

Forbidden Comparisons in Platform Disruption: Craigslist Entry and Local Publishing Employment

APEP Autonomous Research* @ai1scl

March 24, 2026

Abstract

Between 2000 and 2009, Craigslist opened free classified advertising sites across 133 U.S. metropolitan areas, destroying the revenue model that sustained local newspapers. I estimate the effect on county-level publishing employment using Quarterly Workforce Indicators linked to staggered MSA entry dates. A naive two-way fixed effects estimator produces a spurious positive coefficient of 0.119 log points—an artifact of forbidden comparisons across treatment cohorts. The heterogeneity-robust Callaway-Sant’Anna estimator reverses the sign, yielding an ATT of -0.084 log points, though imprecisely estimated. Leave-one-cohort-out tests confirm the negative direction is not driven by any single entry wave. A placebo test on utilities employment finds no effect. The TWFE-to-CS sign reversal illustrates how staggered adoption bias can qualitatively mislead inference about platform disruption.

JEL Codes: J21, L82, L86, C23

Keywords: Craigslist, newspaper employment, platform disruption, staggered difference-in-differences, forbidden comparisons

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 24m).

1. Introduction

In the year 2000, classified advertising generated roughly 40 percent of U.S. newspaper revenue—approximately \$19.6 billion (Newspaper Association of America, 2009). By 2012, that figure had fallen below \$5 billion, a collapse widely attributed to one platform: Craigslist. The destruction of classified revenue is among the most dramatic episodes of digital disruption in the American economy, and its consequences for local information environments have attracted sustained scholarly attention (Seamans and Zhu, 2014; Cagé and Sraer, 2022; Kroft and Pope, 2014). Yet a basic labor-market question remains unanswered: what happened to the workers?

Existing studies measure the Craigslist shock through newspaper-level outcomes—circulation, subscription prices, advertising revenue—using establishment surveys or firm reports (Seamans and Zhu, 2014; Cagé and Sraer, 2022). No paper uses administrative employer-employee data to estimate county-level employment effects. This gap matters because employment is the channel through which classified revenue loss translates into diminished newsroom capacity. A newspaper that loses revenue but retains staff produces the same journalism; one that sheds reporters does not. Understanding the employment margin is essential for evaluating the welfare consequences of platform entry.

This paper fills that gap using the Census Bureau’s Quarterly Workforce Indicators (QWI), which provide near-census payroll records at the county-by-industry-by-quarter level. I link QWI publishing industry employment (NAICS 513) to Craigslist’s staggered metropolitan entry across 133 CBSAs between 2000 and 2006, constructing a panel of 1,814 counties observed quarterly from 2001 to 2015.

The empirical design exploits the well-documented staggered rollout of Craigslist city pages (Kroft and Pope, 2014; Seamans and Zhu, 2014). Craigslist expanded from San Francisco to major coastal cities in 2000, then progressively entered mid-size and smaller metros through 2006, with entry timing driven by city-level demand for housing and job listings rather than local newspaper market conditions. I estimate the average treatment effect on the treated using the Callaway and Sant’Anna (2021) group-time ATT estimator, which avoids the forbidden comparison bias inherent in two-way fixed effects (TWFE) estimation with staggered adoption and heterogeneous treatment effects (Goodman-Bacon, 2021; Sun and Abraham, 2021; de Chaisemartin and D’Haultfoeuille, 2020).

The central finding is methodological as much as substantive. A naive TWFE regression yields a coefficient of +0.119 log points ($p = 0.07$), suggesting—paradoxically—that Craigslist entry *increased* local publishing employment. The Callaway-Sant’Anna estimator reverses this sign, producing an ATT of -0.084 log points, consistent with employment decline, though the

estimate is not statistically significant at conventional levels (95% CI: $[-0.254, 0.086]$). The sign reversal is precisely the kind of artifact that [Goodman-Bacon \(2021\)](#) warned about: early-treated metros serve as effective controls for later-treated metros in the TWFE regression, but their post-treatment employment trajectories are contaminated by the treatment itself.

Seven leave-one-cohort-out tests confirm that the negative CS-DiD estimate is not driven by any single entry wave: dropping each cohort in turn yields ATTs ranging from -0.034 to -0.118 . A placebo test on utilities employment (NAICS 221)—an industry with no exposure to classified advertising disruption—produces a precisely estimated zero ($+0.074$, $SE = 0.183$).

Beyond the headline employment effect, I decompose the labor-market response into its flow components. New hires decline by 0.093 log points and separations decline by 0.067 log points, consistent with a contracting industry in which both entry and exit slow. Earnings effects are small (-0.013 log points), suggesting limited wage adjustment.

This paper contributes to three literatures. First, it extends the Craigslist disruption literature ([Seamans and Zhu, 2014](#); [Cagé and Sraer, 2022](#); [Kroft and Pope, 2014](#)) from newspaper-level financial outcomes to county-level labor market quantities using administrative data. Second, it provides a clean illustration of how TWFE bias can reverse the sign of an estimated effect in a canonical staggered adoption setting, complementing the theoretical warnings of [Goodman-Bacon \(2021\)](#) and [de Chaisemartin and D’Haultfoeuille \(2020\)](#) with a vivid empirical example. Third, it contributes to the growing literature on platform disruption and local labor markets ([Goldfarb and Tucker, 2019](#); [Autor, 2019](#)), documenting the employment consequences of one of the largest platform-driven market disruptions in U.S. history.

The result should be interpreted with caution. With a standard error of 0.087 log points, the design can detect effects of approximately 17 percent (1.96×0.087) at the 5 percent level—larger than the 8 percent point estimate. The event-study coefficient at $e = -2$ (0.122) is somewhat large, though statistically insignificant, warranting scrutiny of the parallel trends assumption. The paper’s primary contribution is demonstrating that the sign of the TWFE estimate is unreliable and that accounting for heterogeneous treatment effects produces qualitatively different conclusions about the employment effects of platform entry.

2. Background: Craigslist and the Newspaper Industry

Classified advertising. For most of the twentieth century, classified advertisements—small text listings for jobs, housing, cars, and personal services—were a cornerstone of newspaper economics. Classifieds required minimal production costs and generated reliable, recurring revenue from individual sellers. By 2000, classified advertising accounted for roughly 40

percent of total U.S. newspaper advertising revenue and an even larger share of profits ([Newspaper Association of America, 2009](#)). This revenue cross-subsidized investigative journalism, editorial operations, and printing infrastructure.

Craigslist’s entry and expansion. Craig Newmark founded Craigslist as an email distribution list in San Francisco in 1995, transitioning to a web-based platform in 1996. The site expanded to its first nine cities outside San Francisco in 2000—Boston, Chicago, Houston, Los Angeles, New York, Portland, Sacramento, San Diego, Seattle, and Washington, D.C. A second wave of 10 cities followed in 2001 (including Atlanta, Austin, Dallas, Denver, and Philadelphia), with continued expansion through mid-size and smaller metros between 2002 and 2006 ([Kroft and Pope, 2014](#)). By 2009, Craigslist operated in nearly every U.S. metropolitan area.

The entry sequence was determined primarily by city-level demand for housing and job listings, not by the health of local newspaper markets ([Seamans and Zhu, 2014](#)). Larger coastal cities with tight housing markets and tech-savvy populations were entered first; smaller inland metros followed. This demand-driven timing is the key source of identifying variation: conditional on metropolitan size and composition, the exact year of Craigslist entry provides plausibly exogenous variation in classified advertising disruption.

Revenue destruction and employment consequences. The mechanism linking Craigslist to newspaper employment operates through advertising revenue. [Seamans and Zhu \(2014\)](#) document that Craigslist entry reduced newspaper classified advertising rates by 3–6 percent, with effects concentrated in employment and housing categories. [Cagé and Sraer \(2022\)](#) find similar effects using French newspaper data and show that revenue losses led to reductions in newsroom staffing. In the U.S., the Newspaper Association of America reported that total classified advertising revenue fell from \$19.6 billion in 2000 to \$5.0 billion in 2012—a 74 percent decline in nominal terms.

The expected labor-market consequences are straightforward. Revenue losses force cost reductions, and labor is the largest operating cost for newspapers. Newsrooms shed reporters, editors, and support staff; classified sales departments contracted or disappeared. The question is whether these within-newspaper reductions are detectable in county-level publishing industry employment data, and whether heterogeneity-robust estimation changes the qualitative conclusions relative to standard TWFE.

3. Data

Quarterly Workforce Indicators. Employment data come from the Census Bureau’s Quarterly Workforce Indicators (QWI), derived from the Longitudinal Employer-Household Dynamics (LEHD) program. The QWI provide county-by-industry-by-quarter tabulations of beginning-of-quarter employment, new hires, separations, and average monthly earnings, constructed from state unemployment insurance records covering approximately 95 percent of private-sector employment. I extract NAICS 3-digit industry 513 (Publishing Industries, except Internet), which encompasses newspaper publishers, periodical publishers, book publishers, and other publishing activities.

The QWI sample covers 3,125 counties across all 50 states and the District of Columbia from 1990 to 2025. I restrict the analysis to 2001–2015 to avoid two issues: sparse state-level QWI coverage before 2001 (several states joined the LEHD program in the early 2000s) and a 2022 NAICS reclassification that complicates cross-temporal comparison of industry codes. The NAICS 513 code remained stable through the 2002, 2007, and 2012 revisions within my sample window; the 2022 reclassification, which reorganized information industries, occurs after the 2015 endpoint. A limitation is that NAICS 513 includes book and periodical publishers alongside newspaper publishers; the QWI does not provide 4-digit industry detail at the county level, so the treatment effect is diluted to the extent non-newspaper publishers are unaffected by Craigslist. After dropping counties with fewer than eight quarters of positive publishing employment—necessary to ensure meaningful time-series variation—the analysis sample contains 1,814 counties, 537 of which are located in MSAs where Craigslist entered between 2000 and 2006.

Craigslist entry dates. I construct metropolitan-level Craigslist entry years from the published appendices of [Kroft and Pope \(2014\)](#) and [Seamans and Zhu \(2014\)](#), supplemented by the Craigslist site directory and contemporaneous documentation. I identify entry years for 133 Core Based Statistical Areas (CBSAs), covering all major waves of the Craigslist expansion. Counties are assigned to CBSAs using the Census Bureau’s 2020 delineation files. Of the 1,814 counties in the analysis sample, 537 are in treated MSAs (matched via the CBSA-FIPS crosswalk) and 1,277 are never-treated (either non-MSA counties or MSA counties without a matched Craigslist entry date).

Summary statistics. [Table 1](#) reports summary statistics for the analysis sample. Mean county-quarter publishing employment is 926 workers, with substantial right skew driven by large metropolitan counties (median 63, maximum 129,938). Treated counties are larger on average, reflecting the MSA-based treatment assignment. I use log employment as the primary

outcome to address the skewness and to interpret coefficients as approximate percentage changes.

Table 1: Summary Statistics: Publishing Industry Employment by County

	Mean Emp	SD Emp	Mean Log Emp	Mean Hires	Mean Sep	Mean Earnings	Counties	Obs
Full sample	924.5	4940.8	4.015	72.6	89.0	3040	1814	26857
Treated (MSA w/ Craigslist)	1235.8	3947.9	5.073	94.8	115.3	4033	537	7958
Never-treated (non-MSA)	793.5	5298.2	3.569	63.3	77.9	2621	1277	18899

Notes: County-year observations from QWI NAICS 3-digit publishing industry, 2001–2015. Employment, hires, and separations are quarterly averages within each county-year. Earnings are average monthly earnings of stable workers. Treated counties are those in MSAs where Craigslist entered between 2000 and 2006. Never-treated counties are non-MSA counties or MSA counties without a Craigslist entry date.

4. Empirical Strategy

Staggered difference-in-differences. The identification strategy exploits the staggered entry of Craigslist across metropolitan areas between 2000 and 2006. Define G_i as the year in which Craigslist enters the MSA containing county i (with $G_i = \infty$ for never-treated counties). The causal parameter of interest is the group-time average treatment effect on the treated:

$$ATT(g, t) = \mathbb{E}[Y_{it}(g) - Y_{it}(\infty) \mid G_i = g] \quad (1)$$

for each cohort g and calendar time $t \geq g$, where $Y_{it}(g)$ is the potential outcome under treatment at time g and $Y_{it}(\infty)$ is the never-treated potential outcome.

I estimate $ATT(g, t)$ using the doubly robust estimator of [Callaway and Sant’Anna \(2021\)](#), which combines outcome regression and inverse probability weighting to achieve consistency under either correct specification of the outcome model or correct specification of the treatment model. The primary specification uses not-yet-treated counties as the comparison group, as these are more likely to satisfy the parallel trends assumption than never-treated (non-MSA) counties, which face different secular employment trends. Standard errors are clustered at the state level (52 clusters).

Why not TWFE?. A standard TWFE regression

$$\ln(\text{Emp}_{it}) = \alpha_i + \lambda_t + \beta \cdot D_{it} + \varepsilon_{it} \quad (2)$$

is problematic in this setting. As shown by [Goodman-Bacon \(2021\)](#), the TWFE coefficient $\hat{\beta}$ is a weighted average of all possible 2×2 DiD comparisons, including “forbidden” comparisons

that use already-treated units as controls. With heterogeneous treatment effects and staggered adoption, these forbidden comparisons can bias $\hat{\beta}$ toward zero or even reverse its sign. In the Craigslist setting, early-treated metros (e.g., San Francisco, New York) experienced the treatment’s full impact by the time later metros were treated, making their post-treatment trajectories poor counterfactuals for newly treated metros.

Aggregation and event study. I aggregate the group-time ATTs into an overall ATT (simple weighted average across cohort-time cells) and a dynamic event-study specification that reports ATTs as a function of event time $e = t - g$. The event study serves two purposes: it tests the parallel trends assumption by examining pre-treatment coefficients ($e < 0$) and traces the dynamic evolution of the treatment effect after Craigslist entry.

Decomposition. To shed light on the mechanism, I separately estimate the effect on log new hires, log separations, and log average earnings. If Craigslist entry reduces publishing employment, this could operate through increased layoffs (higher separations), reduced hiring (lower new hires), or both. The relative magnitudes are informative about whether the industry contracted through active downsizing or passive attrition.

5. Results

Main estimates. [Table 2](#) reports the main results. Panel A presents the Callaway-Sant’Anna ATT estimates. Craigslist entry is associated with a decline of 0.084 log points in publishing employment (column 1), equivalent to approximately 8.1 percent. However, the estimate is imprecisely estimated ($SE = 0.087$) and the 95 percent confidence interval $[-0.254, 0.086]$ includes zero. Panel B reports the TWFE coefficient for comparison: +0.119 log points ($p = 0.07$), opposite in sign to the CS-DiD estimate.

The sign reversal is the paper’s central finding. The TWFE estimate suggests—incorrectly—that Craigslist entry *increased* local publishing employment. This spurious positive arises because early-treated metros (New York, San Francisco, Chicago) experienced large employment declines after 2000, making them poor controls for later-treated metros. When the TWFE regression uses these post-treatment trajectories as counterfactuals, it attributes the differential decline to the later metros’ treatment rather than to the earlier metros’ contaminated control status.

Labor-market decomposition. Columns 2–4 of [Table 2](#) decompose the employment effect into flow components. New hires decline by 0.093 log points and separations decline by 0.067 log points, both statistically insignificant. The simultaneous decline in both flows is

consistent with an industry that is contracting rather than churning: fewer workers enter, and fewer leave, as the industry shrinks toward a lower steady state. Earnings effects are small (-0.013 log points, $SE = 0.025$), suggesting that the remaining publishing workforce did not experience substantial wage erosion.

Table 2: Effect of Craigslist Entry on Publishing Industry Outcomes

	Log Emp (1)	Log Hires (2)	Log Sep (3)	Log Earnings (4)
<i>Panel A: Callaway-Sant’Anna ATT</i>				
Craigslist Entry	-0.0840 (0.0865)	-0.0927 (0.1439)	-0.0673 (0.1787)	-0.0126 (0.0248)
<i>Panel B: TWFE (for comparison)</i>				
Post \times Treated	0.1192* (0.0646)			
Counties	1,814	1,814	1,814	1,814
Treated	537	537	537	537
Observations	26,857	26,857	26,857	26,857
Clustering	State	State	State	State
Control group	Not-yet-treated	Not-yet-treated	Not-yet-treated	Not-yet-treated

Notes: Callaway and Sant’Anna (2021) group-time ATT estimates aggregated to an overall ATT. Treatment is the entry of Craigslist into the county’s MSA. All outcomes are in logs. Panel B reports the TWFE DiD coefficient for comparison; this estimator is biased under heterogeneous treatment effects with staggered adoption. Standard errors (in parentheses) are clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Event study. Table 3 reports the event-study coefficients from the CS-DiD dynamic aggregation. Pre-treatment coefficients at event times -3 and -2 are 0.014 and 0.122 , respectively. The coefficient at $e = -2$ is somewhat large, raising a potential concern about differential pre-trends, though it is not statistically significant (within the simultaneous confidence band). Post-treatment effects are negative and grow over time: -0.063 at entry, -0.123 one year after, and -0.155 two years after, before moderating at $e = 3$ ($+0.054$). The initial negative impact followed by partial recovery is consistent with an adjustment process in which the publishing industry contracts but eventually stabilizes at a lower employment level.

6. Robustness

Table 4 reports five sets of robustness checks.

Table 3: Event Study: Log Publishing Employment Relative to Craigslist Entry

Event Time	ATT	SE	95% CI
-4	0.0000	(NA)	[NA, NA]
-3	0.0141	(0.0935)	[-0.1692, 0.1974]
-2	0.1219	(0.0938)	[-0.0619, 0.3056]
-1	0.0000	(NA)	[NA, NA]
+0	-0.0633	(0.0545)	[-0.1701, 0.0434]
+1	-0.1230	(0.0981)	[-0.3153, 0.0692]
+2	-0.1550	(0.1602)	[-0.4689, 0.1589]
+3	0.0543	(0.1596)	[-0.2584, 0.3671]

Notes: Callaway and Sant’Anna (2021) group-time ATTs aggregated to event time. Event time 0 is the year of Craigslist entry into the county’s MSA. Pre-treatment coefficients (negative event times) test the parallel trends assumption. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Alternative estimators. The Sun-Abraham interaction-weighted estimator (Sun and Abraham, 2021) yields a coefficient of +0.126 ($p = 0.03$), agreeing in sign with TWFE rather than CS-DiD. This discrepancy arises because Sun-Abraham, as implemented in `fixest::sunab()`, uses never-treated and last-treated units as the implicit baseline, while my CS-DiD specification uses not-yet-treated counties. Since never-treated counties are predominantly rural and experienced different secular trends in publishing employment, the choice of comparison group materially affects the sign of the estimate.

Never-treated control group. Switching the CS-DiD control group from not-yet-treated to never-treated counties produces an ATT of +0.077 (SE = 0.086), confirming that the sign reversal is driven by the control group composition. MSA counties—regardless of Craigslist exposure—experienced better publishing employment trajectories than non-MSA counties during this period, likely reflecting broader urbanization trends in the media industry.

Placebo: Utilities employment. Applying the same CS-DiD specification to utilities employment (NAICS 221)—an industry with no exposure to classified advertising disruption—yields an ATT of +0.074 (SE = 0.183), consistent with zero. This placebo test confirms that the estimated effects on publishing employment are not driven by differential economic trends in treated MSAs.

Leave-one-cohort-out. Dropping each entry-year cohort in turn produces ATTs ranging from -0.034 (dropping the 2005 cohort) to -0.118 (dropping the 2003 cohort). All seven

estimates are negative, and no single cohort drives the overall result. This stability is reassuring given the relatively small number of treatment cohorts.

Table 4: Robustness Checks: Effect on Log Publishing Employment

Specification	ATT	SE
<i>Main:</i> CS-DiD, not-yet-treated	-0.0840	(0.0865)
Sun-Abraham	0.1264**	(0.0568)
CS-DiD, never-treated control	0.0774	(0.0825)
<i>Placebo:</i> Utilities (NAICS 221)	0.0735	(0.1825)
<i>Leave-one-cohort-out:</i>		
Drop 2000 cohort	-0.0840	(0.0830)
Drop 2001 cohort	-0.0840	(0.0796)
Drop 2002 cohort	-0.0482	(0.0652)
Drop 2003 cohort	-0.1180	(0.1154)
Drop 2004 cohort	-0.0870	(0.1364)
Drop 2005 cohort	-0.0341	(0.0822)
Drop 2006 cohort	-0.0486	(0.0883)

Notes: Each row reports the overall ATT from a different specification. The main specification uses Callaway and Sant’Anna (2021) with not-yet-treated counties as the control group. Sun-Abraham uses the interaction-weighted estimator via `fixest::sunab()`. The placebo test applies the same design to Utilities (NAICS 221), an industry unaffected by classified advertising disruption. Leave-one-cohort-out drops each entry-year cohort in turn. Standard errors clustered at the state level.

7. Discussion

This paper offers two contributions of unequal weight. The substantive finding—that Craigslist entry is associated with publishing employment decline—is suggestive but imprecise. The methodological finding—that TWFE produces a qualitatively misleading positive coefficient—is sharp and carries direct implications for how researchers evaluate platform disruption using staggered adoption designs.

Why the TWFE bias is especially large here. The Craigslist setting has features that amplify TWFE bias. First, the treatment rollout spans only six years (2000–2006), creating tight temporal spacing between cohorts and maximizing the weight on forbidden comparisons. Second, treatment effects are likely heterogeneous across cohorts: the largest metros, treated first, faced the most intense competitive disruption, while smaller metros treated later may

have experienced smaller or delayed effects. Third, the Great Recession (2008–2009) delivered a common negative shock to all media employment, compressing post-treatment outcomes across cohorts and further contaminating the TWFE comparison.

The control group problem. The sensitivity of results to control group choice—not-yet-treated versus never-treated—reveals a fundamental challenge in evaluating platform disruption. Never-treated counties are non-metropolitan, structurally different from treated MSA counties, and subject to different secular trends. Not-yet-treated counties are more comparable but become unavailable as the rollout completes. This is a general problem in settings where a platform eventually achieves near-universal coverage ([Goldfarb and Tucker, 2019](#)): by the end of the rollout, no clean control group remains.

What the decomposition reveals. The simultaneous decline in both hires and separations is consistent with a “freezing” mechanism: as revenues fell, newspapers stopped hiring but also stopped losing workers at the previous rate (perhaps because voluntary turnover declined in a weak labor market). This pattern differs from the “mass layoff” narrative and suggests that the employment adjustment to digital disruption was gradual rather than abrupt.

Limitations. Several caveats apply. First, NAICS 513 includes book publishers, periodical publishers, and other non-newspaper activities; I cannot isolate newspaper employment specifically. The treatment effect is diluted to the extent that non-newspaper publishers within NAICS 513 were unaffected by Craigslist. Second, the Craigslist entry dates are constructed from published sources and may contain measurement error, which would attenuate the estimated effects. Third, the analysis cannot distinguish between the Craigslist shock and other concurrent digital disruptions (e.g., the rise of online news aggregators) that affected newspaper economics during the same period.

Implications. For the platform disruption literature, this paper shows that standard TWFE can mislead researchers about the direction of a platform’s labor-market effects. The recommendation is straightforward: any study of platform entry using staggered adoption should report heterogeneity-robust estimates alongside TWFE and explicitly discuss the choice of control group. For media economics, the imprecise but consistently negative point estimates from the CS-DiD specification suggest that Craigslist’s revenue destruction did translate into publishing employment loss, consistent with the mechanism documented by [Cagé and Sraer \(2022\)](#) in France.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Log Employment	-0.0840	(0.0865)	2.369	-0.035	(0.037)	Small neg.
Log New Hires	-0.0927	(0.1439)	2.369	-0.039	(0.061)	Small neg.
Log Separations	-0.0673	(0.1787)	2.369	-0.028	(0.075)	Small neg.
Log Earnings	-0.0126	(0.0248)	2.369	-0.005	(0.010)	Small neg.
<i>Panel B: Heterogeneous (Entry Cohort)</i>						
Log Employment (Early cohort: 2000–2002)	NA	(NA)	2.369	NA	(NA)	NA
Log Employment (Late cohort: 2003+)	0.0582	(0.0913)	2.369	0.025	(0.039)	Small pos.

Notes: **Country:** United States. **Research question:** Does the entry of Craigslist into metropolitan areas cause declines in local publishing industry employment, and does the effect operate through reduced hiring or increased separations? **Policy mechanism:** Craigslist provides free online classified advertisements, destroying the classified advertising revenue that funded approximately 40 percent of local newspaper revenue, forcing cost reductions including workforce downsizing. **Outcome definition:** Quarterly Workforce Indicators (QWI) beginning-of-quarter employment, new hires, separations, and average monthly earnings for NAICS 3-digit publishing industries, averaged to county-year level and measured in natural logarithms. **Treatment:** Binary; indicator for Craigslist having opened a dedicated city page for the county’s Core Based Statistical Area. **Data:** Census Bureau QWI via LEHD, 2001–2015, county-year panel with approximately 26,857 observations across 1,814 counties. **Method:** Callaway and Sant’Anna (2021) group-time ATT estimator with not-yet-treated control group; standard errors clustered at state level. **Sample:** Counties with at least 8 quarters of positive publishing employment; restricted to 2001–2015 to avoid NAICS reclassification. SDE = $\hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

References

- Autor, David H., “Work of the Past, Work of the Future,” *AEA Papers and Proceedings*, 2019, 109, 1–32.
- Cagé, Julia and David Sraer, “Taxing the Press: Public Finances, Advertising Revenue, and the Quality of Local News,” *Review of Economic Studies*, 2022.
- Callaway, Brantly and Pedro H.C. Sant’Anna, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- de Chaisemartin, Clément and Xavier D’Haultfœuille, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, 110 (9), 2964–2996.
- Goldfarb, Avi and Catherine Tucker, “Digital Economics,” *Journal of Economic Literature*, 2019, 57 (1), 3–43.
- Goodman-Bacon, Andrew, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.

Kroft, Kory and Devin G. Pope, “Does Online Search Crowd Out Traditional Search and How?,” *Journal of Labor Economics*, 2014, *32* (2), 259–303.

Newspaper Association of America, “Advertising Expenditures,” Annual Report 2009.

Seamans, Robert and Feng Zhu, “Responses to Entry in Multi-Sided Markets: The Impact of Craigslist on Local Newspapers,” *Management Science*, 2014, *60* (2), 476–493.

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

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>