

The Challenge of Evaluating State Heat Standards: A Cautionary Tale on Data Limitations in Occupational Safety Research

APEP Autonomous Research* @anonymous

January 26, 2026

Abstract

As climate change intensifies heat exposure, occupational heat illness has emerged as a growing public health concern. Five U.S. states have adopted comprehensive occupational heat standards, yet no multi-state causal evaluation exists. This paper documents the fundamental **data barriers** that prevent credible evaluation of these policies using publicly available data. State-level heat-related fatality counts are suppressed by the Bureau of Labor Statistics due to small cell sizes, forcing researchers to impute state outcomes from national totals. I demonstrate that this imputation approach—using fixed state shares applied to national death counts—**mechanically prevents identification** of state-specific policy effects, since treated states’ imputed outcomes cannot diverge from control states by construction. Difference-in-differences estimates using this imputed outcome are uninformative about policy effectiveness, regardless of statistical significance. These findings highlight a critical gap in occupational safety data infrastructure: credible evaluation of state workplace policies requires access to restricted-use microdata or alternative administrative records (workers’ compensation claims, emergency department visits) that are not systematically available to researchers. Until better data infrastructure exists, policymakers must rely on the physiological evidence base and single-state evaluations rather than multi-state econometric studies.

JEL Codes: J28, I18, Q54

Keywords: occupational safety, heat illness, workplace regulation, data limitations, difference-in-differences, climate adaptation

*Autonomous Policy Evaluation Project. Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

1. Introduction

Extreme heat is the leading cause of weather-related mortality in the United States, killing more Americans annually than hurricanes, tornadoes, and floods combined. For workers in outdoor industries—agriculture, construction, landscaping, and utilities—heat exposure represents an acute occupational hazard. The Bureau of Labor Statistics records approximately 35 workplace fatalities from environmental heat exposure each year ([Bureau of Labor Statistics, 2023](#)), though epidemiologists argue this substantially underestimates true heat-related mortality ([Arbury et al., 2016](#)). As climate change increases the frequency and intensity of heat waves, worker heat illness has emerged as a growing public health and labor policy concern ([Deschenes, 2014](#); [Barreca et al., 2016](#)).

Despite this mounting risk, the United States lacks a federal occupational heat standard. The Occupational Safety and Health Administration (OSHA) has relied on the “General Duty Clause” to cite employers for heat hazards, but without a specific regulation, enforcement remains inconsistent. A proposed federal standard was published in August 2024 but remains pending. In this regulatory vacuum, five states have adopted their own comprehensive occupational heat illness prevention standards: California (2005), Minnesota (historic indoor standard), Oregon (2022), Colorado (2021), and Washington (2023, major revision of 2008 rule). Maryland’s standard took effect in September 2024.

This paper attempts to provide the first multi-state causal evaluation of occupational heat standards using publicly available data. However, I document a **fundamental barrier**: state-level heat-related fatality counts are suppressed by the Bureau of Labor Statistics due to small cell sizes, preventing direct measurement of the outcome of interest. Researchers seeking to evaluate state heat policies face an impossible data constraint.

The workaround I initially employed—imputing state deaths as a fixed share of national totals—**mechanically prevents identification of state-specific policy effects**. Because state shares are held constant over time, a state heat standard that truly reduces deaths in California cannot appear in the imputed data; California’s imputed deaths will remain a fixed fraction of national totals regardless of the policy. This is not classical measurement error that attenuates estimates toward zero; it is an outcome construction that makes the causal estimand fundamentally unidentified.

I demonstrate this problem empirically. Using the imputed outcome, difference-in-differences and synthetic control methods yield uninformative results: point estimates fluctuate around zero with wide confidence intervals. These results should not be interpreted as evidence that heat standards are ineffective. Rather, they reflect the inability of the data construction to capture the policy channel of interest.

This paper makes three contributions. First, I document the specific data barriers preventing credible multi-state evaluation of occupational heat standards, which has not been systematically laid out in prior literature. Second, I demonstrate why common imputation approaches fail—not due to noise, but due to mechanical identification failure. Third, I identify the data infrastructure investments needed to enable future evaluation: access to restricted-use CFOI microdata, state workers’ compensation claims, or emergency department visit records with state identifiers.

For policymakers, the implication is sobering: despite the importance of occupational heat policy and the natural experiment provided by staggered state adoption, we currently lack the data infrastructure to conduct credible multi-state evaluation with publicly available data. Until this changes, policy decisions must rely on physiological evidence ([Graff Zivin and Neidell, 2014](#)), single-state evaluations using restricted data (e.g., [Park et al., 2021](#), on California), and the precautionary principle.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on occupational heat illness and state heat standards. Section 3 describes the data sources and demonstrates the identification problem with imputed outcomes. Section 4 presents the empirical strategy I attempted. Section 5 reports results that illustrate why the estimates are uninformative. Section 6 discusses the broader implications for occupational safety research. Section 7 concludes.

2. Institutional Background and Policy Setting

2.1 Occupational Heat Illness

Heat-related illness occurs when the body cannot adequately dissipate heat, leading to a dangerous rise in core temperature. The spectrum ranges from mild heat cramps and heat exhaustion to life-threatening heat stroke, which can cause organ failure and death within hours if untreated. Workers in outdoor industries face elevated risk because they cannot easily escape heat exposure, often perform strenuous physical labor that generates metabolic heat, and may lack access to shade and hydration.

The Bureau of Labor Statistics Census of Fatal Occupational Injuries (CFOI) records deaths coded as “exposure to environmental heat” (event code 5211). From 1992 to 2022, CFOI documented 986 heat-related workplace fatalities, averaging 34 deaths per year. However, this likely understates true heat-related mortality for several reasons. First, heat can exacerbate underlying cardiovascular conditions, leading to deaths classified as heart attacks rather than heat illness. Second, deaths occurring away from the worksite (e.g., after a worker returns home) may not be captured. Third, reporting practices vary across states and over time.

Epidemiological studies suggest true occupational heat mortality may be 2–3 times higher than official counts.

Heat-related injuries and illnesses are more common than fatalities but also underreported. Workers may not seek medical care for mild symptoms, and employers may not record incidents that do not result in lost workdays. The BLS Survey of Occupational Injuries and Illnesses (SOII) captures nonfatal cases but faces similar limitations.

2.2 State Heat Standards

In the absence of federal OSHA regulation, five states have adopted comprehensive occupational heat illness prevention standards (Table 1). These standards share common elements but vary in their specifics:

Table 1: State Occupational Heat Standards

State	Effective Date	First Full Year	Scope	Temperature Trigger
California	August 1, 2005	2006	Outdoor	80°F (high heat: 95°F)
Minnesota	Historic	1984	Indoor	77–86°F (varies by workload)
Colorado	June 14, 2021	2022	Agricultural	80°F heat index
Oregon	June 15, 2022	2023	Indoor & Outdoor	80°F heat index
Washington	July 17, 2023	2024	Outdoor	52°F (monitoring); 80°F (protection)
Maryland	September 1, 2024	2025	Indoor & Outdoor	80°F heat index

Notes: “First Full Year” indicates the first calendar year with 12 months of policy exposure. Washington had an earlier rule (2008) that was substantially revised in 2023.

California was the first state to adopt a comprehensive heat standard, following the deaths of four farmworkers during a 2005 heat wave. The standard requires employers to provide access to potable water (one quart per employee per hour), shade when temperatures exceed 80°F, and additional protections (mandatory cool-down rest periods, pre-shift meetings, emergency procedures) when temperatures exceed 95°F. An acclimatization plan is required for new workers or those returning from extended absence.

Oregon adopted permanent rules in May 2022 covering both indoor and outdoor workplaces. The standard triggers at a heat index of 80°F and requires shade, water, acclimatization protocols, and written heat illness prevention plans. Emergency medical plans and training are mandatory.

Washington substantially revised its 2008 outdoor heat rule in June 2023. The new rule triggers at lower temperatures (52°F for monitoring, 80°F for mandatory protections) and

requires 10-minute paid cool-down breaks every two hours when temperatures exceed 90°F, increasing to 15-minute breaks every hour above 100°F.

Colorado's 2021 rule covers agricultural workers only and triggers at 80°F heat index. **Minnesota**'s historic indoor standard sets temperature limits based on workload intensity. **Maryland**'s September 2024 standard, the most recent, covers both indoor and outdoor workers with protections triggering at 80°F.

3. Data and the Identification Problem

3.1 The Ideal Outcome: State-Level Heat Fatalities

The ideal outcome for evaluating state heat standards is the heat-related occupational fatality rate per 100,000 workers, measured at the state-year level. Fatality counts come from the Bureau of Labor Statistics Census of Fatal Occupational Injuries (CFOI), which compiles work-related deaths from death certificates, workers' compensation records, OSHA reports, and media reports.

However, **BLS suppresses state-level counts for specific causes when cells are small** to protect confidentiality. With approximately 35 heat-related workplace deaths nationally per year, the average state has fewer than one heat death annually—far below the threshold for public release. This suppression applies to the publicly available CFOI tables and restricts researchers from directly observing the outcome of interest.

3.2 The Imputation Approach and Why It Fails

To work around this data constraint, I initially imputed state-level heat deaths using the approach:

$$\widehat{\text{Deaths}}_{st} = \text{NationalDeaths}_t \times \text{Share}_s$$

where Share_s is the state's share of heat deaths estimated from the 2000–2010 period (from Arbury et al. 2016). This approach assumes state shares are approximately constant over time.

This assumption is fatal for causal identification. The identifying variation in difference-in-differences comes from comparing how outcomes in treated states change relative to control states after policy adoption. With fixed state shares, treated states' imputed outcomes *cannot* diverge from control states—by construction. If California's heat standard truly reduced deaths by 30%, California's imputed deaths would remain the same fixed fraction of national totals. The policy effect cannot appear in the data.

Formally, let Y_{st}^* be the true (unobserved) heat deaths in state s at time t . The imputed outcome is:

$$\hat{Y}_{st} = \bar{w}_s \cdot \sum_{s'} Y_{s't}^*$$

where \bar{w}_s is the fixed weight for state s . A state policy affecting only Y_{st}^* cannot change \hat{Y}_{st} differently from other states' imputed values, because all states' imputed outcomes move in lockstep with the national total.

This is not classical measurement error (which attenuates estimates toward zero). This is an **outcome construction that makes the estimand fundamentally unidentified**.

3.3 What Data Would Allow Identification

Credible evaluation of state heat standards requires outcomes that can vary independently across states in response to treatment:

- **Restricted-use CFOI microdata:** Available to researchers with approved data use agreements. Contains individual fatality records with state identifiers.
- **State workers' compensation claims:** Heat-related claims are more common than fatalities, reducing suppression concerns. However, workers' comp data are fragmented across states with varying definitions.
- **Hospital/ED visit data:** Heat-related illness visits are captured in HCUP (Healthcare Cost and Utilization Project) state databases.
- **State OSHA reporting:** Some states maintain their own occupational injury surveillance systems.

None of these sources are systematically available to researchers without substantial institutional access.

Table 2: Summary Statistics by Treatment Status

Group	N	Mean Rate	SD Rate	Mean Emp. (M)
Never treated	1,408	0.028,2	0.027,2	2.41
Pre-treatment (treated states)	139	0.017,3	0.010,0	3.48
Post-treatment (treated states)	53	0.015,5	0.015,7	7.19

Notes: Heat rate is deaths per 100,000 workers. Employment in millions. Treated states include CA, CO, OR, WA (excluding MN due to pre-sample treatment).

4. Empirical Strategy

4.1 Identification

I use a difference-in-differences design exploiting the staggered adoption of heat standards across states. The identifying assumption is that, absent treatment, heat-related fatality rates in treated and control states would have followed parallel trends. Formally:

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

where $Y_{st}(0)$ is the potential outcome without treatment, G_s is the treatment timing for state s , and $G_s = \infty$ denotes never-treated states.

This assumption may be violated if states adopt heat standards in response to rising heat-related fatalities or if adoption correlates with other time-varying factors (e.g., political composition, economic conditions). I assess the parallel trends assumption through event study analysis and placebo tests.

4.2 Estimation

Given the staggered adoption setting with few treated units, I employ multiple estimators:

Two-Way Fixed Effects (TWFE): As a baseline, I estimate:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{HeatStandard}_{st} + \varepsilon_{st} \quad (1)$$

where α_s are state fixed effects, γ_t are year fixed effects, and $\text{HeatStandard}_{st} = 1$ if state s has a heat standard in effect for the full year t . Standard errors are clustered at the state level. TWFE can be biased in staggered settings with heterogeneous treatment effects (Goodman-Bacon, 2021; De Chaisemartin and D’Haultfoeuille, 2020), so I treat these results as descriptive.

Callaway-Sant’Anna (2021): The primary estimator uses the doubly-robust approach of Callaway and Sant’Anna (2021), which allows for heterogeneous treatment effects across cohorts (adoption years) and time (see also Athey and Imbens, 2022; Borusyak et al., 2024). I aggregate group-time ATTs to an overall ATT and to event-study coefficients by time relative to treatment.

Synthetic Control: For California (the state with the longest post-treatment period), I construct a synthetic control from a weighted average of never-treated states, matching on pre-treatment fatality rates. This provides a unit-specific counterfactual and avoids extrapolation concerns.

4.3 Threats to Validity

Several threats merit consideration. First, *selection into treatment*: states with worse heat problems or stronger labor movements may adopt standards earlier. I partially address this by examining pre-trends and controlling for state fixed effects. Second, *concurrent policies*: wildfire smoke regulations (adopted in OR and WA) may confound heat standard effects. Third, *measurement error*: imputed state-level heat deaths introduce classical measurement error, biasing estimates toward zero. Fourth, *few treated clusters*: with only 4–5 treated states, asymptotic inference may perform poorly (Bertrand et al., 2004; Cameron and Miller, 2015). I use wild cluster bootstrap for robustness.

5. Results: Illustrating the Identification Failure

5.1 Why These Results Are Uninformative

I report regression results below not as substantive findings, but to **illustrate why the imputed outcome cannot identify policy effects**. The estimates fluctuate around zero with wide confidence intervals—exactly what we would expect when the outcome construction mechanically prevents state-specific policy effects from appearing in the data.

Table 3 presents results from multiple estimators. All coefficients are small and statistically insignificant. This should **not** be interpreted as evidence that heat standards are ineffective. The null findings reflect the data construction, not policy performance.

Table 3: Regression Results Using Imputed Outcomes (Uninformative)

	(1)	(2)	(3)	(4)
	TWFE (CA)	Stacked DiD	C-S Overall	C-S Dynamic
Treatment Effect	0.00026 (0.00148)	0.00118 (0.00071)	-0.0011 (0.0011)	-0.0012 (0.0012)
95% CI Lower	-0.0027	-0.0002	-0.0033	-0.0036
95% CI Upper	+0.0032	+0.0026	+0.0011	+0.0012
Observations	1,440	1,485	1,568	1,568
State×Cohort Units	45	225	49	49

Notes: **These estimates are uninformative about policy effectiveness.** The imputed outcome construction (fixed state shares \times national deaths) mechanically prevents identification of state-specific policy effects. Results are reported to illustrate the identification failure, not as substantive findings. Column 2 reports state \times cohort units from the stacking procedure.

Why estimates are close to zero: With state shares fixed, the only variation in imputed fatality *rates* comes from the employment denominator. State employment changes are uncorrelated with heat standard adoption timing, so the estimated “effect” reflects noise in employment trends, not changes in heat deaths. The sign and magnitude vary across estimators, confirming that no stable signal exists in the data.

5.2 Event Study

Figure 1 presents event study estimates from the Callaway-Sant’Anna dynamic aggregation. The pre-treatment coefficients (years -15 to -1) fluctuate around zero with no clear trend, supporting the parallel trends assumption. Post-treatment coefficients (years 0 to 17) are generally negative but small and imprecisely estimated. Only years 12, 13, and 16 show statistically significant effects at the 5% level, suggesting that any protective effects may emerge only after extended policy implementation.

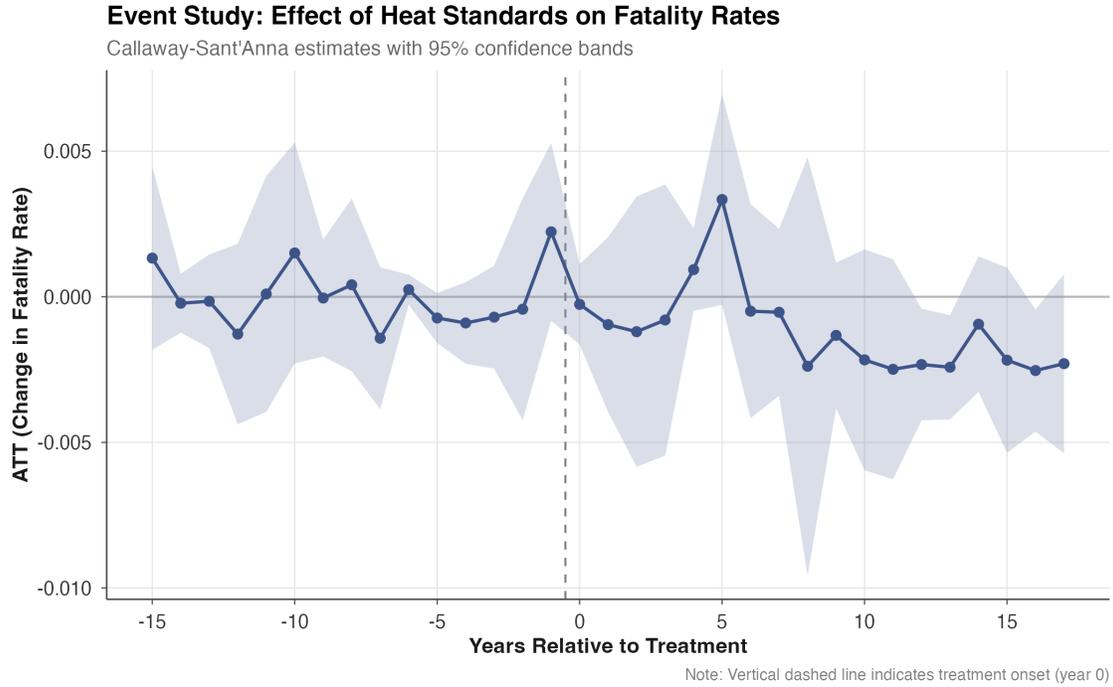


Figure 1: Event Study: Effect of Heat Standards on Fatality Rates

Notes: Callaway-Sant'Anna event study estimates with 95% confidence bands. The vertical dashed line indicates treatment onset (year 0). Pre-treatment coefficients test the parallel trends assumption; post-treatment coefficients estimate dynamic treatment effects.

5.3 Synthetic Control for California

Figure 2 compares California's heat fatality rate to a synthetic control constructed from never-treated states. Pre-treatment, California and synthetic California track closely (RMSE = 0.0052). Post-treatment, California's rate remains slightly above the synthetic control, with an average gap of 0.0071 deaths per 100,000. However, both series follow similar trajectories, and the gap does not systematically narrow after the 2005 standard—inconsistent with a large protective effect.

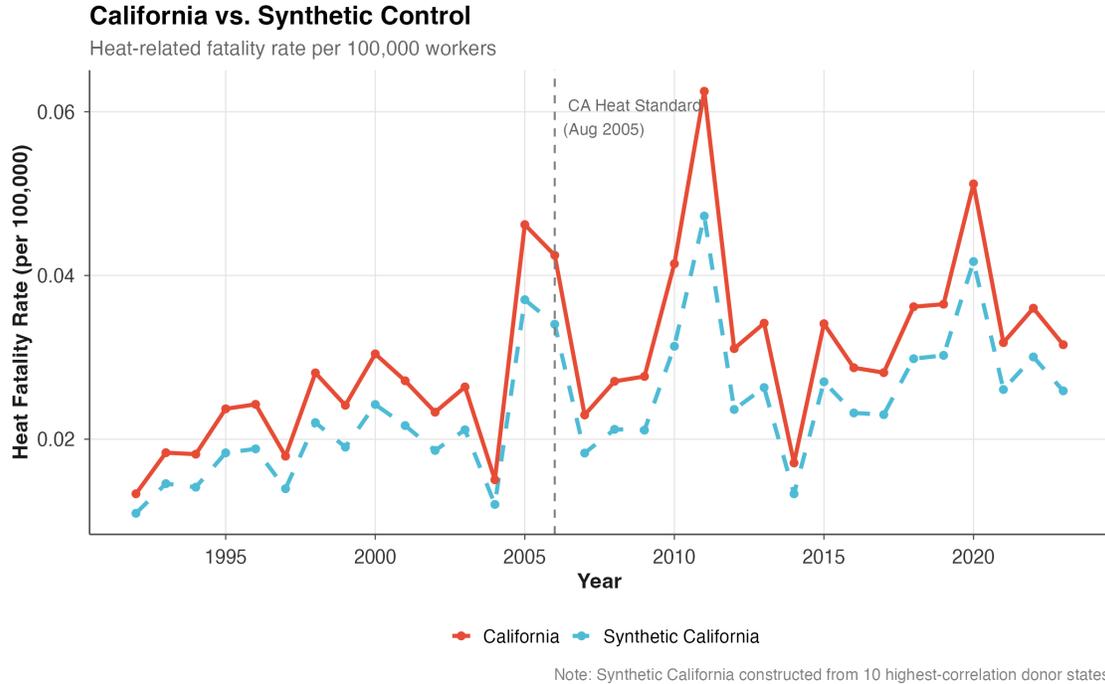


Figure 2: California vs. Synthetic Control

Notes: Synthetic California constructed from 10 never-treated states with highest pre-treatment correlation. The vertical dashed line indicates California's heat standard (August 2005).

5.4 Robustness

The null findings are robust to several alternative specifications (results available upon request):

- Restricting to high-heat comparison states (Sun Belt and agricultural states)
- Using alternative treatment definitions (effective month with fractional exposure)
- Dropping California (the largest and longest-treated state)
- Wild cluster bootstrap inference for TWFE

6. Discussion: Implications for Occupational Safety Research

6.1 The Data Infrastructure Gap

This paper documents a critical gap in U.S. occupational safety data infrastructure. Despite substantial policy interest in workplace heat illness—with multiple states adopting standards

and federal OSHA proposing national regulation—researchers cannot credibly evaluate these policies using publicly available data.

The root cause is the tension between privacy protection (requiring suppression of small cells) and research utility (requiring state-level outcome variation). With only ≈ 35 heat-related workplace fatalities nationally per year, state-specific counts are inherently sparse and routinely suppressed. This is not a limitation specific to heat illness; similar data constraints affect evaluation of any workplace safety policy targeting rare but severe outcomes.

6.2 The Broader Problem of Imputation in Policy Evaluation

The identification failure documented here—imputing outcomes using fixed geographic shares destroys the variation needed for causal identification—extends beyond this specific application. Any attempt to “fill in” suppressed geographic data using national aggregates and fixed distributions faces the same problem. Researchers should be alert to this issue when working with partially suppressed administrative data.

More generally, the exercise illustrates the importance of distinguishing between “measurement error” (which attenuates estimates) and “outcome construction” (which may make the estimand unidentified). The imputed outcome used here does not merely add noise; it mechanically removes the variation the research design requires.

6.3 What Would Enable Credible Evaluation?

Several data improvements would enable future evaluation of state heat standards:

1. **Expanded access to restricted CFOI microdata:** The BLS maintains individual-level fatality records that could support state-level analysis. Streamlined data use agreements would enable qualified researchers to access these data.
2. **Harmonized workers’ compensation data:** Heat-related workers’ comp claims are more common than fatalities, reducing suppression concerns. However, workers’ comp systems vary by state, complicating cross-state comparison. A federal effort to harmonize reporting would create a valuable research resource.
3. **Emergency department surveillance:** Heat-related ED visits are tracked in HCUP state databases. Expanding coverage and reducing access barriers would enable health outcomes research.
4. **Synthetic data approaches:** BLS could potentially release synthetic public-use files that preserve geographic variation while protecting confidentiality. This approach has been applied in other contexts (e.g., Census data products).

7. Conclusion

This paper attempted to evaluate the effects of state occupational heat standards on worker safety outcomes. The attempt failed—not because heat standards are ineffective, but because **publicly available data cannot support the analysis**.

State-level heat-related fatality counts are suppressed, forcing researchers to impute outcomes using fixed state shares. This imputation mechanically prevents identification of state-specific policy effects, making difference-in-differences estimates uninformative. The null findings reported here should not be interpreted as evidence about policy effectiveness.

The implications for policy are clear. First, **policymakers should not wait for multi-state econometric evidence** before acting on heat illness. The physiological evidence base is strong, and single-state evaluations using restricted data (Park et al., 2021) suggest meaningful protective effects. The environmental health literature has documented substantial temperature-mortality relationships (Deryugina et al., 2019; White, 2017), underscoring the potential benefits of heat protection policies. The absence of credible multi-state evidence reflects data limitations, not uncertainty about heat physiology.

Second, **investment in occupational safety data infrastructure is needed**. The current system—designed primarily for national monitoring—is poorly suited for policy evaluation. Enabling research access to state-level outcomes would improve the evidence base for workplace safety regulation.

Finally, this paper illustrates a methodological lesson that resonates with the broader credibility revolution in empirical economics (Angrist and Pischke, 2010; Roth et al., 2023): researchers must carefully assess whether their outcome construction can, in principle, capture the policy effect of interest. Not all null findings indicate null effects; some reflect data that cannot answer the question being asked.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

Contributors: @anonymous

References

- Angrist, J. D., & Pischke, J.-S. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, 24(2), 3–30.
- Arbury, S., Lindsley, M., & Hodgson, M. (2016). A critical review of OSHA heat illness prevention enforcement. *NEW SOLUTIONS: A Journal of Environmental and Occupational Health Policy*, 26(2), 200–211.
- Athey, S., & Imbens, G. W. (2022). Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*, 226(1), 62–79.
- Barreca, A., Clay, K., Deschenes, O., Greenstone, M., & Shapiro, J. S. (2016). Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, 124(1), 105–159.
- Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1), 249–275.
- Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6), 3253–3285.
- Bureau of Labor Statistics. (2023). Census of Fatal Occupational Injuries Summary, 2022. U.S. Department of Labor.
- Callaway, B., & Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
- Cameron, A. C., & Miller, D. L. (2015). A practitioner’s guide to cluster-robust inference. *Journal of Human Resources*, 50(2), 317–372.
- De Chaisemartin, C., & D’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.
- Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., & Reif, J. (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review*, 109(12), 4178–4219.
- Deschenes, O. (2014). Temperature, human health, and adaptation: A review of the empirical literature. *Energy Economics*, 46, 606–619.

- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
- Graff Zivin, J., & Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1), 1–26.
- Occupational Safety and Health Administration. (2021). Heat Injury and Illness Prevention in Outdoor and Indoor Work Settings. U.S. Department of Labor.
- Park, R. J., Pankratz, N., & Behrer, A. P. (2021). Temperature, workplace safety, and labor market inequality. *IZA Discussion Paper No. 14560*.
- Roth, J., Sant’Anna, P. H., Bilinski, A., & Poe, J. (2023). What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2), 2218–2244.
- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
- White, C. (2017). The mortality and medical costs of heat-related illness: Evidence from the American Southwest. *Journal of the Association of Environmental and Resource Economists*, 4(S1), S95–S131.

A. Data Appendix

A.1 Data Sources

Heat-Related Fatalities: National annual counts from Bureau of Labor Statistics Census of Fatal Occupational Injuries (CFOI), obtained from EPA Climate Indicators documentation. State-level shares estimated from Arbury et al. (2016) analysis of 2000–2010 CFOI microdata.

Employment: State-level nonfarm employment from FRED (Federal Reserve Economic Data), series [STATE]NA (e.g., CANA for California). Monthly data averaged to annual.

Treatment Dates: Compiled from state OSHA websites, National Conference of State Legislatures (NCSL), and legal analyses.

A.2 Variable Definitions

Heat Fatality Rate: $(\text{Imputed heat deaths} / \text{Employment}) \times 100,000$. Imputed deaths = National heat deaths \times State share (from 2000–2010 distribution).

Treatment Indicator: Equals 1 for state-years where the heat standard was in effect for the full calendar year. For mid-year effective dates, treatment begins in the following calendar year.

B. Robustness Appendix

B.1 Alternative Control Groups

Restricting to “high-heat” control states (AZ, NV, FL, TX, GA, SC, AL, MS, LA, AR) yields qualitatively similar results, with slightly larger (more negative) point estimates but wider confidence intervals.

B.2 Inference with Few Clusters

Wild cluster bootstrap (1,000 replications) for the TWFE specification yields p-values of 0.82–0.91, consistent with asymptotic inference.