

Membership grew faster than employment: A sixteen-year panel of inside electrical construction in hyperscale data center metros, 2010 to 2025

Research Note, version 1.0. 20 September 2026. This note is deposited under its own concept DOI, 10.5281/zenodo.22866646. It rebuilds and supersedes the findings of the working paper at 10.5281/zenodo.22310706, which remains separately deposited and citable under concept DOI 10.5281/zenodo.22310705 and is not withdrawn. The two are different products: the working paper reports an eight-year panel opening in 2018, this note a sixteen-year panel opening in 2010, and the central estimate’s significance differs between them for the reasons given below.

What this note corrects in the superseded version. Three things.

The working paper could observe one usable pre-treatment period and could not test the parallel-trends assumption its design required. This note rebuilds the panel over sixteen years and reports the test.

The working paper reported “a randomization p-value below 0.001,” computed by permuting treatment across individual locals. Treatment is assigned to a metropolitan area and several locals share one, so permuting locals independently destroys the clustering in the assignment mechanism and yields a null distribution too tight by a factor of about three. The correct units are whole metropolitan areas and whole states, which give 0.0136 and 0.0102 on this panel. Both remain below conventional thresholds, so the direction of the earlier conclusion is unchanged and its stated precision was overstated.

Metropolitan assignment matched a city-name prefix, which returns nothing for a local headquartered in a suburb absent from the prefix list even where the address geocodes to a treated area. Local 150 of Libertyville, Illinois, whose county sits in the Chicago area, was counted as a control. It is treated here. The eight-year estimates the working paper reported are unaffected, because that local does not enter the 2018-base sample; the sixteen-year estimates in this note are computed after the correction. `LM2_CLASSIFICATION_ERROR_2026-09-20.md`, in the dataset at 10.5281/zenodo.22862472, records the sweep that found it and the effect on every specification.

Tommy N. Turner. Independent researcher, Williamsburg, Virginia. tommy@gofast.pro.

Deposit. Data, code and retrieved sources are deposited as a separate citable dataset under concept DOI 10.5281/zenodo.22862472, cited by this note and by the companion measurement paper alike. Every figure in the tables is produced by a script deposited there. Table 1 comes from `lm2_identification_v3.py`, Table 2 from `lm2_v3_0.py`, Table 3’s membership column from `lm2_identification_v3.py` and its employment columns from `lm2_qcew_event.py`, and Table 4 from `lm2_referee_checks_v3.py` and `lm2_identification_v3.py`.

Disclosure. The author serves on the advisory board of an Alfred P. Sloan Foundation-funded study at Emory University and the University of Tennessee on data centers and energy in Virginia, Georgia and Tennessee. This note was produced independently and is not a product of that study. The author has no financial interest in any data center operator, electrical contractor or labor organisation.

Use of generative artificial intelligence. The author used a large language model in preparing this work and is responsible for its content.

Abstract

A prior working paper measured inside electrical union membership in metropolitan areas where hyperscale data center operators built campuses, using an eight-year panel that opened in 2018. Several treated metros already had operating campuses in that base year, so the design could observe only one usable pre-treatment period and could not test the parallel-trends assumption it required. This note rebuilds the panel over sixteen years, 2010 to 2025, on the same specification and the same sources. It adds an independent employment series covering every year of the panel. Seven coefficients for 2010 through 2016 are individually insignificant and jointly indistinguishable from zero, $\chi^2(7) = 9.40$, $p = 0.225$; four treated metros carry a site dated before 2017, so those seven coefficients test whether the two groups were already diverging rather than bounding a clean pre-treatment window. The treated contrast is 27.4 percentage points of membership growth, 95 percent interval -1.4 to $+56.1$, clustered on 33 states, which does not clear conventional significance, against randomization p-values of 0.0136 permuting whole metropolitan areas and 0.0102 permuting whole states, which do. Quarterly Census of Employment and Wages counts of electrical-contractor employment, summed to the same statistical areas over the same years, show a treated contrast of 9.1 to 11.9 percent by 2025 against 17.3 percent in membership. Conditioning the membership estimate on employment growth in the same rows removes 7.3 to 8.9 points of a 26.6 to 27.5 point unconditional contrast on those rows, and leaves a positive residual that does not clear conventional significance on either code. The note is descriptive. It has no instrument and no runner-up control, and the employment series it benchmarks against is not itself flat before treatment: the 2016 coefficient is significant on both industry codes and the joint pre-trend test for the narrower one sits at $p = 0.053$, so part of the ratio may be catch-up rather than treatment and this note cannot separate the two. Within those limits it finds that membership in inside electrical construction grew 1.5 to 1.9 times as fast as electrical-contractor employment in the metros where hyperscale campuses were built, the lower figure on the industry code closest to the work. The two are different populations: the numerator is one branch of one trade, the denominator is every electrical contractor.

Keywords: data centers; construction labor; union membership; LM-2; QCEW; place-based investment

1. Introduction

The working paper this note supersedes reported that inside electrical locals in metros with a published hyperscale campus grew 25.8 percentage points faster between their 2018 and 2025 filings than inside locals elsewhere.[1] A referee accepted the measurement and rejected the identification. The panel opened in 2018; several treated metros had campuses under construction or in operation by then; and with one usable pre-treatment coefficient the event study could not distinguish a pre-existing divergence from a base period. The referee's verdict was that the contribution was measurement rather than causal identification.

Two things follow, and this note does both. The first is to extend the panel backward far enough to test the assumption. The Department of Labor publishes its yearly filing extracts back to 2000, and the extracts for 2010 through 2017 were retrieved in September 2026, giving sixteen years.[2] The second is to supply the external benchmark the referee identified as missing. The Bureau of Labor Statistics publishes county employment of electrical contractors from employer tax filings; the prior paper held three years of that series and compared two endpoints. This note builds all

sixteen and runs the same event study on it.[3]

Neither extension is free. Extending the panel means asserting that the classification of which locals hold inside construction jurisdiction, taken from maps the union revised in February 2018, held in 2010. Section 2.3 reports what the union’s own archived maps say about that, and what it costs when the one local whose classification cannot be carried backward is removed.

The core finding of the note is a comparison the membership series alone cannot make. Inside electrical union membership in metros with a published hyperscale campus grew faster over these sixteen years than electrical-contractor employment in the same metros, by a factor of 1.9 on the broader industry code and 1.5 on the narrower one, which is the code closest to the work a campus generates.

That is a differential in growth rates and not a density. Members per job fell in treated metros over the panel, from 0.338 to 0.298 on the broader industry code and from 0.448 to 0.372 on the narrower; it fell further in untreated metros, by 21.5 and 21.0 percent against 11.8 and 17.0. A reading on which union coverage rose in treated metros is not supported by either code. The membership count is also not a subset of the employment count, since it includes the retired, the unemployed and members working outside the industry or the area, so the ratio is a diagnostic of co-movement rather than a rate. What the comparison establishes is that the two series diverged, and in which direction. The survey-based measures the field normally uses will not support that comparison at this geography. The companion measurement paper does the arithmetic behind that claim; this note rests it on the sample sizes the database publishes and on a caution its authors have published elsewhere. Writing about state-level estimates, Hirsch, Macpherson and Vroman observed that Current Population Survey sample sizes within the smallest states do not allow precise estimates of union density, and added that the argument applies “with even greater force to density estimates for metropolitan areas and detailed industries and occupations.”[4]

1.1 The rival account

The strongest account running against a treatment reading is agglomeration. On that account a trade concentrates where it is already thick, and the concentration is self-reinforcing rather than driven by any particular project. Ellison and Glaeser stated the identification problem in their own terms: they wrote that “we cannot distinguish comparative advantage from spillover theories based only on the mean levels of geographic concentration,” and mean levels are what this note has.[5] It does not attempt the distinction.

What the agglomeration account predicts that a treatment account does not is a divergence already underway before the campuses arrive. That prediction is now testable, and section 3 reports the test. Seven coefficients for 2010 through 2016 are jointly indistinguishable from zero in the membership series; four treated metros carry a site dated before 2017, so they are not a clean pre-treatment window and section 3 says so. The employment series does not reject either, at $p = 0.305$ on the broader industry code and $p = 0.053$ on the narrower, but it is visibly less flat: treated metros sat below controls through 2016 and converged, which is the pattern an agglomeration account would expect and a treatment account would not. The two series disagree, and section 6 says what that costs.

A second rival runs through subsidy rather than agglomeration. Bartik put the annual value of United States business incentives at \$46.3 billion in 2018 dollars for 2015.[6] Data centers receive state tax exemptions in many states, though no count for this sample is established here. A subsidy

that moves construction would move the membership series without any campus-specific mechanism. This note cannot separate the two, because it has no instrument and no variation in subsidy regime that is independent of the campus siting it would explain.

A third possibility is not a rival but a confound of timing. Four metros in the campus file also host a large semiconductor or battery plant of the kind the CHIPS and Science Act and the Inflation Reduction Act, both enacted in August 2022, were written to subsidise; no award list was retrieved, so whether each plant received an award is not established here, and four is a floor rather than a census. Three of the four enter the sample under the 2010 base. Section 5 reports what removing them costs.

2. Data and classification

2.1 The filings

Labor organisations with annual receipts of \$250,000 or more file a Form LM-2 with the Department of Labor’s Office of Labor-Management Standards. Item 20 reports the number of members at the end of the reporting period, computed from Schedule 13, where the filer lists the membership categories it tracks and the number in each.[7] Agency fee payers are entered on a separate line and excluded from Item 20; work permit holders are nonmembers and fall outside it.[8] What “member” means is whatever the filing local defines it to mean in Item 69.

The department publishes each year’s filings as pipe-delimited relational tables. The `lm_data` table carries one row per filing; the `ar_membership` table carries the Schedule 13 rows.[9] The extracts for 2010 through 2017 were retrieved on 19 September 2026 and added to those for 2018 through 2025 already on deposit. Each of the eight added archives carries `lm_data` and `ar_membership` metadata files byte-identical to the 2018 versions, so the column layout is unchanged across the extension. Schedule 13 is populated throughout, with between 13,731 and 14,813 category rows per year against 14,273 in 2018.[10] Rerunning the prior paper’s script over the enlarged directory reproduced its results file byte for byte, which establishes that the additional years changed nothing about the 2018 to 2025 estimates.[11]

2.2 The sample and the base year

The sample is every IBEW local union reporting 1,000 or more members in its base-year filing and a non-blank membership figure in its 2025 filing. Unamended filings are preferred to amendments, and a blank membership field is treated as missing rather than as zero. A local is dropped if in any consecutive pair of years its membership falls below 100 from 1,000 or more, or rises more than tenfold.[12] That specification is unchanged from the working paper. Only the base year moves.

The base year moves from 2018 to 2010. Selecting on membership in 2018 conditions the sample on a year that already contains eight years of the effect being measured, which selects on the outcome. The 2018 base runs as a robustness cut beside the event study in Table 1 and every row of Table 4. The apprentice series, the employment event study in Table 3 and the conditioning exercise in section 4.3 are reported on the 2010 base only.

The two base years produce 84 and 89 inside locals respectively, and the difference is entirely a membership-floor crossing. Nine locals stood below 1,000 members in 2010 and above it by 2018; four stood above in 2010 and between 904 and 917 by 2018. No local in either difference is explained

by the movement screen or by a missing filing.[13] Two locals lack a membership value in a single year, Local 175 in 2012 and Local 1186 in 2013, and are dropped from the balanced panel, leaving 82.

2.3 Branch classification, and whether it holds in 2010

A local is classified as inside construction if its designation number appears on the IBEW’s own jurisdictional maps.[14] Those maps were retrieved in September 2026, and 50 of the 53 carry a February 2018 revision date. Carrying that classification back to 2010 is an assumption, and the Internet Archive holds the evidence that tests it.

The union published the same maps at an earlier address, and captures of that address were retrieved and parsed by the method used for the current set.[15] The archived files carry their own preparation dates, running from May 2003 to July 2007, so they are the maps the union had in force at the panel’s opening year rather than maps drawn in 2010. Across 51 United States jurisdictions the archived set carries 340 state-local pairs against 329 in the current set. Seventeen pairs differ, in ten jurisdictions. The specification reads the roster as a set of local numbers with state discarded, so the operative difference is twelve numbers. Eight of the twelve file no report of any kind in 2018 at all, and three more report between 263 and 851 members.

One differing local clears the 1,000-member floor in both base years: Local 111 in Denver, at 3,069 members in 2010 and 3,830 in 2018. It appears on the current Colorado map and not on the archived one, which carries Local 969 in that role. Local 969 is recorded as terminated on 1 July 2019, and the current Colorado map is the July 2019 revision, so the two maps are not inconsistent; each is correct for its own date.[16] Four captures of the earlier Colorado map spanning June 2006 to September 2014 are identical byte for byte and all show Local 969, so the classification of Local 111 as inside is anachronistic for at least the first five panel years. No Colorado map between September 2014 and 2019 was located.[17]

Local 111 sits in a control group, because no operator publishes a campus in Colorado. Removing it from the roster raises the treated coefficient from 27.4 to 27.8 points and the control-group-C coefficient from 40.9 to 43.5, and moves the joint pre-trend test from $p = 0.225$ to $p = 0.251$. [18] The anachronism therefore works against the result rather than for it.

A second test does not depend on the maps. Of the 89 inside locals in the 2018-base sample, **all 89 file a report in every one of the nine years from 2010 through 2018, with no gaps**, and 87 report a non-blank membership figure in every one. Among IBEW locals generally, 778 file in 2010 and 724 in 2018 on a raw read of the extracts; the released panel holds 777 for 2010, because its builder drops one malformed row that the churn check does not; restricted to locals at or above the 1,000-member threshold, one of 154 disappears and one of 156 is new, 0.6 percent in each direction.[19] The population the roster classifies is close to static.

2.4 Treatment, and the employment benchmark

A local is treated if its base-year headquarters city lies in a metropolitan area where Google, Meta, Microsoft or Amazon Web Services publishes an owned campus. The campus file holds 124 sites, of which 95 fall in the 48 metros; the crosswalk from headquarters city to metro was validated against the Office of Management and Budget’s July 2023 delineations by geocoding each local’s filed address to a county.[20] Four locals enter the sample only under the 2010 base year and were geocoded and crosswalked by the same method for this note. Three are untreated. The fourth,

Local 150 of Libertyville, Illinois, sits in the Chicago metropolitan area and is treated; it was recorded as untreated until 20 September 2026, because the crosswalk’s city-prefix match does not carry Libertyville and the county geocode that does was not consulted where the prefix match returned nothing. The front matter above records what the correction cost.[21]

Of the 84 inside locals, 28 are treated. The 56 untreated divide into 40 whose state holds a published campus somewhere, Group B, and 16 whose state neither holds a site nor is named in a treated area’s title, Group C. The label matters: a state with no campus of its own still counts as a campus state if a treated metropolitan area’s title names it, so Group C is narrower than “no campus in the state”. The division matters because a campus a hundred miles away draws on the same regional labour pool. Feyrer, Mansur and Sacerdote found that the employment effect of new oil and gas production reached its maximum at 100 miles, measured as a radius from the county centroid rather than from any well, and therefore well past the county line. Rosenthal and Strange write that labour market pooling “is likely to have effects within commuting areas, which is to say at the metropolitan level.” Summarising a literature that includes two of their own studies, they write that “agglomeration economies attenuate rapidly,” and the work they summarise measures that attenuation in miles rather than at the scale of a state.[22] Group B is therefore a weak analogue of an attenuation band rather than a clean control, and its intermediate position between A and C is what attenuation predicts rather than evidence against the design. Group C is the control.

The design follows Greenstone, Hornbeck and Moretti, who compared counties that won a large new plant against the runner-up counties that lost it. They found the spillover concentrated in plants sharing the new plant’s own industry, a significant 17 percent for same-industry plants against an insignificant 3.3 percent for plants in other industries in their first model; in their second the same contrast weakens to a poorly determined 33 percent against a marginally significant 8.9 percent.[23] The operators published no site-selection finalists on any page examined for this note or for the prior paper, so no runner-up list was assembled and that design cannot be run in its original form. The substitute the prior paper used, a comparison of inside locals against IBEW locals in the union’s other branches in the same metros, is reported there and is not relied on here, because it does not survive state fixed effects.

The employment benchmark is the Quarterly Census of Employment and Wages, which counts employment from employer tax filings rather than from a survey. Two industry codes are used: 23821, electrical contractors and other wiring installation contractors, and 238212, nonresidential electrical contractors. The prior paper held three years. This note holds all sixteen, assembled from two routes, because the Bureau’s Open Data Access industry slices returned HTTP 404 for every year probed from 1990 to 2013, seven years in all, so this note takes 2014 as that endpoint’s first usable year without asserting that no earlier slice exists.[24] A metropolitan area’s employment is summed over the counties disclosed in every year of the panel, so a county suppressed in one year is dropped from all years for that area and disclosure cannot manufacture growth.

3. The pre-trend

Table 1 reports an event study on the balanced panel of 82 inside locals: log membership, local fixed effects, year fixed effects, treatment interacted with every year, omitted year 2017, standard errors clustered on 32 states. The omitted year is the last before the panel’s main treatment wave, so every coefficient for 2010 through 2016 is a test of whether the two groups were already diverging.

It is not a clean parallel-trends test, because no treatment begins in 2017. Four of the treated metros in the sample carry a dated site before that year — Des Moines in 2013, Columbus and Dallas-Fort Worth in 2015, Albuquerque in 2016 — so for the locals in those metros some of the years being read as pre-treatment are post-treatment. The omitted year is a normalisation chosen as the last before the main wave. A flat gap through 2016 is evidence that the two groups were not already diverging, which is worth having, and not evidence that a design assumption holds.

Table 1. Event study, inside locals, treated against untreated, percent

Year	Base 2010	Base 2018
2010	+1.22	+4.68
2011	-0.60	+0.27
2012	-0.87	-0.27
2013	-1.77	-0.93
2014	+2.69	+3.15
2015	+1.19	+1.10
2016	-0.45	-0.07
2017	omitted	omitted
2018	+1.17	+1.71
2019	+1.89	+2.91
2020	+2.42	+3.58
2021	+4.07	+5.09
2022	+6.87	+7.33
2023	+10.96 (p 0.082)	+11.07 (p 0.055)
2024	+13.84 (p 0.059)	+14.78 (p 0.028)
2025	+17.29 (p 0.045)	+18.89 (p 0.015)
Joint test, 2010 to 2016	chi2(7) = 9.40, p = 0.225	chi2(7) = 7.68, p = 0.362

No pre-period coefficient is significant at conventional levels in either specification; the smallest pre-period p-value is 0.159. The post-period path rises monotonically from 2018 and clears conventional significance only at the end of it: in 2025 alone on the 2010 base, $p = 0.045$, and from 2024 on the 2018 base. The 2023 coefficient is $p = 0.082$ and $p = 0.055$ respectively, and clears on neither.

A joint test that fails to reject is not evidence that trends are parallel. It is the absence of evidence that they are not, on 82 locals in 32 clusters, and the pre-period intervals are wide: the 2011 interval runs from -9.0 to $+7.8$ percent. The claim the extended panel supports is narrower than it looks. The design was put to a test it could have failed, on seven coefficients rather than one, and it did not fail.

4. Results

4.1 Group medians and regression

Table 2 reports median membership growth from 2010 to 2025 for the 84 inside locals, split by whether an operator publishes a campus in the local’s own metro, elsewhere in its state, or nowhere in its state. Medians are used because the filing does not separate a merger from organic growth.

Table 2. Median growth, inside locals, 2010 to 2025, percent

Group	n	Median
A, campus published in the metro	28	+42.13
B, campus in the state, metro untreated	40	+24.45
C, state neither holds a site nor is named in a treated area's title	16	+3.96
B and C combined	56	+18.92

A minus C is 38.2 points; A minus B is 17.7; A minus the combined comparison is 23.2. In aggregate membership, treated locals grew from 81,641 to 109,417, or 34.0 percent, and untreated locals from 149,027 to 169,406, or 13.7 percent.[25]

The 28 treated locals do not move together. Six grew by more than 100 percent, led by Local 640 in Phoenix at 148.7 and Local 60 in San Antonio at 142.6. Six contracted, four of them Chicago-area locals, of which Local 701 in Warrenville fell 23.2 percent; a fifth Chicago-metro local, Local 697 in Merrillville, grew 15.3 percent. Local 134 in Chicago is the largest treated local in every year of the panel and is one of the six that contracted, from 14,254 members to 12,161. Local 26, which testified in Loudoun County, grew 56.3 percent from 7,727 to 12,078 and ranks thirteenth of 28. The treated locals sit in 23 distinct metropolitan areas.

Regression on the same population gives a treated coefficient of **27.35 percentage points, 95 percent interval -1.35 to $+56.05$, $p = 0.062$** , clustered on 33 states.[26] Restricting the comparison to Group C raises it to 40.89, interval 12.38 to 69.39, $p = 0.005$, and gives Group B a coefficient of 18.95 in its own right, which is what an attenuation band rather than a control should look like. On the 2018 base the same two estimates are 25.83 and 35.47.

The headline coefficient does not clear conventional significance and both randomization tests do. The clustered regression rests on 33 state clusters, where a cluster-robust standard error is least reliable. Randomization inference clears at both permutation units, the contrast against Group C clears at $p = 0.005$, and the leave-one-out range, which carries no p-value, never crosses zero. The estimate is stable in sign and magnitude and fragile in significance, and this note reports it that way rather than choosing the route that reads best.

4.2 Apprentices

Schedule 13 itemises the membership categories a local tracks, and some inside locals track apprentices separately. That gives a second series inside the same filings, and it speaks to composition rather than level: a local absorbing a neighbouring local gains journeymen, while a local staffing a build-out expands its training pipeline.

Among treated locals reporting an apprentice category in both 2010 and 2025, the aggregate count rose from 5,417 to 10,828, or 99.9 percent. Among untreated inside locals the comparable figures are 4,491 to 6,249, or 39.1 percent. Median change per local is +124.5 percent in the treated group against +85.8 percent elsewhere.[27] Local 26 alone went from 2,033 apprentices to 3,471.

The bound on this is heavy and the series should not carry weight on its own. Only 12 of 28 treated locals and 13 of 56 untreated locals report an apprentice category in both endpoint years. A local reporting none is excluded rather than zeroed, because Schedule 13 categories are free text and a local may fold apprentices into another category. The aggregate is also not a per-local average: Local 26 supplies 37.5 percent of the treated total in 2010 and 32.1 percent in 2025.

Two extraction rules were wrong when this series was first built and are recorded here because they moved the corrected figures materially. The magnitude of the worst case is not reported by any deposited script and is not quoted here; the single negation case arises only under the 2018 base. Locals reporting more than one apprentice category had only one of them counted, and a category named “all active members with exception of apprentices” was counted as apprentices. Of the fifteen treated series reporting an apprentice category, seven were affected by the first rule; the second rule rejected no category on this base, and the one case it catches enters only under the 2018 base.[28]

4.3 The employment benchmark

Two constructions of this correlation appear across the two deposits and they are not the same quantity. The figures below sum employment for each local in the estimation sample, so a metro holding several locals contributes its employment several times. The companion measurement paper reports 0.553 on 77 and 80 locals, counting each metro’s employment once across the locals with a complete series. Neither is wrong and they should not be quoted interchangeably. On this construction, membership growth and QCEW employment growth in the same metro over the same fifteen-year growth span correlate at 0.538 for the broader industry code and 0.549 for the narrower, across 81 and 83 locals.[29] Over an eight-year window the prior paper reported 0.58. The correlation is the measurement result: a count assembled from union filings tracks an independent count assembled from employer tax filings.

Running the event study of Table 1 on employment rather than membership gives the comparison in Table 3. The unit is the metropolitan area rather than the local, because five sample locals sit in Chicago and two in Minneapolis and there is one employment figure for each metro.

Table 3. Event study on membership and on employment, percent. The membership column comes from `lm2_identification_v3.py`; both employment columns come from `lm2_qcew_event.py`, which is not the script behind section 4.3’s correlations and conditioning.

Year	Membership, 82 locals	QCEW 23821, 63 metros	QCEW 238212, 66 metros
2016	-0.45	-3.64 (p 0.008)	-4.73 (p 0.001)
2018	+1.17	-0.11	+0.67
2020	+2.42	-0.84	+1.30
2022	+6.87	+1.44	+3.94
2023	+10.96 (p 0.082)	+4.45	+6.98 (p 0.080)
2024	+13.84 (p 0.059)	+4.01	+6.37
2025	+17.29 (p 0.045)	+9.10 (p 0.061)	+11.92 (p 0.022)
Joint test, 2010 to 2016	chi2(7) = 9.40, p = 0.225	chi2(7) = 8.33, p = 0.305	chi2(7) = 13.92, p = 0.053

Membership grows 1.9 times as fast as employment against the broader code and 1.5 times as fast

against the narrower one. Read directly, membership per electrical-contractor job fell more slowly in treated metros than in untreated ones. It fell in both, so this is a differential and not a rise in coverage.

Conditioning membership growth on employment growth within the same rows tests how much of the membership result is a labour-market result. The treated coefficient falls from 26.64 to 19.31 against the broader code, $p = 0.084$, and from 27.47 to 18.53 against the narrower, $p = 0.159$.^[30] On this base, between a quarter and a third of the membership effect runs through employment in the trade generally. The residual is positive on both industry codes. Neither conditional coefficient clears $p = 0.05$, at 0.084 and 0.159, and the two unconditional coefficients on the same rows sit at 0.067 and 0.060. The exercise establishes the sign of the residual and not its significance.

5. Robustness, and what does not survive

A staggered-adoption estimator is reported here rather than in the results, because the dating will not carry it. The referee objected that treatment is undated, with pre-2018 campuses and 2025 announcements pooled, and named Callaway and Sant’Anna among the literature the prior paper had not engaged. It can be run on a subset: 21 of the 28 treated locals in the balanced panel sit in a metro carrying an announcement year usable as a cohort, across ten cohorts from 2013 to 2025, and the estimator returns a cohort-weighted aggregate over post-treatment periods of +13.8 percent.^[31] Across five passes — four base periods against a never-treated comparison and one against a not-yet-treated comparison — that aggregate ranges from +11.02 to +14.09 on this base and from +9.98 to +13.26 on the 2018 base.^[32]

It is not reported as a finding, for the reason the companion measurement paper gives for not reporting it at all. Seventy-eight of the 124 sites in the campus file carry no announcement year, so a cohort structure built on the 46 that do will not bear a claim about timing. Seven treated locals drop out of the estimator, six of them in metros with no dated site; one of those is Northern Virginia, which holds eight campuses in the file and not one date, so Local 26 contributes nothing here. Two pre-period cells exclude zero at the standard base period, one of them resting on two locals, and the other does not survive a change of base period. The number is in the deposited log and in this paragraph. Nothing in this note’s conclusions rests on it.

Table 4 reports the checks the prior referee requested, on both base years.

Table 4. Robustness of the treated coefficient, inside locals, percent

Specification	Base 2010	Base 2018
Inside only, treated	+27.35 (p 0.062)	+25.83 (p 0.005)
With state fixed effects, treated $\times$ inside	+21.56 (p 0.332)	+15.42 (p 0.370)
A minus B, Wald test	+21.94 (p 0.140)	+22.02 (p 0.016)
Drop metros hosting a CHIPS- or IRA-era fab or battery plant	+18.37 (p 0.209)	+17.11 (p 0.048)
Drop metros whose only sites are dated 2025 or later	+31.02 (p 0.055)	+28.74 (p 0.004)

Specification	Base 2010	Base 2018
Weighted by base-year membership	+20.35 (p 0.241)	+22.80 (p 0.009)
Leave-one-out range	+23.66 to +30.03	+21.24 to +27.68
Randomization inference, permuting metropolitan areas	p 0.0136	p 0.0022
Randomization inference, permuting states	p 0.0102	p 0.0044

The sixteen-year window is the weaker of the two on conventional significance, and the note states that plainly. Five specifications clear $p = 0.05$ on the eight-year window and do not clear it on the sixteen-year one: the inside-only estimate itself, the A-minus-B contrast, the removal of the semiconductor-and-battery metros (four in the campus file, of which three enter this sample), the removal of late-announcement metros, and weighting by base-year membership. State fixed effects absorb the interaction into insignificance on both windows, and that row changes the estimand as well as the specification, so it does not measure the cost of the fixed effects alone.

Two checks clear $p = 0.05$ on both windows, and both are randomization tests: permuting whole metropolitan areas, and permuting whole states. The contrast against Group C clears on both as well, at $p = 0.005$ on the sixteen-year window, and is a different comparison group rather than a different specification. The leave-one-out range carries no p-value and never crosses zero on either window.

What the longer window buys is the pre-trend test, and what it costs is precision. The point estimate is close on the two windows, 27.35 against 25.83, and the sixteen-year standard error is larger: 14.64 against 9.18. Nothing deposited here identifies why, and this note does not assert a reason. A reader who wants the tighter intervals should use the eight-year panel and accept that its parallel-trends assumption is untested. That trade is the note’s substantive contribution and it runs against the note’s own headline.

The list of five is a statement about significance crossings and not about magnitudes. Several coefficients are larger on the longer window, among them the state-fixed-effects interaction, the late-announcement cut and both ends of the leave-one-out range, so the crossings reflect the wider standard errors on the longer window as well as any change in the estimates themselves.

The fab and battery cut is the one that matters most, because the confound is real and the project list that would bound it has not been retrieved. Four metros in the campus file host such a plant: Phoenix, Columbus, Boise and Kansas City. Three of the four enter the 2010-base sample, because Boise’s local stood below the membership floor in 2010, so that cut removes three treated locals on the longer window and four on the shorter one. Four is a floor rather than a census. Removing them costs nine points and conventional significance on the longer window.

What would bound the confound is a project list: the set of semiconductor and battery facilities receiving federal support, with locations and announcement dates. No such list was retrieved for this note. With one, the four metros become a measured set rather than a floor, the cut becomes a stated sample restriction rather than a sensitivity check, and the overlap between subsidised fabrication and hyperscale siting becomes estimable in its own right. That retrieval is the single

change that would do most for the design, and it is a retrieval problem rather than an econometric one.

6. Limitations

Membership is not employment. Item 20 counts members, not hours or jobs. A local’s roll moves with organising, with jurisdictional transfers and with the local’s own classification choices. Section 4.3 benchmarks the series against a count that does measure employment, and the two diverge by a factor the note reports rather than explains.

The employment series is not cleanly flat before treatment. The 2016 coefficient is significant in both industry codes and the 2015 coefficient is also significant on the narrower one, at -4.42 , $p = 0.043$, with 2014 at $p = 0.062$, so the pattern runs over more than a single year. The joint pre-trend test for the narrower code sits at $p = 0.053$. Dropping each treated metro in turn leaves that test between $p = 0.013$ and $p = 0.083$ and the 2016 coefficient between -4.46 and -4.97 , so no single metro carries it.[33] The sign means treated metros sat below controls before 2017 and converged. Part of the membership-to-employment ratio in section 4.3 may therefore be catch-up rather than treatment, and the note cannot separate the two.

The mechanism is bounded to one trade and one branch of it. The note measures inside electrical construction and nothing else. It does not measure the other trades on a data center site. Virginia’s Joint Legislative Audit and Review Commission lists electricians alongside pipefitters and HVAC technicians among the occupations on a data center construction site.[34] It does not measure non-union electrical construction, which the filings cannot see at all and whose share of the trade in the metros this note covers is not established here. A membership gain in an inside local is consistent with the trade growing, with union share rising within a fixed trade, or with both, and the employment benchmark in section 4.3 separates those only in aggregate.

There is no instrument and no runner-up control. No site-selection finalists were found on any operator page examined, so the Greenstone, Hornbeck and Moretti design cannot be run in its original form.[35] The within-union placebo the prior paper used does not survive state fixed effects and is not relied on here.

The control group is concentrated. Group C is seven states — Alaska, Colorado, Florida, Hawaii, Massachusetts, New York and Vermont — and twelve of its sixteen locals sit in three of them, from `lm2_prose_figures.py`, block [6c]. It is also defined by the absence of a published campus in a file this note records as incomplete, so membership of the group is a statement about what the operators have published rather than about where campuses are.

Treatment is assigned from a headquarters address, not from a jurisdiction. A local is placed in a metropolitan area by the address on its filing, and an IBEW inside jurisdiction is not coterminous with a statistical area. No machine-readable jurisdiction boundary file was found on the union’s own site, which is the only source searched, so nothing here corrects for the mismatch.

Most published sites carry no date. Seventy-eight of the 124 sites in the campus file have no announcement year, and twenty-two of the twenty-eight treated locals sit in a metro carrying at least one dated site; twenty-one sit in a metro whose earliest date falls inside the panel, the exception being Indianapolis, dated 2026. Both counts come from `lm2_prose_figures.py`, block [6b]. The annual path in section 3 is therefore normalised on a convention rather than on an event,

and a flat pre-2017 gap is evidence that the two groups were not already diverging rather than evidence that a design assumption holds.

Travellers place part of the treatment in the comparison group. A member working on a permit in another local's jurisdiction is counted by the home local and not by the local whose territory holds the work, and a large project is exactly the occasion for that movement. A campus in a treated metro drawing electricians from locals elsewhere raises membership, if it raises it at all, on the rolls of the untreated locals supplying them, which would bias the contrast toward zero to the extent it operates. No field in any released panel row records the jurisdiction worked as against the member's home local, and none of the LM-2 sections examined for this note asks for one, so the magnitude is not recoverable here.

Two Schedule 13 anomalies are on the record. Local 124 reports 44 apprentices in fiscal 2021 between 447 and 563 in the adjacent years, as filed and never amended, with no other category absorbing the fall. Local 98's fiscal 2020 apprentice count falls by 460 while its two member categories rise by 429, which reads as a reclassification and leaves its total nearly unchanged.[36] Local 124's case is not a reclassification: its reported total falls by 391 in the same year, so that anomaly moves a membership figure the event study uses. Local 124 is one of the 82 locals in the balanced panel.

The classification of Local 111 is anachronistic for at least five panel years, as section 2.3 sets out, and removing it moves the treated coefficient by half a point in the direction that strengthens the result. The headline figure reported here is the one that keeps Local 111.

7. Conclusion

The working paper this note supersedes could not test the assumption its event study rested on. The panel now runs sixteen years and the seven coefficients for 2010 through 2016 are jointly indistinguishable from zero, subject to the four early-dated metros section 3 names. The treated contrast is close on the two windows, 27.4 against 25.8, and is measured less precisely on the longer one: it does not clear conventional significance in the clustered regression, where the eight-year estimate did.

The finding the longer panel makes available is not the treatment coefficient. It is the comparison between two independently assembled counts of the same trade in the same places over the same years. Union membership in inside electrical construction grew 1.5 to 1.9 times as fast as electrical-contractor employment in metros where hyperscale operators built campuses, the lower figure on the narrower industry code. Conditioning on employment removes between a quarter and a third of the membership effect and leaves a positive residual.

What that residual represents is not settled here. It is consistent with the union holding a larger share of a growing trade where the work concentrated, though members per job fell in treated metros as well as elsewhere and no level reading supports a rise. It is also consistent with treated metros converging on controls from below, which the employment series shows and the membership series does not. Separating them requires a source that measures hours or jobs inside a union local's own jurisdiction, and none was found among the agency series examined for this note. The Bureau of Labor Statistics publishes a metropolitan occupational series for electricians, which is organised by statistical area rather than by jurisdiction and so does not close the gap.

Endnotes

- [1] Turner, “Where the electricians went,” working paper v0.1, 5 September 2026, § 4.3. DOI 10.5281/zenodo.22310706.
- [2] OLMS yearly data download, retrieved 19 September 2026. The servlet’s report codes rotate; three distinct code stems were observed across fourteen days, so the codes are not a durable citation and the page must be re-read. See `LM2_Extension_Gate_Log_2026-09-05.md`.
- [3] Referee adjudication, 5 September 2026. The referee’s point reads “QCEW 238210 / OEWS 47-2111 benchmark; both series exist; not retrieved.”
- [4] Hirsch, Macpherson and Vroman, “State-Level Estimates of Union Density, 1964 to Present,” author manuscript, footnote 2. The condensed published version, “Estimates of Union Density by State,” *Monthly Labor Review*, July 2001, does not carry the footnote.
- [5] Ellison and Glaeser, “Geographic Concentration in U.S. Manufacturing Industries,” NBER working paper 4840, § 2.2.2, in the discussion following Proposition 4. They add in a footnote that they “do not feel that attempts to distinguish the theories on such grounds will be fruitful.”
- [6] Bartik, “Should Place-Based Jobs Policies Be Used to Help Distressed Communities?,” Upjohn Institute working paper 19-308, August 2019, doi 10.17848/wp19-308. The retrieved text carries no PDF pagination, so the figure is cited by section rather than by page. The figure is stated in 2018 dollars for 2015. Only the working paper is held on disk; the published *Journal of Economic Perspectives* version was not retrieved, so nothing here compares the two.
- [7] Form LM-2 instructions, Item 20 and Schedule 13.
- [8] Form LM-2 instructions, Schedule 13 and Item 21.
- [9] OLMS yearly data extracts, `lm_data` and `ar_membership` tables.
- [10] `LM2_Extension_Gate_Log_2026-09-05.md`, Gate 1.
- [11] `lm2_v2_1.py` rerun with the enlarged extract directory reproduces `lm2_v2_1_results.json` at md5 dfe44784f47b0b8ec908a2202edcf730.
- [12] `lm2_v3_0.py`. The crosswalk object is read out of the prior paper’s script by abstract syntax tree at run time rather than retyped, so the two cannot drift.
- [13] `lm2_base_year_diagnostics.py`, pass 2.
- [14] IBEW Inside jurisdictional maps. A label is a text item that is numeric and set at 16 points or more; the size gap between highway shields and local labels is clean in both map vintages.
- [15] `lm2_ibew_roster_archived.py`; `IBEW_Inside_Locals_Roster_archived.csv`, 340 rows.
- [16] OLMS records Local 969 with a termination date of 1 July 2019 and its terminate flag set in the 2019 filing; the local’s last report covers 2019 with 274 members.
- [17] `ibew_src/C0_Inside_map_captures_2026-09-19.txt`. Four captures, 3 June 2006 to 9 September 2014, 548,932 bytes and one SHA-256. Three Archive queries are recorded there, one of which returned a server error and therefore establishes nothing.

- [18] Run with `LM2_ROSTER_DROP=111`; output saved at `lm2_roster_drop111_base2010.log`.
- [19] `lm2_churn_2010_2018.py`.
- [20] `lm2_omb_crosswalk.py`; `Hyperscaler_Campus_Locations.csv`.
- [21] `lm2_omb_crosswalk_extend.py`. Locals 300, 150, 816 and 82, geocoded by the Census Bureau Geocoder Geographies API from each local's 2010 filed address.
- [22] Feyrer, Mansur and Sacerdote, NBER working paper 21624, § 7: "The employment effects reach a maximum at 100 miles." Rosenthal and Strange, author manuscript April 2020, § I and § III.
- [23] Greenstone, Hornbeck and Moretti, "Identifying Agglomeration Spillovers," NBER working paper 13833, pp. 31 and 32. They also report that a one standard deviation increase in worker transitions between the incumbent and new plant industries is associated with a larger spillover; the derived files in this project record that magnitude as 7 percent in one place and 7 percentage points in another, and the source text held on disk is truncated before its tables, so the figure is not restated here.
- [24] Probed 19 September 2026: 1990, 2000, 2005, 2007, 2010, 2012 and 2013 all return HTTP 404 at `data.bls.gov/cew/data/api/{year}/a/industry/`; 2014 returns 3,453 rows. Years before 2014 were taken from the annual singlefile bulk archives. Provenance is in the header of each file under `qcew_src/`.
- [25] `lm2_referee_checks_v3.py`, check 6.
- [26] `lm2_identification_v3.py`, specifications R3 and R4.
- [27] `lm2_apprentice_groups.py`.
- [28] `lm2_apprentice_flag_check.py` opened the filings that exposed both rules as wrong; `lm2_v3_0_tables67.py` carries the corrected rules and prints the multi-category locals and every rejected category. The seven-of-fifteen count is on the 2010 base; the single negation case is Local 1579, which is a treated local only under the 2018 base.
- [29] `lm2_qcew_benchmark_v3.py`.
- [30] Same script. The unconditional coefficient is estimated on the same rows as the conditional one, so the comparison is within-sample.
- [31] Cohort is the earliest announcement year among a metro's sites in `Hyperscaler_Campus_Locations.csv`. The estimator follows Callaway and Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics* 225 (2021). **That paper is not held on disk.** The estimator is implemented from its published definition as given in the secondary literature this note holds, and the citation has not been checked against the article itself. Treat it as an unverified citation.
- [32] `lm2_staggered_diagnostics.py`, passes C and D.
- [33] `lm2_qcew_pretrend_diag.py`. Leave-one-out over all 23 treated metropolitan areas.
- [34] Virginia Joint Legislative Audit and Review Commission, *Data Centers in Virginia*, Report 598, 9 December 2024, held on disk. Listing the occupations on a data center construction site, it names "electricians installing cabling, equipment, and generators" and "pipefitters and HVAC technicians". Verified against the report rather than through the prior working paper.

[35] Greenstone, Hornbeck and Moretti, “Identifying Agglomeration Spillovers,” compared counties that won a large new plant against runner-up counties that lost it.

[36] `olms_src/failed_reports_sched13_2026-09-19.txt`. Both reports were retrieved from the Online Public Disclosure Room and read against the extracts.

Provenance and open items

Three items are open and none is closed by this note. No Colorado inside map between September 2014 and 2019 has been located, so the year Local 111 took the jurisdiction is not established. The Callaway and Sant’Anna citation at the endnote in section 5 rests on the estimator’s published definition and not on a copy held on disk, and is therefore unverified against the article. And the project list that would bound the semiconductor and battery confound in section 5 has not been retrieved, so the four metros named there are a floor rather than a census.