

The Resilient Grid: Why the Largest Power Failure in U.S. History Left No Trace in the Labor Market

APEP Autonomous Research* @ailscl

March 24, 2026

Abstract

Winter Storm Uri knocked out power for 4.5 million Texas customers in February 2021, killing 246 people and causing an estimated \$80–130 billion in damages. I exploit the boundary between Texas’s isolated ERCOT grid and the nationally connected SPP/MISO/WECC grids—determined decades earlier by regulatory arbitrage, not economic conditions—to estimate the causal effect on county-level employment. The headline result is a precise null: ERCOT counties show zero differential employment change in Q1 2021 ($\hat{\beta} = 0.001$, $SE = 0.012$). Positive post-period coefficients in baseline specifications are absorbed by county-specific trends, reflecting pre-existing growth differentials rather than storm effects. The quarterly labor market was remarkably resilient to an infrastructure catastrophe that dominated insurance markets, housing repairs, and mortality statistics for years.

JEL Codes: Q54, R11, H84, J21

Keywords: infrastructure failure, electrical grid, ERCOT, Winter Storm Uri, labor market resilience, natural disaster

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 49m).

1. Introduction

On February 14, 2021, the temperature in Amarillo, Texas dropped to -11°F . The lights stayed on. Two hundred miles southeast, Dallas reached 2°F —thirteen degrees warmer—and the city went dark. The difference was not geography or weather preparedness. It was a regulatory decision made decades earlier: Amarillo’s power comes from the Southwest Power Pool, connected to the national grid; Dallas depends on the Electric Reliability Council of Texas (ERCOT), deliberately isolated from interstate transmission to avoid federal jurisdiction. When ERCOT’s generators failed under extreme cold, there was nowhere to import power from. The result was the largest controlled load-shedding event in American history: 4.5 million customers without electricity for up to five days, 246 deaths, and damage estimates between \$80 and \$130 billion (Busby et al., 2021; Doss-Gollin et al., 2021).

This paper asks whether that catastrophe left a mark on the labor market. The answer, surprisingly, is no.

I exploit the ERCOT/SPP grid boundary as a natural experiment. Texas counties served by ERCOT (213 counties, including Dallas, Houston, Austin, and San Antonio) experienced massive blackouts during Uri. Counties served by SPP, MISO, or WECC (41 counties, including the Panhandle, parts of East Texas, and El Paso) maintained power through interstate connections despite experiencing comparable or colder temperatures. Grid membership was determined by utilities’ historical decisions to avoid Federal Energy Regulatory Commission (FERC) oversight—a regulatory arbitrage choice orthogonal to county-level economic fundamentals (Rhodes, 2021).

Using quarterly county-level employment data from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW) for 254 Texas counties over 2018–2023, I estimate difference-in-differences specifications comparing ERCOT to non-ERCOT counties before and after February 2021. The main finding is a precise zero on persistent quarterly disruption: the coefficient on the ERCOT \times Uri Quarter interaction is 0.001 (SE = 0.012), meaning I can rule out employment effects larger than 2.5 percent in either direction. This null persists across specifications—including county-specific linear trends, rural-only samples, and size-matched comparisons—and across outcomes including wages and establishment counts. I emphasize that this result speaks to quarterly persistence, not to the absence of any short-run dislocation: the five-day blackout occurred in mid-February, and quarterly QCEW averages incorporate March and April, during which businesses reopened. The finding is that the labor market recovered so rapidly that the disruption left no residual in the lowest-frequency administrative data available.

This finding contributes to three literatures. First, it speaks to the economics of infras-

structure resilience (Rose, 2004; Hallegatte et al., 2019). While the engineering and insurance literatures document Uri’s massive costs through property damage, burst pipes, and mortality (Busby et al., 2021; Lee and Maron, 2022), the labor market tells a different story: businesses reopened, workers returned, and quarterly employment showed no discontinuity. The contrast between catastrophic physical damage and zero labor market disruption suggests that the economic costs of infrastructure failures concentrate in capital losses rather than persistent labor reallocation.

Second, the paper informs the ongoing debate over ERCOT’s grid isolation. Multiple bills have proposed connecting ERCOT to the national interconnection since Uri, with proponents arguing that isolation imposes unacceptable economic costs (Rhodes, 2021; King et al., 2021). My results suggest that while grid isolation clearly contributed to the blackout and its associated mortality and property damage, the quarterly employment costs are negligible—implying that the case for interconnection rests on mortality and property protection rather than labor market insurance.

Third, this paper adds to the growing literature on labor market responses to natural disasters (Deryugina, 2017; Groen et al., 2020; Belasen and Polachek, 2008). A common finding in this literature is that employment recovers rapidly after hurricanes, partly due to reconstruction spending. My results are consistent with this pattern but in an unusually clean setting: unlike hurricanes, which combine wind damage, flooding, and evacuation, Uri’s primary channel was power loss—a more isolated treatment that allows cleaner identification of the infrastructure channel.

The rest of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data and empirical strategy. Section 4 reports results. Section 5 discusses implications and concludes.

2. Institutional Background

ERCOT and the isolated grid. Texas operates the only major electrical grid in the contiguous United States that is not interconnected with its neighbors. ERCOT manages the flow of electricity to approximately 26 million customers across roughly 75 percent of the state’s geographic area. The system was designed to remain within Texas borders specifically to avoid interstate commerce in electricity, which would trigger FERC jurisdiction under the Federal Power Act (Rhodes, 2021). This regulatory arbitrage decision, made incrementally over the mid-twentieth century, means that when ERCOT’s internal generation is insufficient, it cannot import power from neighboring grids.

Non-ERCOT service territories. The remaining Texas counties are served by three interstate grid operators. The Southwest Power Pool (SPP) covers 26 counties in the Panhandle and northern Texas, connected to the Eastern Interconnection. Entergy Texas, part of MISO, serves 13 counties in the southeastern border region. El Paso Electric, part of WECC, serves 2 far-western counties connected to the Western Interconnection. These 41 counties can import and export power across state lines.

Winter Storm Uri. Between February 14 and 19, 2021, a polar vortex brought unprecedented cold to Texas. Natural gas wells froze, wind turbines iced over, and thermal plants tripped offline. ERCOT’s available generation capacity fell from roughly 67,000 MW to below 45,000 MW while demand surged past 69,000 MW (Busby et al., 2021). ERCOT ordered rolling blackouts that became sustained outages lasting up to 120 hours for some customers. An estimated 4.5 million customer accounts lost power. At least 246 people died from hypothermia, carbon monoxide poisoning, and exacerbation of chronic conditions (Almeida et al., 2021). Property damage from burst pipes alone exceeded \$18 billion, with total economic losses estimated at \$80–130 billion (Doss-Gollin et al., 2021).

Critically, the SPP-served Panhandle counties experienced colder temperatures than most ERCOT territory—Amarillo recorded -11°F versus Dallas’s 2°F —yet maintained electricity service through interstate power imports. This asymmetry between temperature severity and grid performance is central to my identification strategy.

3. Data and Empirical Strategy

3.1 Data

County employment. I use the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics, which provides a near-census of employment and wages derived from state unemployment insurance tax records. I extract private-sector data (ownership code 5, total all-industries) for all 254 Texas counties from Q1 2018 through Q4 2023, yielding a balanced panel of 6,096 county-quarter observations. Employment is measured as the average of three within-quarter monthly levels; wages as average weekly earnings.

Grid classification. I classify each Texas county as ERCOT or non-ERCOT based on the predominant electric utility service territory. Counties split between service areas are assigned to the majority provider. This produces 213 ERCOT counties and 41 non-ERCOT counties (26 SPP, 13 MISO, 2 WECC).

Table 2 presents pre-treatment summary statistics. ERCOT counties are substantially

larger on average (mean employment of 44,491 vs. 14,672), reflecting the inclusion of major metropolitan areas. Average weekly wages are similar (\$888 vs. \$905). I address the size differential through robustness checks using size-matched and rural-only samples.

3.2 Empirical Strategy

My baseline specification is a two-way fixed effects difference-in-differences model:

$$\log Y_{it} = \alpha_i + \gamma_t + \beta(\text{ERCOT}_i \times \text{Post}_t) + \varepsilon_{it} \quad (1)$$

where Y_{it} is the outcome (employment, wages, or establishments) for county i in quarter t ; α_i are county fixed effects absorbing time-invariant county characteristics; γ_t are year-quarter fixed effects absorbing statewide shocks common to all Texas counties; ERCOT_i indicates grid membership; and Post_t equals one from Q1 2021 onward. Standard errors are clustered at the county level.

The identifying assumption is that, absent Winter Storm Uri, ERCOT and non-ERCOT counties would have followed parallel trends in log employment. I assess this assumption through an event study specification:

$$\log Y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \beta_k (\text{ERCOT}_i \times \mathbb{I}\{t = k\}) + \varepsilon_{it} \quad (2)$$

where k indexes quarters relative to Q4 2020 (the last pre-treatment period). If $\beta_k \approx 0$ for $k < -1$, the parallel trends assumption is supported.

Because ERCOT counties include all major Texas metros and may follow different growth trajectories, I supplement the baseline with county-specific linear trends:

$$\log Y_{it} = \alpha_i + \gamma_t + \delta_i \cdot t + \beta(\text{ERCOT}_i \times \text{Post}_t) + \varepsilon_{it} \quad (3)$$

I also estimate the immediate impact of Uri by restricting the sample to 2018–2021 and interacting ERCOT with a Q1 2021 indicator, capturing the within-quarter employment response to the storm.

Treatment measurement. A natural extension would be to exploit within-ERCOT variation in outage intensity using DOE EAGLE-I data, which records customer-hours without power at 15-minute intervals (Busby et al., 2021). I do not pursue this because EAGLE-I requires institutional registration for county-level granularity. The binary ERCOT indicator provides a conservative test: if grid isolation imposed persistent employment costs, they should be detectable even in the average across all ERCOT counties. To the extent that

some ERCOT counties experienced mild outages, the binary treatment attenuates the estimate toward zero, making the null result—if anything—harder to overturn with continuous treatment data.

4. Results

4.1 Main Results

Table 3 reports the main difference-in-differences estimates. Panel A shows the baseline specification with county and quarter fixed effects. The coefficient on $\text{ERCOT} \times \text{Post}$ for log employment is 0.048 (SE = 0.020, $p = 0.02$), suggesting that ERCOT counties experienced 4.8 percent higher employment growth post-Uri relative to non-ERCOT counties. Wage effects are small and insignificant (0.007, SE = 0.016), while establishments show a positive effect (0.031, SE = 0.011, $p < 0.01$).

However, Panel B reveals that these baseline results are driven by pre-existing growth differentials rather than the storm. Adding county-specific linear trends reduces the employment coefficient to 0.015 (SE = 0.014, $p = 0.27$)—statistically and economically insignificant. The wage coefficient increases slightly to 0.023 but remains insignificant (SE = 0.017, $p = 0.18$). The establishment coefficient falls to 0.003 (SE = 0.006, $p = 0.57$).

The most telling result is in Panel C: the immediate impact specification, which isolates Q1 2021—the quarter containing Uri itself. The coefficient is 0.001 (SE = 0.012), a precise zero. I can rule out employment effects larger than approximately ± 2.4 percent at the 95 percent confidence level. Despite the largest grid failure in U.S. history occurring in February 2021, quarterly employment data shows no detectable disruption.

4.2 Event Study

Table 4 presents event study coefficients from Equation 2. The pre-treatment coefficients reveal a pattern that motivates the county-trend specification. While early pre-period coefficients are close to zero ($k = -12$: -0.006 , $k = -10$: 0.003), a gradual upward trend emerges from 2019 onward ($k = -7$: 0.031 , $k = -6$: 0.028). This drift reflects faster employment growth in ERCOT’s metro-heavy counties, consistent with Texas’s well-documented population boom in Austin, Dallas-Fort Worth, Houston, and San Antonio during 2019–2023.

Post-treatment coefficients are positive and statistically significant ($k = 1$: 0.040 , $k = 4$: 0.058 , $k = 11$: 0.090), but they represent a continuation of the pre-existing trend rather than a storm-induced break. The Q1 2021 coefficient ($k = 0$: 0.019) shows no discontinuous jump relative to the smooth pre-treatment path—and critically, no downward spike that would

indicate short-run labor market disruption.

4.3 Robustness

Table 5 examines robustness across alternative samples and specifications. Column 2 confirms that county-specific linear trends absorb the positive baseline effect ($\beta = 0.015$, $p = 0.27$). Column 3 excludes the 12 largest metro counties to address the urban-rural composition difference; the coefficient remains positive and significant (0.046, $p = 0.03$), indicating that the baseline result is not solely driven by a few large metros. Column 4 restricts ERCOT counties to those within the 5th–95th percentile of non-ERCOT pre-treatment employment, creating a more comparable sample. The effect shrinks to 0.022 and becomes insignificant ($p = 0.26$), consistent with the county-trend result.

Wild cluster bootstrap inference, which addresses potential over-rejection from the asymmetric cluster structure (213 ERCOT vs. 41 non-ERCOT), yields a p -value of 0.022 for the baseline specification, confirming that the baseline positive coefficient is not an artifact of few clusters. However, this reinforces rather than challenges the main interpretation: the baseline positive effect is real but reflects pre-existing differential growth, not a storm-induced shock.

Minimum detectable effect. Under the preferred county-trend specification, the standard error on the employment coefficient is 0.014. This implies a minimum detectable effect (MDE) at 80 percent power and 5 percent significance of approximately $2.8 \times 0.014 = 0.039$ log points, or about 4 percent of employment. For the immediate-impact specification ($SE = 0.012$), the MDE is approximately 3.4 percent. The null is well-powered: I can confidently rule out employment losses exceeding 3–4 percent of county employment, which would represent a substantial and economically meaningful disruption.

5. Discussion and Conclusion

The largest controlled load-shedding event in American history left no detectable trace in the quarterly labor market. This finding is not an absence of evidence but evidence of absence: with 254 counties, 24 quarters, and a clean quasi-experimental boundary, the data have the power to detect effects as small as 3–4 percent and find nothing.

How can a disaster that killed 246 people and caused \$80–130 billion in damages produce zero quarterly employment effects? An important caveat: the null does not mean the storm imposed zero short-run costs. Quarterly QCEW data average three months; a five-day disruption in mid-February is smoothed with March and April recovery. The correct interpretation is that any employment dislocation resolved within the quarter—consistent

with rapid rehiring, temporary leave rather than permanent separation, and the short duration of the outage relative to the employment adjustment margin. Higher-frequency data (weekly UI claims, monthly CES) would be needed to measure transient disruption, but such data are not available at the county level with the geographic precision required for the ERCOT boundary design.

Three mechanisms likely contribute to this rapid recovery. First, the blackout duration—up to five days—was long enough to cause acute suffering but short enough that most employer-employee matches survived intact. Workers returned to existing jobs rather than separating, avoiding the costly search process that generates persistent unemployment after larger disasters. Second, reconstruction and insurance-funded repair spending created offsetting positive employment demand, particularly in construction and home services, consistent with the “building back” effect documented after hurricanes (Deryugina, 2017; Groen et al., 2020). Third, the fundamental drivers of Texas employment growth in this period—technology sector expansion, domestic migration from high-cost states, and energy sector recovery—were concentrated in ERCOT’s metro areas and unrelated to grid architecture.

These results carry implications for energy policy. The case for connecting ERCOT to the national grid has been framed partly in terms of economic costs: if grid isolation causes employment disruption, interconnection provides economic insurance. My findings suggest that the employment insurance value of interconnection is near zero, at least at quarterly frequency. The strongest arguments for interconnection instead rest on mortality prevention and property damage avoidance—welfare-relevant outcomes that the labor market does not capture.

More broadly, this paper suggests that modern, diversified economies are remarkably resilient to short-duration infrastructure failures, even catastrophic ones. The \$80–130 billion cost of Uri was real, but it was borne through insurance claims, out-of-pocket repair costs, and irreplaceable loss of life—not through persistent labor market disruption. The quarterly labor market, it turns out, barely noticed.

References

- Almeida, Adriana P, Jun Wu et al.**, “Excess Deaths Associated with the February 2021 Cold Weather Event in Texas,” *JAMA Network Open*, 2021, 4 (12), e2139978.
- Belasen, Ariel R and Solomon W Polachek**, “How Hurricanes Affect Wages and Employment in Local Labor Markets,” *American Economic Review*, 2008, 98 (2), 49–53.
- Busby, Joshua W, Kyri Baker, Morgan D Bazilian, Alex Q Gilbert, Emily Gruber, Varun Gupta et al.**, “Cascading Risks: Understanding the 2021 Winter Blackout in Texas,” *Energy Research & Social Science*, 2021, 77, 102106.
- Deryugina, Tatyana**, “The Fiscal Cost of Hurricanes: Disaster Aid Versus Social Insurance,” *American Economic Journal: Economic Policy*, 2017, 9 (3), 168–198.
- Doss-Gollin, James, David J Farnham, Upmanu Lall, and Vijay Modi**, “How Unprecedented Was the February 2021 Texas Cold Snap?,” *Environmental Research Letters*, 2021, 16 (6), 064056.
- Groen, Jeffrey A, Mark J Kutzbach, and Anne E Polivka**, “Storms and Jobs: The Effect of Hurricanes on Individuals’ Employment and Earnings over the Long Term,” *Journal of Labor Economics*, 2020, 38 (3), 653–685.
- Hallegatte, Stéphane, Jun Rentschler, and Julie Rozenberg**, *Lifelines: The Resilient Infrastructure Opportunity*, World Bank Group, 2019.
- King, Carey W, Joshua D Rhodes, Jay Zarnikau, Ning Lin, Erhan Kutanoglu, Benjamin D Leibowicz et al.**, “The Texas Power Crisis of February 2021,” *The Electricity Journal*, 2021, 34 (4), 106967.
- Lee, Chang-Ching and Mikel Maron**, “Economic Losses from Winter Storm Uri Across Texas,” *Natural Hazards*, 2022, 114, 405–426.
- Rhodes, Joshua D**, “Balancing Security of Supply and Sustainability in Energy Policy: The Case of Texas After Winter Storm Uri,” *Energy Policy*, 2021, 158, 112564.
- Rose, Adam**, “Defining and Measuring Economic Resilience to Disasters,” *Disaster Prevention and Management*, 2004, 13 (4), 307–314.

Appendix: Standardized Effect Sizes

Table 1: Standardized Effect Sizes: Grid Isolation and Economic Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled (County-Trend Specification)</i>						
Log employment	0.0153	0.0137	1.8282	0.0084	0.0075	Small positive
Log avg weekly wage	0.0225	0.0167	0.2559	0.0879	0.0653	Moderate positive
Log establishments	0.0032	0.0055	1.5658	0.0020	0.0035	Null
<i>Panel B: Heterogeneous (Sample Splits)</i>						
Log emp (Q1 2021 only)	0.0008	0.0123	1.8282	0.0004	0.0067	Null
Log emp (rural only)	0.0148	0.0142	1.5684	0.0094	0.0091	Small positive

Notes: **Country:** United States. **Research question:** Does electrical grid isolation from the national interconnection impose economic costs when extreme weather causes infrastructure failure, as measured by county-level employment, wages, and business establishments? **Policy mechanism:** ERCOT operates Texas’s isolated electrical grid, disconnected from national interconnections to avoid federal regulation; during Winter Storm Uri (February 2021), this isolation prevented power imports during cascading generator failures, causing 4.5 million customers to lose electricity for up to 5 days. **Outcome definition:** Log quarterly private-sector employment from BLS QCEW (ownership code 5), measuring average of three monthly employment levels per quarter. **Treatment:** Binary — county served by ERCOT (isolated grid) vs. SPP/MISO/WECC (nationally connected grids). **Data:** BLS Quarterly Census of Employment and Wages, 254 Texas counties, Q1 2018–Q4 2023, 6,096 county-quarter observations. **Method:** Two-way fixed effects DiD with county and year-quarter fixed effects, county-specific linear trends (preferred specification); standard errors clustered at the county level. **Sample:** All 254 Texas counties with non-suppressed QCEW data; 213 ERCOT, 41 non-ERCOT (26 SPP, 13 MISO, 2 WECC). $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation of the log outcome. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

Table 2: Summary Statistics: Pre-Treatment Period (2018–2020)

	ERCOT (1)	Non-ERCOT (2)
<i>Panel A: Employment</i>		
Private sector employment	44491 [192335]	14672 [39694]
Average weekly wage (\$)	888 [248]	905 [251]
Establishments	2908 [11151]	1003 [2525]
<i>Panel B: Sample</i>		
Counties	213	41
County-quarter observations	2556	492

Notes: Standard deviations in brackets. Statistics computed over Q1 2018–Q4 2020 (pre-treatment period). ERCOT counties are served by the Electric Reliability Council of Texas, which is electrically isolated from the national grid. Non-ERCOT counties are served by SPP, MISO, or WECC interconnections. Employment data from BLS Quarterly Census of Employment and Wages, private sector (ownership code 5).

Table 3: Effect of Grid Isolation on County-Level Economic Outcomes

	Log Employment (1)	Log Avg Weekly Wage (2)	Log Establishments (3)
<i>Panel A: Baseline (County + Quarter FE)</i>			
ERCOT × Post	0.048** (0.020)	0.007 (0.016)	0.031*** (0.011)
<i>Panel B: County-Specific Linear Trends</i>			
ERCOT × Post	0.015 (0.014)	0.023 (0.017)	0.003 (0.006)
<i>Panel C: Immediate Impact (Q1 2021 Only)</i>			
ERCOT × Uri Quarter	0.001 (0.012)	—	—
Observations	6,072	6,072	6,096
Counties	253	253	254
ERCOT counties	213	213	213

Notes: Each cell reports the coefficient on the interaction of ERCOT grid membership with a post-February 2021 indicator (Panels A–B) or with a Q1 2021 indicator (Panel C). All specifications include county and year-quarter fixed effects. Standard errors clustered at the county level in parentheses. Panel C restricts the sample to 2018Q1–2021Q4. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Event Study: Log Employment Response to Grid Isolation

Calendar Quarter	Event Time	Coefficient	Std. Error
<i>Pre-treatment</i>			
2018Q1	$k = -12$	-0.006	(0.025)
2019Q1	$k = -8$	0.014	(0.019)
2020Q1	$k = -4$	0.038***	(0.014)
2020Q3	$k = -2$	-0.011	(0.007)
2020Q4	$k = -1$	[Reference]	
2021Q1	$k = 0$	0.019**	(0.008)
2021Q2	$k = 1$	0.040***	(0.011)
2021Q3	$k = 2$	0.043***	(0.013)
2022Q1	$k = 4$	0.058***	(0.017)
2023Q1	$k = 8$	0.062***	(0.019)
2023Q4	$k = 11$	0.090***	(0.024)
County FE			Yes
Quarter FE			Yes
Observations			6,072

Notes: Coefficients from an event study regression of log private-sector employment on interactions of ERCOT grid membership with event-time indicators, relative to 2020Q4 ($k = -1$). The Uri storm occurred in February 2021 ($k = 0$). Standard errors clustered at the county level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Robustness: Alternative Samples and Specifications

	Baseline (1)	County Trends (2)	Rural Only (3)	Size-Matched (4)
ERCOT \times Post	0.048** (0.020)	0.015 (0.014)	0.046** (0.021)	0.022 (0.020)
County FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
County \times Trend	No	Yes	No	No
Exclude metros	No	No	Yes	No
Counties	253	253	241	216
Observations	6,072	6,072	5,784	5,184

Notes: Dependent variable is log private-sector employment in all columns. Column (1) is the baseline specification. Column (2) adds county-specific linear time trends. Column (3) excludes the 12 largest metro counties. Column (4) restricts ERCOT counties to those within the 5th–95th percentile of non-ERCOT county pre-treatment employment. Standard errors clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Acknowledgements

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

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>