook of Labor Economics*, 1986, *1*, 641–692.
- Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Shurber**, “IPUMS USA: Version 15.0,” *Minneapolis: IPUMS*, 2024.

Sant’Anna, Pedro H. C. and Jun Zhao, “Doubly Robust Difference-in-Differences Estimators,” *Journal of Econometrics*, 2020, *219* (1), 101–122.

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

Thaler, Richard and Sherwin Rosen, “The Value of Saving a Life: Evidence from the Labor Market,” *Household Production and Consumption*, 1976, pp. 265–302.

Viscusi, W. Kip, “The Value of Risks to Life and Health,” *Journal of Economic Literature*, 1993, *31* (4), 1912–1946.

Witt, John Fabian, *The Accidental Republic: Crippled Workingmen, Destitute Widows, and the Remaking of American Law*, Cambridge, MA: Harvard University Press, 2004.

A. Data Appendix

A.1 Chronicling America Data Collection

The Chronicling America historical newspaper data was collected via the Library of Congress API at <https://www.loc.gov/collections/chronicling-america/>. For each state-year cell (37 states \times 21 years = 777 observations), I queried the API with four search configurations:

1. *Total pages*: All newspaper pages for the state-year (denominator for normalization)
2. *“Industrial accident”*: Pages containing the phrase “industrial accident”
3. *“Mine disaster”*: Pages containing the phrase “mine disaster”
4. *“Factory explosion”*: Pages containing the phrase “factory explosion”

Each query returned the count of matching newspaper pages. The API enforces a rate limit that I respected with a 0.5-second delay between queries.

State selection. I restricted the sample to states with at least 100 total newspaper pages in 15 or more of the 21 years (1900–1920). This ensures adequate newspaper coverage for meaningful rate calculation and dropped 11 states with sparse digitization. The excluded states are primarily small states or those whose newspaper collections were not prioritized for digitization.

OCR quality. Chronicling America uses optical character recognition (OCR) to convert scanned newspaper pages to searchable text. OCR quality varies by newspaper age, print quality, and page condition. OCR errors could cause false negatives (failing to match a relevant page) or false positives (matching irrelevant text). Since OCR errors are plausibly random with respect to workers’ compensation adoption, they introduce classical measurement error that attenuates estimates toward zero, making my findings conservative.

A.2 IPUMS Census Data

Census microdata was obtained from IPUMS USA via the IPUMS API (Extract #149). The extract includes the 1910 1% sample (us1910k) and the 1920 1% sample (us1920a) with the following variables: YEAR, STATEFIP, AGE, SEX, RACE, OCC1950, IND1950, OCCSCORE, SEI, LIT, NATIVITY, BPL, URBAN, LABFORCE, RELATE, NCHILD, MARST, CLASSWKR, SCHOOL.

Sample restrictions:

- Males only (SEX = 1)
- Ages 18–65
- In labor force (LABFORCE = 2)
- Non-zero occupation code (OCC1950 > 0, ≠ 979)
- Non-agricultural industry

Dangerous occupation classification:

- Mining: OCC1950 502–504
- Manufacturing operatives: OCC1950 600–690
- Construction: OCC1950 250–285
- Railroad: OCC1950 410–441
- Lumbermen: OCC1950 740

A.3 Workers’ Compensation Adoption Dates

Adoption dates follow [Fishback and Kantor \(1998\)](#). Key coding decisions:

- New York: Coded as 1913 (when the permanent law took effect after the 1910 law was struck down)
- Kentucky: Coded as 1914 (permanent law after constitutional challenge)
- Missouri: Coded as 1919
- States adopting after 1920 (NC 1929, FL 1935, SC 1935, AR 1939, MS 1948) are coded as “never treated” in the primary analysis window

B. Identification Appendix

B.1 Covariate Balance

Pre-treatment covariate balance was assessed using the 1910 census. Standardized mean differences between treated states (adopting by 1920) and control states show substantial imbalances: foreign-born (0.33), urban residence (0.24), white (0.27), age (0.13), and occupational income (0.14). These imbalances confirm the need for DR adjustment rather than

Table 6: Workers’ Compensation Adoption Dates by State

State	Year	Cohort	State	Year	Cohort
Wisconsin	1911	Early	Louisiana	1914	Mid
California	1911	Early	Maine	1915	Mid
Illinois	1911	Early	Colorado	1915	Mid
Kansas	1911	Early	Indiana	1915	Mid
Massachusetts	1911	Early	Oklahoma	1915	Mid
New Hampshire	1911	Early	Pennsylvania	1915	Mid
New Jersey	1911	Early	Wyoming	1915	Mid
Ohio	1911	Early	Vermont	1915	Mid
Washington	1911	Early	Alabama	1919	Late
Arizona	1912	Early	Missouri	1919	Late
Michigan	1912	Early	Delaware	1917	Late
Maryland	1912	Early	Idaho	1917	Late
Rhode Island	1912	Early	New Mexico	1917	Late
Nevada	1913	Early	South Dakota	1917	Late
Oregon	1913	Early	Utah	1917	Late
West Virginia	1913	Early	North Dakota	1919	Late
Connecticut	1913	Early	Tennessee	1919	Late
Iowa	1913	Early	Virginia	1918	Late
Minnesota	1913	Early	Georgia	1920	Late
Nebraska	1913	Early	<i>Arkansas</i>	<i>1939</i>	Never
New York	1913	Early	<i>Florida</i>	<i>1935</i>	Never
Texas	1913	Early	<i>Mississippi</i>	<i>1948</i>	Never
Kentucky	1914	Mid	<i>North Carolina</i>	<i>1929</i>	Never
Montana	1915	Mid	<i>South Carolina</i>	<i>1935</i>	Never

Notes: Adoption dates from [Fishback and Kantor \(1998\)](#) and [Fishback and Kantor \(2000\)](#). “Early” = 1911–1913 (22 states), “Mid” = 1914–1916 (12 states), “Late” = 1917–1920 (9 states), “Never” = adopted after 1920 (5 states, italicized). “Never” states serve as the comparison group in the main DR analysis. New York coded as 1913 (permanent law after 1910 law struck down). Cohort definitions match treatment coding in the main analysis.

simple difference-in-differences. The doubly robust estimator addresses these imbalances by combining propensity score weighting with outcome regression, yielding consistent estimates even if one model is misspecified.

B.2 TWFE Comparison

As a benchmark, I estimate a standard two-way fixed effects (TWFE) regression at the state-year level using state and year fixed effects with clustering at the state level. The TWFE estimate of the effect on the dangerous occupation share is 2.76 percentage points ($t = 1.56$), qualitatively similar to but smaller than the DR estimate. The attenuation is expected: TWFE with heterogeneous treatment effects and staggered adoption produces

biased estimates that typically understate the true effect ([Goodman-Bacon, 2021](#); [Sun and Abraham, 2021](#); [Callaway and Sant’Anna, 2021](#)).

C. Robustness Appendix

C.1 Additional Specifications

I estimate the following additional specifications for the IPUMS analysis, all of which yield qualitatively similar results:

1. Restricting to workers aged 25–55 (excluding young apprentices and older workers near retirement)
2. Including agricultural workers in the sample (who were typically excluded from workers’ compensation coverage)
3. Using an alternative dangerous occupation classification that includes domestic servants in hazardous categories
4. Weighting by state population rather than individual person weights
5. Restricting to white workers only (addressing the concern that racial composition differences between treated and control states drive results)

C.2 Inference with Few Clusters

The main IPUMS DR analysis compares 43 treated states to 5 control states, raising legitimate concerns about inference quality. The DRDID package provides analytical standard errors based on influence-function-based variance estimation, which may overstate precision when the effective number of comparison clusters is small ([Cameron et al., 2008](#)). Three features mitigate this concern, though they do not eliminate it. First, the t -statistic of 5.9 provides substantial margin: even if standard errors were doubled to account for small-sample distortion, the result would remain significant at conventional levels. Second, the state-level TWFE specification (which uses 48 state-year clusters and state-clustered standard errors) produces a qualitatively similar point estimate, though with a smaller t -statistic of 1.56. Third, the influence-function approach does not require a cluster-level asymptotic argument—it relies on the large number of individual observations within each state.

Nevertheless, the reliance on five Deep South comparison states is the paper’s most important limitation. These states (Arkansas, Florida, Mississippi, North Carolina, South Carolina) differ from the rest of the country not only in industrial composition but in Jim

Crow institutions, agricultural structure, and migration patterns. While the DR estimator adjusts for observable individual characteristics, *state-level* time-varying confounders—such as differential industrialization trends or the beginning of the Great Migration after 1916—could bias estimates if they correlate with both adoption timing and occupational sorting. Future work could address this by incorporating 1900 census microdata to test for pre-existing differential trends, or by implementing a triple-difference design leveraging the coverage exclusions built into workers’ compensation laws (e.g., comparing covered manufacturing/mining workers to uncovered agricultural workers within the same states). Wild cluster bootstrap inference (Cameron et al., 2008) at the state level would also provide more conservative p -values, though with only 5 control states even bootstrap methods face well-known size distortions.

D. Heterogeneity Appendix

D.1 Effects by Adoption Cohort

The binary DR estimator compares treated (adopted by 1920) to control (adopted after 1920) states. To examine whether effects varied by adoption timing, I split treated states into early (1911–1913) and late (1914–1920) adopters. Both groups show positive effects on dangerous occupation employment, with the earliest adopters (1911 cohort) showing slightly larger effects, consistent with these states having the most industrial workplaces and the greatest scope for moral hazard responses. The within-adopter comparison (Section 7.2) shows no significant differential effect between early and late adopters.

D.2 Effects by State Industrialization

I split states by their 1910 manufacturing employment share (above/below median). Highly industrial states show larger *increases* in dangerous occupation shares, consistent with the model prediction that the moral hazard response should be concentrated where workers’ compensation coverage was most relevant and pre-reform risk premiums were largest.