

Where the electricians went: Hyperscale data center construction and building trades membership in federal union filings, 2018 to 2025

Working paper, version 0.1. 5 September 2026.

Every figure in the tables is produced by the scripts named in the endnotes, from the sources deposited with it; the counts and classification history that are the author's own are labelled as such in their notes. Section 3 sets out the three classification variables and what in each is judgement; the closing note on provenance says what the apparatus does not cover.

Author note

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

Deposit. The data, the code and the retrieved sources are deposited with this paper, except where a source is copyrighted and cannot be redistributed; the replication package and its run order are described in its README. Version DOI: [10.5281/zenodo.22310706](https://doi.org/10.5281/zenodo.22310706). Concept DOI: [10.5281/zenodo.22310705](https://doi.org/10.5281/zenodo.22310705). Both are reserved and resolve on publication.

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 paper 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 for data extraction, code and retrieval. The author is responsible for the content.

Abstract

Published employment figures for data center construction come from the industry, from agencies relaying an industry figure, or from unions testifying on projects; this paper could locate no independent series at the metropolitan level. It builds one from federal filings. Labor organisations above a receipts threshold report their membership annually to the Department of Labor under penalty of perjury, and the International Brotherhood of Electrical Workers publishes maps of which locals hold inside construction jurisdiction. Joining the two gives an eight-year panel of inside electrical construction membership by metropolitan area for 89 locals. Treatment is a hyperscale campus published by one of four operators in the local's metropolitan area. The median treated local grew 33.7 percent between its 2018 and 2025 filings against 13.2 percent elsewhere, and the regression difference is 25.8 percentage points, 95 percent interval 7.8 to 43.8, clustered on 36 states, with a randomization p-value below 0.001. An event study with local and year fixed effects finds no differential trend in the two pre-period years the panel can see and a divergence significant from 2023, reaching 17.8 log points by 2025. IBEW locals in the same metros that are not on the Inside maps diverged from their own comparison group by less than half as much on medians, and the regression form of that comparison does not survive state fixed effects. The effect is concentrated: removing the five fastest-growing locals leaves 9.0 points and no significance, and removing the

four treated metros known to host a federally subsidised semiconductor or battery plant leaves 17.1. Seventeen of the twenty-nine treated locals sit in metros the industry’s capacity report does not list in its H2 2025 edition. County employment of nonresidential electrical contractors, from employer tax filings, correlates with the membership series at 0.58 across 86 locals and shows the same treatment contrast independently, 19.6 points, $p = 0.003$. The paper is descriptive. It has no instrument and no runner-up control, membership is not employment, and within those limits it finds the trade grew where the campuses were built.

Keywords: data centers; construction labor; union membership; LM-2; place-based investment; agglomeration; electrical workers

1. Introduction

A hyperscale data center is a concrete shell with an unusual bill of materials. It needs little of what most commercial buildings need, which is finish carpentry, paint and interior fit-out, and a great deal of what most do not, which is medium-voltage switchgear, generator plant, chilled-water piping and air handling at industrial scale. Virginia’s legislative auditors, listing the occupations on a site, name electricians “installing cabling, equipment and generators” alongside pipefitters and HVAC technicians.[1] If the build-out since 2018 has moved local construction labor markets, the trace should appear in the electrical trade, in the metros where the largest facilities went up.

Testing that has not been done with the published data this paper could locate. The employment figures it located for data center construction come from three kinds of source, and none of them is independent. Industry associations produce economic impact studies; the Northern Virginia Technology Council’s 2024 report, prepared by Mangum Economics, puts direct construction employment in Virginia at 14,240 jobs.[2] State audit agencies carry industry self-reports; Virginia’s Joint Legislative Audit and Review Commission, in its 2024 study, states that “approximately 1,500 workers are on site” at the height of construction, and attributes the figure in the same sentence to “data center representatives.”[1] And union officials testify at public hearings; the political coordinator of IBEW Local 26 told a Loudoun County town hall in July 2026 that his local had gone from 9,000 members to 17,000 over the same eight years.[3] Each of these may be accurate. None can be checked against the others within the record they appear in. Section 7 checks them against a federal employment series at the same geography and industry, which no source located for this paper had brought to the question.

This paper builds one. Every labor organisation subject to the Labor-Management Reporting and Disclosure Act, the Civil Service Reform Act or the Foreign Service Act with annual receipts of \$250,000 or more, the threshold in force for every fiscal year in this paper, files a Form LM-2 with the Department of Labor’s Office of Labor-Management Standards, and Item 20 of that form reports the number of members at the end of the reporting period, with the categories of membership itemised in Schedule 13.[4] The filing is signed by two officers under penalty of perjury.[5] The department publishes each year’s filings as a relational dataset; the reporting history held for one local runs back to 2000.[6] Separately, the IBEW publishes a jurisdictional map for every state showing which of its locals holds the inside construction jurisdiction for each area.[7] Joining the two yields a count of inside electrical construction membership by local, and by extension by metropolitan area, for every year of the build-out.

The measure has limits that are stated now and repeated where they bind. Membership is not

employment. A local’s roll moves with organising, with jurisdictional transfers and with the local’s own classification choices, and no item on the form asks for hours.[5] The CPS-based estimates that the field uses for union density are imprecise at this geography, for a reason the authors of the standard database stated themselves, in the manuscript of their methods note: sample sizes that will not support state-level estimates apply “with even greater force to density estimates for metropolitan areas and detailed industries and occupations.”[8] Union-reported membership is the older source, displaced by the CPS in the 1970s, and where the two overlap they agree closely.[9] This paper returns to the older source at a finer geography than the directories built from it published.

The design follows Greenstone, Hornbeck and Moretti, who compared counties that won a large new plant against the runner-up counties that lost it.[10] The operators publish no site-selection finalists on any page examined for this paper, so no runner-up list is available. The substitute is a placebo inside the same metros. The IBEW allocates construction jurisdiction by area, and publishes a separate set of maps headed “Inside Jurisdiction” for the wiring of buildings.[7] Locals in the union’s other fields, utilities, telecommunications, manufacturing, railroads, broadcasting, government and construction outside the Inside maps, draw on the same regional labor market but do not hold that jurisdiction. If the divergence appears in Inside locals and is smaller in the others, the explanation has to run through construction, though not necessarily through data center construction alone, a point section 10 takes up. Greenstone and coauthors found the same structure in their own data: productivity spillovers of 17 percent for incumbent plants in the new plant’s own two-digit industry against a statistically insignificant 3.3 percent for plants in other industries, and a one standard deviation increase in worker flows between the industries associated with a 7 point larger spillover.[11]

The result is that Inside locals in metros where a hyperscale operator published an owned campus grew 25.8 percentage points faster over 2018 to 2025 than Inside locals elsewhere, with a 95 percent interval of 7.8 to 43.8 and standard errors clustered on state. An event study with local and year fixed effects shows no differential trend in the two pre-period years the panel can see and a significant divergence from 2023. The effect survives restriction to December fiscal-year filers, removal of Chicago, removal of Phoenix, weighting by size, and estimation at the median. It does not survive removal of the five fastest-growing locals, which cuts it to 9.0 points. Four treated metros known to the author, two of them among those five, also host a semiconductor or battery plant subsidised by federal legislation of August 2022; no project list was retrieved, so four is a floor. Removing them leaves 17.1 points at $p = 0.048$. Federal employment counts for nonresidential electrical contractors, summed to the same statistical areas, grew 19.6 points faster in treated metros, $p = 0.003$, and correlate with the membership series at 0.58.

Four things the paper does not establish are stated here so that no later sentence can be read as claiming them. It does not establish that data center construction caused the membership growth; there is no instrument and no runner-up control, and the within-metro placebo is suggestive rather than decisive. It does not establish employment or hours; the filing reports neither, and section 7 benchmarks it against the federal employment series that does. It does not establish that every treated metro grew, and section 5 shows that most of the average effect comes from a minority of them. And it does not separate a demand shock to one trade from a self-reinforcing concentration of that trade’s workers. Ellison and Glaeser showed that natural advantage and spillovers produce the same observed concentration and “cannot” be distinguished “based only on the mean levels of geographic concentration,”[12] and mean levels are what this paper has.

2. The filings and what they measure

2.1 Form LM-2

The Labor-Management Reporting and Disclosure Act of 1959 requires labor organisations to file annual financial reports with the Secretary of Labor.[13] Organisations with total annual receipts of \$250,000 or more file Form LM-2, the form required of the largest filers; a rule effective 1 July 2026 raises the threshold to \$350,000 for fiscal years beginning on or after that date, which touches no filing in this paper.[14] The form carries a statement of assets, liabilities, receipts and disbursements, a schedule of officers and employees, itemised disbursements above a threshold, and, at Item 20, “How many members did the labor organization have at the end of the reporting period?”[15] Two officers sign a declaration “under penalty of perjury and other applicable penalties of law” that the report “is, to the best of the undersigned individual’s knowledge and belief, true, correct and complete.”[5]

Item 20 is computed from Schedule 13, Membership Status Information, where the filer lists “the categories of membership tracked by the reporting labor organization” in Column A and the number in each category in Column B.[16] The instructions require each category to be defined in Item 69, Additional Information, with a statement of whether its members pay full dues.[17] The software sums Column B into Item 20. Agency fee payers, who pay in lieu of dues without joining, are entered on a separate line and excluded from the total: “This total in Column (B) is not the total number of members of the labor organization reportable in Item 20.”[18] Work permit holders are defined in Item 21 as nonmembers working within the local’s jurisdiction, and Schedule 13 counts members, so they fall outside Item 20.[19]

Three consequences follow for the measure. Apprentices are counted if the local tracks them as a membership category, and in the case examined in section 9 they are. Travelers on permits are not counted. And what “member” means is whatever the filing local says it means in Item 69.

2.2 The yearly extract

OLMS publishes each year’s filings as a set of pipe-delimited relational tables.[20] The `lm_data` table carries one row per filing, on any of the LM-2, LM-3, LM-4 and LM-5 forms, with the union’s affiliation, designation, headquarters city and state, fiscal period, total membership, total receipts and an amendment flag. The `ar_membership` table carries the Schedule 13 rows, joined on report identifier. The extracts for 2018 through 2025 were retrieved in September 2026 and are the only data source for the membership figures in sections 4 through 8; section 9 adds one local’s filed report and its reporting history.[21]

The extracts carry known defects, and the specification in section 3.1 is built around them. Some locals file amended reports, and a local-year may have only an amendment on file. Some filings carry a blank membership field. Some carry keying errors, visible as a roll that collapses or multiplies between adjacent years and returns. Category labels in Schedule 13 vary in spelling across years; one local’s 2023 filing lists “APPRETICES AND TRAINEES,” and an exact-string join on that field returns silent zeros for the year.[22] The 2025 extract holds 19,080 filings against 21,277 in 2018 across all forms, and 4,607 against 4,530 on the LM-2 alone, so the decline is in the smaller forms; how much of it is attrition and how much late filing is undetermined.[23]

2.3 Why not the CPS

The standard measure of union membership in the United States is survey-based. The Current Population Survey asks employed wage and salary workers whether they are members of a labor union, and the database that Hirsch and Macpherson maintain aggregates the answers to national, state, industry and occupation estimates.[24]

It is imprecise here, and the reason is in the authors’ own account of the state database, in a footnote of their manuscript that the published summary omits. Because CPS sample sizes within the smallest states will not support precise density estimates, the Bureau of Labor Statistics has been “understandably reluctant to publish official estimates of state union density, an argument that applies with even greater force to density estimates for metropolitan areas and detailed industries and occupations.”[25] A count of inside electricians in the Phoenix metropolitan area in a given year is the kind of cell that argument reaches.

Union-reported membership is the older source. Before the CPS carried a union question, state density estimates relied “on the financial reports made by labor unions to the Department of Labor, along with supplemental information obtained from unions and employee associations not reporting.”[26] Where the two sources overlap, for 1974, 1976 and 1978, the median ratio of CPS to directory estimates across states was 1.02, and cross-section regressions of one on the other gave adjusted R-squared values of 0.865, 0.859 and 0.839.[27] The directories appear to have overstated membership, and survey respondents may have understated their affiliation with employee associations.[28] Those financial reports are the annual reports this paper uses, filed then as now with the Department of Labor. This paper uses them one level down, at the local rather than the state, for a single branch of a single trade.

3. Classification

Three variables classify each local, and none of the three was constructed by the author from judgement alone. Each comes from a published external source, and where a crosswalk from that source to the filing data required judgement, the judgement is documented in the replication script and its limits are stated.

3.1 The sample

The sample is every IBEW local union that reported 1,000 or more members in its 2018 filing and a non-blank membership figure in its 2025 filing. “Local union” is the designation in the filing; districts, state councils and the international body are excluded. Where a local-year has an unamended filing, that filing is used; where it has only amendments, the highest-numbered amendment is used. A membership field that is blank or not a whole number is treated as missing, never 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, on the ground that such moves are transcription errors rather than events.[29] The sample is 154 locals.

3.2 Branch: which locals do construction

The IBEW describes its members’ work as “utilities, construction, telecommunications, broadcasting, manufacturing, railroads and government,” and organises its departments on the same lines.[30] Within construction, the union’s own jurisdictional maps distinguish Inside work, the

wiring of buildings, from the rest, and it is the Inside maps this paper uses. A local may hold more than one kind of jurisdiction, and many do.

The IBEW publishes a map for each state showing the Inside construction jurisdiction of each local.[31] The 53 PDFs linked from the union’s index page, which cover the fifty states, the District of Columbia, the Caribbean Islands and Panama, were retrieved and parsed. Among purely numeric text items the two kinds separate cleanly by size: highway shields sit at 12 and below, local-union labels at 16.5 and above, and nothing falls between. City names and other non-numeric text run up to 26.4 units and overlap both ranges, so the rule is a conjunction: a label is a text item that is numeric and set at 16 or more.[32] Two locals appear on a state’s map only in a 12-point note, Local 24 on Virginia’s and Local 1925 on Kentucky’s, and are included on that basis; a third such note, to Local 654 on Delaware’s map, names a Pennsylvania local already on the roster. The result is a roster of 262 distinct Inside locals across the 51 United States jurisdictions.[33] Rendered-image validation, comparing the count of extracted labels against the count of polygons on the map, was done for Virginia, Ohio, Connecticut and Delaware, four of the 51; the rendered images were not kept.[34]

A local is classified Inside if its designation number appears on the roster. Of the 154 sample locals, 89 do and 65 do not. The roster establishes which local holds Inside jurisdiction for an area. It does not establish that a listed local is exclusively Inside, and it does not establish that an unlisted local does no construction.[35]

3.3 Treatment: which metros have hyperscale campuses

Three of the four hyperscale operators used here publish locations of campuses they own and operate. Google publishes a list by city and county with an in-development flag; Meta publishes a list by city, from which the file records a year for each site; Microsoft publishes state fact sheets and community pages naming cities and counties.[36] Amazon Web Services publishes no consolidated site list on the four pages examined; its campuses were assembled from fifteen of the company’s own investment announcements, fact sheets and news pages, and its coverage is incomplete by construction.[37] The combined file holds 124 sites: Google 42, Meta 28, Microsoft 30, AWS 24.[38]

Each site carries a metropolitan area assignment; three do not. The set of treated metros is read from that column, excluding blank entries and sites annotated as rural. It contains 48 metros. A local is treated if its 2018 headquarters city lies in one of them. The crosswalk from headquarters city to metro is a list of city-name prefixes per metro, built on metropolitan statistical area membership and extended after a blind replication identified suburbs the first pass had missed.[39] It was then checked against the Office of Management and Budget’s July 2023 delineations: each local’s 2018 headquarters address was geocoded to a county, the county mapped to its core based statistical area, and the 48 metro strings mapped to statistical areas by name. The two classifications agree for all 154 locals.[40] Each of the two earlier corrections to the prefix list had lowered the estimate.

Of the 89 Inside locals, 29 are treated. The 60 untreated divide into 38 whose state has at least one published campus somewhere, Group B, and 22 whose state has none, Group C. Group B is kept separate because a campus a hundred miles away draws on the same regional labor pool: Feyrer, Mansur and Sacerdote found the employment effect of new oil and gas production more than three times larger within 100 miles of the well than within the producing county,[41] and Rosenthal and Strange report that agglomeration effects reach across broad distances even as they weaken with it.[42] Group B is therefore an attenuation band rather than a control, and section 6 tests whether it behaves like one. Group C is the control, and its bound is: no campus published by Google,

Meta or Microsoft, and no Amazon announcement located, in that state.[43] The control states are Alaska, Colorado, Florida, Hawaii, Massachusetts, Montana, New Jersey, New York and West Virginia, and fourteen of the 22 control locals are in New York, New Jersey and Florida. Section 10 says what that composition costs.

The treated set is the third the project has classified. A hand-built list of treated metros gave an A-minus-C median contrast of 32.31 points; reading the 48 metros from the operators' file added six locals and gave 27.32; extending the crosswalk to suburbs a blind replication identified added three more and gave 24.75. The nine added locals have a median growth of 4.84 percent.[44]

3.4 Colocation market: the industry's own measure

CBRE publishes a semi-annual report on North American data center markets with operating inventory, under-construction capacity and vacancy by market.[45] The H2 2025 and H1 2026 editions each list eight primary markets and eight secondary; for H2 2025 both lists were read from the published figures on 5 September 2026 and agree with the file. The figures used here are from the H2 2025 edition. Across the three editions held, the transcribed market figures sum to the totals CBRE states at four points, including an H1 2026 primary-market inventory summing to 10,902.9 megawatts against a stated 10,903.[46] CBRE publishes no geographic definition of any market in the overview text held for the three editions, so the crosswalk from its sixteen market names to headquarters cities is again judgement, and the question of whether its inventory counts owner-operated hyperscale capacity is answered in that text neither way.[47]

4. Result

4.1 Three groups

Table 1 gives median growth in membership from the 2018 filing to the 2025 filing for the 89 Inside locals, split three ways. Group figures are medians because the filing does not separate a merger from organic growth, and a mean would let one absorption of a neighbouring local carry a group. The regressions of sections 4.3 and 5 estimate conditional means and are reported beside a median regression for that reason.

Table 1. Inside construction locals, median membership growth, 2018 to 2025[48]

Group	n	Median growth	Growing more than 25 percent
A. Treated metro	29	+33.74%	16 of 29
B. Untreated metro, campus published elsewhere in the state	38	+16.37%	12 of 38
C. Untreated metro, no campus published in the state	22	+8.99%	1 of 22
B and C together	60	+13.24%	13 of 60

The contrast between A and C is 24.75 points. Between A and the two untreated groups together it is 20.50. The ordering across the three groups is monotone on the median and on the share

growing more than a quarter, which falls from sixteen of twenty-nine to twelve of thirty-eight to one of twenty-two.

In members rather than percentages, the 29 treated locals reported 85,728 members in 2018 and 111,370 in 2025, a rise of 29.9 percent; the 60 untreated reported 163,832 and 175,475, a rise of 7.1 percent.[49] The treated group added 25,642 members from a base half the size of the untreated group, which added 11,643.

4.2 The treated locals

Table 2 lists the 29 treated locals individually so that the group can be reconstructed and the concentration reported in section 5 can be seen before it is measured. Growth is the 2025 filing over the 2018 filing.

Table 2. Inside locals in treated metros, ranked by growth[50]

Local	Headquarters	Metro	Operators publishing	2018	2025	Growth
640	Phoenix AZ	Phoenix	3	1,643	4,410	+168.4%
60	San Antonio TX	San Antonio	1	1,416	3,125	+120.7%
666	Highland Springs VA	Richmond	3	1,055	2,300	+118.0%
683	Columbus OH	Columbus	4	1,918	3,909	+103.8%
20	Grand Prairie TX	Dallas-Fort Worth	2	2,702	5,192	+92.2%
369	Louisville KY	Louisville	1	3,016	5,280	+75.1%
291	Boise ID	Boise	1	1,075	1,848	+71.9%
474	Memphis TN	Memphis	1	1,495	2,247	+50.3%
481	Indianapolis IN	Indianapolis	2	3,236	4,671	+44.3%
429	Nashville TN	Nashville	1	1,612	2,312	+43.4%
613	Atlanta GA	Atlanta	4	4,401	6,125	+39.2%
347	Des Moines IA	Des Moines	2	1,621	2,238	+38.1%
153	South Bend IN	South Bend	1	1,486	2,048	+37.8%
26	Lanham MD	Northern Virginia	3	8,960	12,078	+34.8%
124	Kansas City MO	Kansas City	2	3,400	4,547	+33.7%
22	Omaha NE	Omaha	2	1,857	2,457	+32.3%
357	Las Vegas NV	Las Vegas	1	3,354	4,061	+21.1%
697	Merrillville IN	Chicago	2	1,314	1,539	+17.1%
611	Albuquerque NM	Albuquerque	1	2,426	2,836	+16.9%
292	Minneapolis MN	Minneapolis	1	4,249	4,680	+10.1%
58	Detroit MI	Detroit	1	4,796	5,028	+4.8%
332	San Jose CA	San Jose	1	3,365	3,484	+3.5%
8	Rossford OH	Toledo	1	2,103	2,161	+2.8%
110	St. Paul MN	Minneapolis	1	2,433	2,469	+1.5%
176	Joliet IL	Chicago	2	1,245	1,262	+1.4%
134	Chicago IL	Chicago	2	12,288	12,161	−1.0%
98	Philadelphia PA	Philadelphia	1	4,610	4,544	−1.4%
701	Warrenville IL	Chicago	2	1,565	1,398	−10.7%
1579	Augusta GA	Augusta	2	1,087	960	−11.7%

Twenty-nine locals in twenty-five metros; Chicago contributes four and Minneapolis two. Four

locals more than doubled. Four declined. The Chicago metro, whose Local 134 is the largest treated local at 12,288 members in 2018, is flat across three of its four locals, with the fourth at 17 percent, below the treated median; its treatment in section 5 matters for the direction of the estimate.

4.3 Regression

Table 3 puts standard errors on the contrast. The dependent variable is growth in percentage points from 2018 to 2025. Standard errors are clustered on state, which gives 36 clusters among Inside locals and 41 in the full sample, because locals in one state share a labor market and a regulatory environment.[51]

Table 3. Growth 2018 to 2025 in percentage points, OLS, standard errors clustered by state[52]

Specification	Coefficient	95 percent interval	p	n
Inside locals only: treated	+25.83	[+7.84, +43.82]	0.005	89
Inside only: treated, with Group B as separate term, control is C	+35.72	[+15.95, +55.49]	<0.001	89
Group B coefficient in that model	+15.61	[+5.78, +25.44]	0.002	
All IBEW locals: treated $\times$ Inside interaction	+19.68	[−2.03, +41.40]	0.076	154
All IBEW locals: treated $\times$ Inside, controlling log 2018 membership	+20.69	[−1.71, +43.09]	0.070	154

The treated coefficient among Inside locals is 25.8 points with an interval that excludes zero at the 1 percent level. A randomization test that reassigns treatment to 29 of the 89 locals at random 5,000 times finds a difference in means at least this large in 0.06 percent of draws.[49] Weighting locals by 2018 membership gives 22.8 points, $p = 0.009$. Against the clean control alone the coefficient is 35.7, Group B sits 15.6 points above that control, and the difference between the treated and Group B coefficients, 20.1 points, is itself significant at $p = 0.026$.[49] The three-group gradient of Table 1 holds under a standard error. The interaction in the full sample, the difference-in-differences against locals not on the Inside maps, is 19.7 points without state fixed effects and does not reach 5 percent. With state fixed effects added it is 15.4 points with a standard error of 17.2, $p = 0.37$.[49] Section 6 takes that up; the short version is that the branch comparison is informative on medians and underpowered as a regression. One further coefficient from the full-sample model belongs on the record. In untreated metros Inside locals grew 10.0 points faster than the others, $p = 0.020$; in

treated metros the same premium is 29.7 points, $p = 0.005$.^[53] The treatment effect sits on top of a construction premium that is present without it.

5. Timing and robustness

5.1 Event study

The cross-section says treated locals grew more. It does not say when. Every one of the 89 Inside locals has a membership figure in all eight filing years, so the balanced panel is the full Inside sample. An event study on that panel, with local fixed effects, year fixed effects and a treated-by-year interaction relative to 2018, places the divergence in time with error bars. The outcome is log membership; the coefficients approximate percentage differences from the 2018 base.

Table 4. Event study, log membership, local and year fixed effects, treated $\times$ year, 2018 base^[54]

Year	Effect $\times 100$	95 percent interval	p
2019	+1.17	$[-2.34, +4.68]$	0.513
2020	+1.83	$[-2.85, +6.51]$	0.443
2021	+3.32	$[-1.90, +8.54]$	0.212
2022	+5.75	$[-1.17, +12.68]$	0.104
2023	+9.80	$[+0.74, +18.87]$	0.034
2024	+13.71	$[+2.82, +24.61]$	0.014
2025	+17.80	$[+4.94, +30.66]$	0.007

712 observations, 89 locals, 29 treated, 36 state clusters. A joint test that the 2019 and 2020 effects are zero gives $p = 0.739$.

Treated and untreated Inside locals moved together through 2020, with point estimates under two points. The estimate exceeds three points from 2021, is significant from 2023, and reaches 17.8 log points by 2025, about 19.5 percent in levels. It is monotone across all seven years. Greenstone, Hornbeck and Moretti reported the absence of differential pre-trends as a check on their design, and it is the check this paper can run.^[55] The absence here is bounded: two pre-period years, both estimated near zero with intervals that would admit a differential of four to six points, on a panel that begins in 2018 and cannot see the years before it.

The table is also not an event study in the strict sense, because the treatment has no single date. The campus file carries a year for each of Meta’s 28 sites and for 17 of Amazon’s 24, one for Microsoft’s 30, and none for Google’s 42.^[38] The Meta years run from 2010, and Altoona is dated 2013, Fort Worth 2015, Los Lunas 2016, and Henrico, New Albany and Sarpy 2017; Amazon’s New Albany site is dated 2015. So for Des Moines, Dallas-Fort Worth, Albuquerque, Richmond, Columbus and Omaha the site predates the panel, and 2019 and 2020 are not pre-period years there. At the other end, six treated metros have every site dated and dated 2025 or later, and two of them hold sample locals: Philadelphia, whose only entry is an Amazon announcement of 2025, and Toledo, whose Meta site is dated 2025.^[49] Eighteen treated metros carry no dated site at all, and two more carry a dated 2025 site alongside an undated one. Those two locals grew 0.7 percent at the median, and dropping them raises the cross-section coefficient to 28.7. The binary treatment

therefore pools metros treated before the panel and metros barely treated within it, and the 2023 divergence in Table 4 is the average over that mix. A staggered design keyed to break-ground dates is the right estimator and cannot be built from the file as published, because the file dates none of Google’s sites and only one of Microsoft’s.

The cumulative median series tells the same story with one wrinkle the regression smooths. On raw medians the gap between treated and untreated is 1.37 points in 2019, 1.02 in 2020, 2.19 in 2021, 8.01 in 2022, 9.20 in 2023, 13.62 in 2024 and 20.50 in 2025. It narrows by a third of a point between 2019 and 2020 and widens in every other interval. The three largest one-year increments are 6.88 in 2024 to 2025, 5.82 in 2021 to 2022 and 4.42 in 2023 to 2024.[56] An earlier pass on this data, run before any Inside classification existed and with every IBEW local in both groups, found a roughly flat divergence and read it as a trend predating the current build-out. That reading is withdrawn. On classified locals the gap is about two points through 2021 and reaches eight in 2022.

5.2 Robustness, and the cut that fails

Table 5 re-estimates the Inside-only treated coefficient under seven cuts. The one that fails is placed in bold and discussed first.

Table 5. Coefficient on treated, Inside locals, alternative samples[57]

Cut	Coefficient	95 percent interval	p	n, treated
All Inside locals	+25.83	[+7.84, +43.82]	0.005	89, 29
Drop the five fastest-growing locals	+9.03	[−2.09, +20.15]	0.112	84, 24
Drop five locals with any year-over-year move above 40 percent	+15.75	[+0.45, +31.05]	0.044	84, 25
December fiscal-year filers only	+25.36	[+5.65, +45.07]	0.012	60, 24
June fiscal-year filers only	+46.09	[−19.15, +111.33]	0.166	23, 4
Drop Chicago metro	+31.95	[+13.36, +50.55]	0.001	85, 25
Drop Phoenix metro	+21.24	[+5.76, +36.73]	0.007	88, 28
Median regression	+20.07	[+5.15, +34.98]		89

Removing the five fastest-growing locals, which are Phoenix, San Antonio, Richmond, Columbus and Dallas-Fort Worth, cuts the coefficient to 9.0 points and the interval crosses zero. That cut trims one tail of the treated distribution and nothing else, so it shrinks any estimate mechanically; it is reported because a reader will ask for it, and the honest gloss is that the average treatment

effect is carried by a minority of treated metros. Less mechanical checks give a narrower range. Leaving out any single local moves the coefficient between 21.2 and 27.7. Removing the five locals with any single-year move above 40 percent, two of which are in the top five, gives 15.8 at $p = 0.044$. The median regression at 20.1 shows the centre of the treated distribution moved as well as its tail.[49]

A second cut names a confound the placebo cannot reach. A semiconductor fab or a battery plant is wired by the same locals that wire a data center, and four treated metros known to the author to host one are Phoenix (TSMC), Columbus (Intel), Boise (Micron) and Kansas City (Panasonic, De Soto). No project list was retrieved, so the four are a floor and not a census; Micron’s Clay fab sits in the Syracuse metro, whose Local 43 is in the control group, so the same confound runs on both sides of the contrast. The federal legislation subsidising semiconductor and battery manufacturing was enacted in August 2022, the year of the second-largest one-year widening in the median gap and the year from which the event-study coefficient more than doubles. Removing the four treated metros leaves 17.1 points, interval 0.2 to 34.1, $p = 0.048$.[49] The effect survives and is a third smaller. Section 10 states what would separate the two channels.

The remaining cuts hold. Restricting to December fiscal-year filers, 60 of the 89, gives 25.4 at $p = 0.012$, which closes the concern that the result is an artefact of comparing filings that close in different months. June filers are four treated locals and uninformative. Chicago pulls the estimate down rather than up: its four treated locals are all below the treated median and removing them raises the coefficient to 32.0. Removing Phoenix, the fastest-growing local, leaves 21.2 at $p = 0.007$. Removing the two 2025-dated metros of section 5.1 gives 28.7 at $p = 0.004$.

6. The placebo

The design substitutes a within-metro branch placebo for the runner-up control it cannot have. The 65 sample locals not on the Inside roster work in the union’s other fields, utilities, telecommunications, manufacturing, railroads, broadcasting, government and construction outside the Inside maps, and the roster does not say which. They sit in the same metros, file on the same forms, and do not hold the Inside jurisdiction under which the buildings are wired. Table 6 applies the same crosswalk to them.

Table 6. Non-Inside IBEW locals, median membership growth, 2018 to 2025[58]

Group	n	Median growth
Treated metro	23	+12.89%
Elsewhere	42	+4.17%
Difference		+8.72 points

Non-Inside locals in treated metros grew 8.7 points faster than non-Inside locals elsewhere. Inside locals in treated metros grew 20.5 points faster than Inside locals elsewhere. The difference of differences is 11.8 points on medians, and 19.7 points in the regression interaction of Table 3, where it does not reach conventional significance, $p = 0.076$.[59]

Two readings are open. The first is that the 8.7 points is a general metro effect, which is what a placebo is meant to absorb, and the excess for Inside locals is the construction channel. The second

is that non-Inside locals also benefit from the build-out through work the roster does not label as Inside: utility locals on the substations and transmission that a campus requires, line locals on the interconnection. On the second reading the placebo is partly treated and the 11.8 to 19.7 point excess understates the construction effect. The paper does not choose between them. What Table 6 weighs against is a reading in which the whole of Table 1 is a metro effect unrelated to what the locals do; on that reading the two branches would move together, and on medians they do not. The regression form of the same comparison is weaker. The interaction is 19.7 points at $p = 0.076$ without state fixed effects, and 15.4 at $p = 0.37$ with them. Only twenty-five states hold locals of both kinds, and in eleven of those each side is one or two locals, so the within-state comparison has little to work with.[49] The placebo is evidence on medians and not a test that survives fixed effects.

The three-group split can be run on the non-Inside locals too, and it should be, because if Group B sits above Group C there as well, part of the Group B premium in Table 1 is a state effect and not attenuation. It does: non-Inside medians are 12.9 in treated metros, 10.5 in Group B states and -1.1 in Group C states.[49] The Inside gradient is 33.7, 16.4 and 9.0. The Group B premium over C is about seven points in Inside locals and about twelve in non-Inside locals, so the band carries a state-level component that the design does not separate from distance to a campus. A distance function from each local's county to the nearest published campus would separate them and has not been built.

The non-Inside treated group is small enough that its 23 members are listed in the endnote so the row can be rebuilt.[60] Its spread is wide, from Grove City OH at $+48.2$ percent to Downers Grove IL Local 21 at -51.5 , and the ten suburban locals added in the last crosswalk revision raised its median. That is the number the stated rule produces, and a narrower crosswalk would have produced a larger difference of differences.

Helm, using German trade shocks, found that employment spillovers between local industries are larger where the industries employ similar workers and strongest within the same broad sector.[61] Greenstone and coauthors found the same on worker flows.[11] The prediction for this design is that the non-Inside locals closest in labor to Inside construction, those doing line and utility construction, should show the largest response, and those furthest from it, manufacturing and telecommunications, the smallest. The roster does not separate the non-Inside locals by field, so that prediction is not tested here. It is the first thing a finer classification would test.

7. Benchmark against federal employment data

The objection to any membership series is that it measures organising as well as work. The Quarterly Census of Employment and Wages counts jobs covered by unemployment insurance, by county and six-digit industry, from employer tax filings, and is independent of both the union and the industry.[62] NAICS 238212 covers nonresidential electrical contractors, the classification under which a contractor wiring a data center would report. County employment in that industry was summed to each sample local's statistical area for 2018 and 2025, over the counties disclosed in both years, and compared with the local's membership growth.[63] Eighty-six of the 89 Inside locals have a usable series. The three without are West Frankfort, Illinois, which lies outside every statistical area in the 2023 delineations, and Kalispell, Montana and Medford, Oregon, whose single counties are suppressed in the 2018 file. The five-digit column also loses Memphis, where the counties suppressed in one year or the other held 90 percent of the area's disclosed 2018 employment.

Table 7. Membership growth against electrical contractor employment growth by statistical area, 2018 to 2025[63]

	NAICS 238212, nonresidential	NAICS 23821, all electrical contractors
Locals with a usable series	86	85
Correlation, Pearson	0.58	0.49
Correlation, Spearman	0.61	0.59
Slope of membership growth on employment growth (SE)	0.69 (0.15)	0.65 (0.20)
Employment growth, treated metros, median	+30.9%	+27.4%
Employment growth, untreated metros, median	+14.3%	+11.4%
Employment growth, treated minus untreated, OLS	+19.6 [+6.7, +32.5], p = 0.003	+13.4 [+1.1, +25.6], p = 0.032
Membership growth on treated, controlling for employment growth	+12.6, p = 0.14	+16.9, p = 0.07

Three things are in the table. First, the two series move together: a statistical area whose non-residential electrical contractor employment grew ten points faster saw its Inside local grow about seven points faster, and the rank correlation is 0.61. Second, the employment series reproduces the treatment contrast without any reference to the union. Nonresidential electrical contractor employment grew 30.9 percent at the median across the statistical areas of the 29 treated locals against 14.3 percent across those of the 57 untreated, a regression difference of 19.6 points with an interval that excludes zero at the 1 percent level, clustered by state. That is a second, independent measure of the same divergence, from tax filings rather than union filings, and it is the closest thing this paper has to a replication. Third, adding employment growth to the membership regression cuts the treated coefficient from 25.8 to 12.6 and it loses significance. That is what a membership series that tracks employment should do. About half of the membership effect is employment growth in the same industry and place; the remainder is the union's changing share of that employment, or measurement in one series or the other.

The metros where the two series part company are informative. Columbus, Boise, Indianapolis, Nashville and Des Moines sit near the line: employment up 91, 99, 52, 38 and 35 percent against membership up 104, 72, 44, 43 and 38. The largest gaps on the membership side, over all 86 locals, are Richmond, Phoenix and Louisville, with membership up 118, 168 and 75 percent against employment up 32, 90 and 6. In Richmond, Louisville and Omaha the roll grew more than three times as fast as the covered electrical workforce of the statistical area. That points to organising or to a jurisdiction reaching beyond the statistical area, and the filing cannot say which; it cannot point to travelers, who stay on their home local's roll while their work is counted where it is performed, which pushes membership below employment in the receiving metro wherever the contractor employing them reports there. Among treated locals the largest gaps the other way are South Bend, Augusta and Detroit, with employment up 128, 26 and 39 percent while membership rose 38 percent, fell 12, and rose 5; over all 86, Beaumont, Trenton and Moline are larger still.[63]

The benchmark has its own limits. QCEW counts a job at the employer's establishment, so an

electrical contractor based in one metro and working in another is counted at home, which is the same direction of error as the permit problem in section 9 and does not cancel it. County cells are suppressed where an employer would be identifiable; the consistent-panel sum drops such counties in both years. For the nonresidential code the dropped counties hold at most 9 percent of any area's disclosed 2018 employment. For the five-digit total they exceed 10 percent in six of the 85 usable areas, and restricting that column to the 79 below 10 percent raises its correlation to 0.61 and its treatment difference to 16.0 points, $p = 0.007$.^[63] The 2025 annual file carries a server date of 21 August 2026 and is subject to BLS revision. The five-digit total, which adds residential electrical contractors, gives the same picture with more noise. And the areas compared are the areas where Inside locals sit, so the employment series is independent of the union as a source and not in its geography.

8. The industry's market measure against the trade's growth

CBRE's semi-annual report is widely cited for data center capacity by market.^[45] Twelve of the 29 treated locals sit in one of its sixteen markets. Seventeen do not.

Table 8. Treated Inside locals in CBRE markets, H2 2025 capacity against local growth^[64]

Market	Operating MW	Under construction MW	Local	Growth
Phoenix	807.3	176.0	640	+168.4%
Austin-San Antonio	216.1	469.0	60	+120.7%
Dallas-Fort Worth	1,067.3	698.7	20	+92.2%
Atlanta	1,459.2	2,076.0	613	+39.2%
Northern Virginia	4,039.6	1,905.7	26	+34.8%
Chicago	904.6	715.1	697	+17.1%
Minneapolis	65.1	12.0	292	+10.1%
Silicon Valley	489.2	144.0	332	+3.5%
Minneapolis	65.1	12.0	110	+1.5%
Chicago	904.6	715.1	176	+1.4%
Chicago	904.6	715.1	134	−1.0%
Chicago	904.6	715.1	701	−10.7%

Across the twelve, the Pearson correlation between local growth and log operating inventory is 0.093, and between local growth and log under-construction capacity 0.106.^[65] CBRE capacity does not predict Inside local growth on this sample and this edition. Phoenix has the lowest ratio of under-construction to operating capacity among the primary markets and the highest growth in the dataset. Northern Virginia has more than twice the operating inventory of any other market and its local sits mid-table.

The seventeen treated locals outside every market in the H2 2025 edition are Richmond, Columbus, Louisville, Boise, Memphis, Indianapolis, Nashville, Des Moines, South Bend, Kansas City, Omaha, Las Vegas, Albuquerque, Detroit, Toledo, Philadelphia and Augusta. Four of them grew more than 70 percent. Richmond, Columbus and Louisville, three of the six fastest-growing treated locals, are in metros that the capacity report does not list in that edition.

The likely reason is that CBRE measures a different thing. Its H2 2025 and H1 2026 market tables are headed “Wholesale Primary Market Fundamentals,” its rate commentary refers to “primary wholesale colocation markets,” and its 2018 edition defines wholesale colocation as “Building shell & infrastructure to PDU providing space, power & cooling” while defining a hyperscale cloud provider only as a type of user.[66] A campus that Google or Meta builds for its own use is not wholesale colocation. But in the text held for the three editions no statement says that owner-operated capacity is excluded from inventory, and none says it is included. That the self-build metros are missing because CBRE does not count self-build is an inference from framing. It is the most plausible reading and it is not a sourced fact.

Whatever the reason, the consequence for research is the same. A study that used CBRE market inventory as its treatment variable would place Richmond, Columbus and Louisville in the control group. The operators’ site lists, incomplete as they are, cover metros the market report does not. Augusta, where two operators publish a campus, appears as a market in none of the three CBRE editions held; editions between them were not examined.[67]

9. Virginia, and the local that testified

IBEW Local 26, headquartered in Lanham, Maryland, appears on the IBEW Inside maps for the District of Columbia, Maryland and Virginia. It filed 8,960 members in 2018 and 12,078 in 2025, a rise of 34.8 percent, fourteenth of the 29 treated locals and one point above the treated median.[68] It is a representative case within the group, not an exceptional one, in the metro CBRE reports with the largest operating inventory of the markets it lists.[45]

Its FY2025 Schedule 13 shows what the roll contains.[69]

Category as filed	Number
Journeyman Wireman	6,305
Residential Wireman	1,823
Apprentices and Trainees	3,471
Maintenance Workers	145
Technicians	223
Communications Workers	111
Total, equal to Item 20	12,078

Apprentices are 28.7 percent of the roll. Their count by fiscal year from 2018 is 2,280, 2,448, 2,248, 2,302, 2,514, 2,529, 2,431 and then 3,471: flat within a band of 300 for seven years and up by 1,040 in the eighth.[70] The local’s growth through 2024 was therefore in journeymen and other categories, and the apprentice surge is a 2025 event. If a pipeline effect is coming it has only just entered the filings.

On 6 July 2026, at a Loudoun County town hall, the local’s political coordinator said that the local had worked 14 million hours eight years earlier, 28 million in the past year, and expected to work north of 33 million, and that it had gone from 9,000 members to 17,000 over that period.[71] The base is accurate within forty. The 17,000 endpoint appears in no filing from FY2018 through FY2025, the latest of which is 12,078; the largest single-year gain in the series is 1,697, from 10,381 in FY2024, and no FY2026 report was on the reporting history as of 4 September 2026.[72] The most

likely reconciliation is that the speaker was counting people working under the local's agreements rather than members. Under the LM-2 instructions, work permit fees are "charged to nonmembers of the labor organization who work within its jurisdiction," and Item 20 counts members.[19] At a planning commission hearing seventeen days later the same speaker described his local's position in terms of work rather than membership: "we have 97 [percent] of the work in the data center[s]."[73]

The federal form and the podium are answering different questions. Item 20 asks how many members the organisation has. The testimony describes how many people the industry put to work through the local. Neither need be wrong. The second number has no field of its own on the form, so it enters the public record as the first and becomes uncheckable there. That is an observation about the reporting architecture, and the LM-2 series is what makes the gap visible.

10. Limitations

The design is descriptive. There is no instrument and no runner-up control. The identification is a within-branch comparison across metros and a within-metro comparison across branches, with local and year fixed effects in the panel. It weighs against the readings named in sections 5 and 6. It does not rule out a metro-level shock correlated with campus siting that moves Inside electricians and no other branch, and section 5.2 names the obvious one: the Inside branch also wires the semiconductor and battery plants that the 2022 federal legislation subsidised, in at least four of the treated metros. Whether those plants were under construction on the same timetable as the campuses was not established; no project list was retrieved. Removing those metros leaves two-thirds of the effect. Separating the channels needs a metro-level measure of non-data-center megaproject construction over the window, which this paper does not have.

The control is concentrated in three states. Fourteen of the 22 Group C locals are in New York, New Jersey and Florida, and six are in New York alone, including Local 3. The clean-control contrast therefore sets treated locals, four of whose five fastest-growing had 2018 rolls under 2,000, against a control weighted toward New York and New Jersey, on an unweighted percentage outcome. Group C's 2018 rolls are not tabulated in this paper. The Group B comparison of 20.1 points, $p = 0.026$, is the check that does not depend on Group C, and it holds. Weighting by 2018 membership, which gives 22.8, raises the leverage of the large Northeastern locals rather than removing it, and also holds.

Membership is not employment. The filing reports neither hours nor employment. Organising drives, jurisdictional transfers and a local's own choice of what to count in Schedule 13 all move Item 20 without moving work. The apprentice series of section 9 shows a category shift of 1,040 in one year inside one local. Medians limit the damage a single such event can do to a group figure and do nothing for a local figure. Section 9 also shows the direction of the largest measurement error. If the 17,000 figure counts travelers on permits, the roll undercounts the working electrical workforce most where demand is hottest, because permit workers are where the work is and are counted at home. That error is correlated with treatment and biases the estimate toward zero. It also means the series measures where locals grew their rolls, which is a union market-share question as much as an employment question. Section 7 benchmarks the series against county employment in nonresidential electrical contracting from the Quarterly Census of Employment and Wages and finds it tracks that series with a slope of 0.69; the metropolitan occupational series for electricians, SOC 47-2111 in the Occupational Employment and Wage Statistics, was not used. Richmond, Louisville and Omaha, where membership grew more than three times as fast as covered

employment, and Phoenix, which carries the second-largest gap in the file, are the treated locals where the market-share reading has the most room. Amalgamations of locals were not checked against the IBEW’s notices.

The effect is concentrated. Removing five locals removes significance. The claim the data support is that a minority of treated metros saw very large growth and that the centre of the treated distribution moved as well, against an untreated median of 13.2 percent and 13 of 60 untreated locals growing more than a quarter. Table 2 shows four treated locals that shrank.

Group sizes. 29 treated Inside locals, 22 clean controls, 23 treated non-Inside locals. The full-sample interaction has 154 observations in 41 state clusters, and its interval is wide before state fixed effects and uninformative after them.

Treatment is binary and undated. A campus begun in 2013 and an announcement of 2025 receive the same indicator. Section 5.1 shows what the dated subset of the file allows and why a staggered design cannot be built from it. Table 2 carries a count of operators per metro that could serve as a dose, and it has not been used as one; a regression of growth on that count is the next thing to run.

Group B may hold locals that belong in A. The two largest untreated growth figures are Knoxville, Local 760, at +86.9 percent, and Austin, Local 520, at +61.5; Baltimore, Local 24, is third at +58.5. Knoxville and Baltimore were individually checked on 4 September 2026 and no operator publishes a campus in either metro, bounded to Meta’s full US list, Google’s full worldwide list, Microsoft’s sixteen-state community navigation and twelve-state 2025 fact-sheet index, and Amazon’s data-center communities page and Maryland fact sheet.[74] Austin was not individually checked; the site file assigns Meta’s Temple campus to the Killeen-Temple statistical area, which is separate. Any of the three moving to A would raise the treated coefficient. Amazon’s absence from any metro is an absence of published evidence.[37]

Fiscal years differ. Of the 154 sample locals, 108 close their year in December, 38 in June and 8 in other months, so a 2018-to-2025 comparison spans different real windows across locals.[75] The December-only cut in Table 5 holds. No within-metro pair was checked for matching fiscal years.

The 2025 extract is thinner. 19,080 filings against 21,277 in 2018. How much is attrition and how much is late filing is undetermined, and it bears on every claim resting on the most recent year.[23] The balanced panel of section 5 is unaffected by construction, since every local in it filed in every year, but a local that stopped filing in 2025 is not in it.

The crosswalks are the author’s. Neither the operators’ site pages nor the CBRE editions examined publish a geographic definition of a market or a metro. The mapping from headquarters city to metro is a list of city-name prefixes in the replication script, and the OMB check of section 3.3 reproduces it for every sample local.[40] The check has its own limits: 142 of the 154 addresses were geocoded by the Census Bureau’s service and twelve were assigned to a county by ZIP code, and the 48 metro strings were matched to statistical-area titles by name, which for the Nashville, Salt Lake City and Milwaukee combined statistical areas means the named component area and not the combination. The CBRE market crosswalk has no such check, because CBRE publishes no market boundaries in the text held.

Displacement is not measured. Kline and Moretti, on the Tennessee Valley Authority, found that agglomeration gains in the treated region were offset by losses elsewhere in the country.[76] Nothing here speaks to whether the electricians who joined treated locals came from untreated ones, and the data cannot say, because a traveler on a permit in a treated jurisdiction remains a

member of the local he came from. Bartik’s review of the place-based literature adds that migration into a local labor market moves local demand by a similar or larger amount, and names housing construction as one channel,[77] so an inflow of electricians to a treated metro is itself a small construction shock there; the series cannot separate that from the campus.

11. Conclusion

The paper set out to find a workforce-side measure of the data center build-out that neither the industry nor its opponents produced, among the sources it could locate, and to see whether it moved. It found one in a form that labor organisations have filed since 1959, joined it to the union’s map of who holds construction jurisdiction, and defined treatment from the operators’ published site lists. Inside electrical locals in metros with a published hyperscale campus grew 25.8 points faster than Inside locals elsewhere over 2018 to 2025, with an interval of 7.8 to 43.8. The divergence is absent in the two pre-period years the panel can see, significant from 2023, and monotone. It is concentrated in a handful of metros. Four treated metros known to the author, a floor and not a census, also host a subsidised fab or battery plant, and the effect is two-thirds as large without them. IBEW locals in the same metros that are not on the Inside maps diverged by less than half as much on medians, and the regression form of that comparison does not survive state fixed effects. The industry’s own capacity measure, in its H2 2025 edition, does not list seventeen of the twenty-nine treated locals’ metros and does not predict growth among the twelve it does.

Three things follow for research on this industry. First, a union-side workforce series exists, at the metro and the trade, in a federal extract. Its limits are stated above, and each is a limit on a number that can be checked, and section 7 checks it against a federal employment series that correlates with it at 0.58 and shows the same treatment contrast. Survey-based union estimates at metropolitan level do exist. Kokkelenberg and Sockell published them for 1973 to 1981,[78] and Hirsch and Macpherson’s annual data book has carried metropolitan figures for years beginning with 1983.[24] What the latter’s own methods note says is that the sample will not carry them precisely, “with even greater force” at the metropolitan and detailed-industry level than at the state level.[25] The filings can be tabulated at that geography, and against the federal employment series for the same industry and places they track it with a slope of 0.69 and reproduce its treatment contrast.

Second, treatment should be defined from what the operators publish and not from what the market reports count. The CBRE result of section 8 is a negative with a bound, twelve locals and one edition, but a study that took market inventory as its exposure variable would have had Richmond, Columbus and Louisville in its control group, and those three locals grew 118, 104 and 75 percent.

Third, the question this paper cannot answer is the one Ellison and Glaeser said mean levels cannot answer, whether a concentration reflects an advantage of the place or spillovers among those concentrated there.[12] In this setting that is the question of whether a trade that concentrates around a construction boom stays concentrated when the boom ends. Allcott and Keniston found that resource booms in American counties raised local wages without crowding out manufacturing and left no evidence of foregone learning-by-doing.[79] Kline and Moretti found that manufacturing gains in the Tennessee Valley outlasted the subsidies that produced them.[76] Whether the electricians of Phoenix and Columbus are a resource boom or a big push is a question the next eight years of filings can answer. Among the works retrieved for this paper, the scholarship on data center labor is qualitative and operations-side, concerned with the small permanent workforce and

its uncertainties.[80] The construction workforce is the one these filings count, and they are public.

Endnotes

[1] Virginia Joint Legislative Audit and Review Commission, *Data Centers in Virginia*, Report 598 (2024), ch. 2, p. 14: “Data center representatives indicated that, at the height of construction, approximately 1,500 workers are on site building a facility and installing electrical and cooling systems.” The occupations listed on the same page are site developers and surveyors, equipment operators, steel and concrete erectors, electricians installing cabling, equipment and generators, and pipefitters and HVAC technicians. State primary source; the 1,500 figure is attributed in the report to industry representatives. Transcribed at Loudoun/Loudoun_Research_Origins_Operators_v1_0.md.

[2] Northern Virginia Technology Council, *Data Center Report* (2024), prepared by Mangum Economics, Table 2, Virginia 2023 in 2024 dollars: direct data center construction, 14,240 jobs, \$887.1 million pay and benefits, \$2,854.6 million economic output. Industry-association study. Transcribed at Loudoun/Loudoun_Research_Origins_Operators_v1_0.md § 2.2. Advocacy-adjacent source; the figure is reported as the industry’s own.

[3] Don Slaiman, Loudoun County Data Center Town Hall, 6 July 2026, transcript at [01:00:24]–[01:00:48], held at Loudoun/Loudoun_transcripts/ and quoted in Loudoun_Extract_Case_Hearings_v1_0.md § 2.9. Speaker self-identification as transcribed on 6 July: “I represent the IBW.” The title “political coordinator for the Electricians Union Local 26” is his self-identification at the 23 July hearing, note 73. Advocate-sourced claim.

[4] U.S. Department of Labor, Office of Labor-Management Standards, *Instructions for Form LM-2 Labor Organization Annual Report*, revised 2020, Item 20 and Schedule 13, dol.gov/sites/dolgov/files/olms/regs/compliance/GPEA_Forms/2020/efile/LM-2_instructionsRevised2020.pdf. Text saved at `olms_src/LM-2_Instructions_Revised2020.txt`, retrieved 5 September 2026; the saved text runs through Schedule 16 and is truncated there by the retrieval tool, which covers every passage cited in this paper.

[5] Form LM-2 as printed on the filed report at note 5 is the complete form; no item on it asks for hours worked. Items 70–71, signature declaration, as printed on IBEW Local 26 FY2025 report, file 012-627, held at 012-627 (LM2) 06_30_2025 IBEW.pdf, p. 1: “Each of the undersigned, duly authorized officers of the above labor organization, declares, under penalty of perjury and other applicable penalties of law, that all of the information submitted in this report (including information contained in any accompanying documents) has been examined by the signatory and is, to the best of the undersigned individual’s knowledge and belief, true, correct and complete.”

[6] OLMS Online Public Disclosure Room, Yearly Data Download, olmsapps.dol.gov/olpdr/. Extracts are pipe-delimited relational tables by filing year. The earliest year on the reporting history at note 68 is FY2000; the extract’s own start year was not checked.

[7] International Brotherhood of Electrical Workers, Inside Jurisdictional Maps, ibew.org/inside-jurisdictional-maps, one PDF per state, 53 linked files retrieved 4 September 2026.

[8] Barry T. Hirsch, David A. Macpherson and Wayne G. Vroman, “State-Level Estimates of Union Density, 1964 to Present,” author manuscript (2001), fn. 2, unionstats.com/mlr/MLR_StateUnion_July%202001.pdf. The condensed published version, “Estimates of Union Density by State,” *Monthly Labor Review*

124, no. 7 (July 2001): 51–55, does not carry this footnote. Both texts are saved in full at `lm2_lit/Hirsch_2001_state_union_density_MLR_FULLTEXT.txt`, retrieved 5 September 2026.

[9] Hirsch, Macpherson and Vroman, *MLR* 2001, p. 52. See note 27.

[10] Michael Greenstone, Richard Hornbeck and Enrico Moretti, “Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings,” *Journal of Political Economy* 118, no. 3 (2010). Working-paper version, NBER WP 13833, saved at `lm2_lit/Greenstone_2010_agglomeration_spillovers.txt`; page references are the working paper’s.

[11] Greenstone, Hornbeck and Moretti, WP 13833, p. 31: “the estimated increase in TFP for plants in the same 2-digit industry is a statistically significant 17% in Model 1 ... the estimates for plants in other industries are a statistically insignificant 3.3% in Model 1.” Introduction: “A one standard deviation increase in our measure of worker transitions is associated with a 7 percentage point increase in the magnitude of the spillover.”

[12] Glenn Ellison and Edward L. Glaeser, “Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach,” *Journal of Political Economy* 105, no. 5 (1997). Working-paper version, NBER WP 4840 (1994), is the text held; the passage follows Proposition 4 there: “we cannot distinguish comparative advantage from spillover theories based only on the mean levels of geographic concentration.” Saved at `lm2_lit/Ellison_1997_dartboard.txt`.

[13] Labor-Management Reporting and Disclosure Act of 1959, § 201(b), 29 U.S.C. § 431(b): “Every labor organization shall file annually with the Secretary a financial report signed by its president and treasurer or corresponding principal officers containing the following information in such detail as may be necessary accurately to disclose its financial condition and operations for its preceding fiscal year”. Cornell LII text saved at `cites_src/29USC431.txt`, retrieved 5 September 2026. The perjury declaration is on the form, note 5, not in § 431.

[14] LM-2 Instructions, General Instructions § II: “Every labor organization subject to the LM-RDA, CSRA, or FSA with total annual receipts of \$250,000 or more must file Form LM-2.” Also the form header on the filing at note 5: “MUST BE USED BY LABOR ORGANIZATIONS WITH \$250,000 OR MORE IN TOTAL ANNUAL RECEIPTS.” The simplified-report regulation, 29 C.F.R. § 403.4(a)(1), now reads “less than \$350,000” for Form LM-3, as amended at 91 FR 32622 (1 June 2026), effective 1 July 2026 and applicable “prospectively to labor organizations whose fiscal years begin on or after July 1, 2026.” Saved at `cites_src/29CFR403-4.txt` and `cites_src/FR_2026-10849_LM2_rule.txt`.

[15] Form LM-2, Item 20, as printed on the filed report at note 5, p. 2: “How many members did the labor organization have at the end of the reporting period?” The instructions describe the item rather than reproduce it: “After Schedule 13 is completed and the ‘Save and Calculate’ button is clicked, the software will enter the total number of members into Item 20.”

[16] LM-2 Instructions, Schedule 13: “Enter in Column (A) the categories of membership tracked by the reporting labor organization. ... In Column (B) enter the number of members for each of the membership categories listed in Column (A).”

[17] LM-2 Instructions, Schedule 13: “Define each category of membership in Item 69 (Additional Information). The definition should include a description of the members covered by the category and indicate whether the members pay full dues.”

[18] LM-2 Instructions, Schedule 13, Agency Fee Payers line and Total Members/Fee Payers line:

“Agency fee paying nonmembers are those who make payments in lieu of dues” and “This total in Column (B) is not the total number of members of the labor organization reportable in Item 20.”

[19] LM-2 Instructions, Item 21(e): “Work permit fees are fees charged to nonmembers of the labor organization who work within its jurisdiction.”

[20] OLMS Yearly Data Download file layout: `lm_data` and `ar_membership` among the tables, joined on `RPT_ID`.

[21] Extracts for filing years 2018 through 2025, retrieved from `olmsapps.dol.gov/olpdr/GetYearlyFileServlet` in September 2026 and held in the project folder by year. Author’s data.

[22] IBEW Local 26 FY2023 filing, Schedule 13 category label as filed: “APPRETICES AND TRAINEES”. Noted at `LM2_National_Trades_Analysis_2026-09-04.md`.

[23] Record counts in the `lm_data` tables of the 2018 and 2025 extracts. Author’s count.

[24] Barry T. Hirsch and David A. Macpherson, “Union Membership and Coverage Database from the Current Population Survey: Note,” *ILR Review* 56, no. 2 (2003): 349–354, as cited on `unionstats.com`, which describes coverage of 1973 to 1981 from the May CPS and 1983 to 2025 from the CPS outgoing rotation group files. Site text saved at `cites_src/unionstats_home.txt`, retrieved 5 September 2026. On the metropolitan series, Hirsch, Macpherson and Vroman, *MLR* 2001, endnote 5 and p. 51: the annual *Union Membership and Earnings Data Book* carries “disaggregated union and earnings figures beginning with 1983 for States, metropolitan areas, detailed industries, and detailed occupations,” and the CPS compilations are “the primary source for recent estimates of union density for States, as well as for metropolitan areas, detailed industry, and detailed occupation.”

[25] Hirsch, Macpherson and Vroman, author manuscript, fn. 2, not in the published version. The sentence concerns precision, not availability; see note 24 for the metropolitan series the same authors publish: “Because CPS sample sizes within the smallest states do not allow precise estimates of union density, BLS has been understandably reluctant to publish official estimates of state union density, an argument that applies with even greater force to density estimates for metropolitan areas and detailed industries and occupations.”

[26] Hirsch, Macpherson and Vroman, *MLR* 2001, p. 53 (“State” capitalised there); manuscript p. 7: “Estimates of state union density prior to the CPS rely on the financial reports made by labor unions to the Department of Labor, along with supplemental information obtained from unions and employee associations not reporting.”

[27] Hirsch, Macpherson and Vroman, *MLR* 2001, p. 52; manuscript pp. 5–6: “The median ratio was 1.02, with 22 of 51 being smaller than 1.0. ... Cross section regressions for the three years, with the CPS unionization rate estimates regressed on the Directory estimates, yielded adjusted R2s of 0.865 in 1974, 0.859 in 1976 and 0.839 in 1978 and standard errors of 5.0 to 5.1 percentage points.”

[28] Hirsch, Macpherson and Vroman, *MLR* 2001, p. 54, endnote 5; manuscript fn. 6: “The Directory appears to overstate member and association membership, whereas respondents in the CPS may understate their affiliation with employee associations.”

[29] Specification as coded in `lm2_v2_1.py`; results in `lm2_v2_1_results.json`. Author’s calculations.

[30] International Brotherhood of Electrical Workers, `ibew.org`, home and About pages, retrieved 5 September 2026 and saved at `cites_src/IBEW_branches.txt`: the union “represents approxi-

mately 901,000 active members and retirees who work in a wide variety of fields, including utilities, construction, telecommunications, broadcasting, manufacturing, railroads and government,” and lists departments for Broadcast & Recording, Construction & Maintenance, Government Employees, Manufacturing, Railroad, Telecommunications and Utility. The Inside and Outside distinction within construction is taken from the maps at note 7, whose titles read “Inside Jurisdiction.”

[31] See note 7.

[32] Every text item of size 8 and above from all 53 PDFs, with each PDF’s byte length and SHA-256, is saved at `ibew_src/IBEW_Inside_maps_textitems_2026-09-05.txt`, extracted with `pdf.js` on 5 September 2026 and verified by checksum against the in-browser string. `lm2_ibew_roster.py` rebuilds the roster from that file and reproduces `IBEW_Inside_Locals_Roster.csv` exactly: 329 rows, 262 locals, 327 from labels and two from notes. Author’s calculations.

[33] `IBEW_Inside_Locals_Roster.csv`, 329 rows, 262 distinct locals, 51 jurisdictions.

[34] Rendered-image validation for Virginia, Ohio, Connecticut and Delaware, 4 September 2026, comparing extracted label count against map polygon count. Four of 51 jurisdictions.

[35] Bound on the roster: it records Inside jurisdiction and nothing else. A listed local may also hold other branches; an unlisted local may do construction work not on the Inside map.

[36] Google, `datacenters.google/locations/`; Meta, `datacenters.atmeta.com/` site list; Microsoft state fact sheets and community pages. URLs per row in `Hyperscaler_Campus_Locations.csv`. Operator-sourced.

[37] AWS publishes no consolidated US site list on the pages examined: `aws.amazon.com/about-aws/global-infrastructure`, `aws.amazon.com/trust-center/data-center/`, `aboutamazon.com/news/aws/amazon-web-services-data-center`, `aboutamazon.com/impact/amazon-united-states`. Its 24 rows were assembled from Amazon investment announcements and state fact sheets. An AWS absence anywhere in this paper is an absence of published evidence on those pages.

[38] `Hyperscaler_Campus_Locations.csv`, 124 rows: Google 42, Meta 28, Microsoft 30, AWS 24.

[39] `CROSSWALK` dictionary in `lm2_v2_1.py`. Rows the site file annotates as CSA-only are treated as their own named metro. Not checked against an OMB delineation file.

[40] `lm2_omb_crosswalk.py`, reading `OMB_list1_2023.xlsx` (Census Bureau, List 1, July 2023 delineations, OMB Bulletin 23-01), `lm2_sample_county_geocode.csv` (Census Geocoder, benchmark Public_AR_Current, Counties layer, 5 September 2026; 142 matched, 12 assigned by ZIP and flagged MANUAL in the file) and `Hyperscaler_Campus_Locations.csv`. Output `lm2_local_metro_omb.csv`. Zero of 154 locals differ from the prefix crosswalk. Author’s calculations.

[41] James Feyrer, Erin T. Mansur and Bruce Sacerdote, “Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution,” *American Economic Review* 107, no. 4 (2017). Working-paper version, NBER WP 21624, abstract: “Every million dollars of oil and gas extracted produces ... 0.78 jobs within the county. ... Within 100 miles of the new production, one million dollars generates ... 2.49 jobs.” Saved at `lm2_lit/Feyrer_2017_fracking_dispersion.txt`.

[42] Stuart S. Rosenthal and William C. Strange, “How Close Is Close? The Spatial Reach of Agglomeration Economies,” *Journal of Economic Perspectives* 34, no. 3 (2020). Author manuscript dated 14 April 2020, saved at `lm2_lit/Rosenthal_2020_how_close_is_close.txt`: “agglomera-

tion effects can extend over broad distances, they attenuate, with nearby activity” exerting the strongest effect.

[43] Group C bound: no campus published by Google, Meta or Microsoft, and no Amazon announcement located, in the state. States: AK, CO, FL, HI, MA, MT, NJ, NY, WV.

[44] Version history in `LM2_Findings_CONSOLIDATED_v2_2.md`, Provenance. The superseded runs are retained as `_v2_0.md` and `_v2_1.md`.

[45] CBRE, *North America Data Center Trends*, H1 2018, H2 2025 and H1 2026 editions, cbre.com/insights/books/north-america-data-center-trends-h1-2026 and the corresponding pages for earlier editions. Transcribed in `CBRE_Market_Inventory.csv`, 45 rows. Industry-sourced.

[46] Sums of the transcribed primary-market rows in `CBRE_Market_Inventory.csv` against totals stated in the saved overview texts at note 66: H1 2026 inventory 10,902.9 against “10,903 megawatts (MW)”; H1 2026 under construction 7,481.1 against “7,481.1 MW”; H2 2025 inventory 9,432.6 against “9,432 megawatts (MW)”; H1 2018 474.3 against “more than 474 MW of capacity under development within U.S. primary markets.” Author’s transcription.

[47] See note 66.

[48] `lm2_v2_1_results.json`. Author’s calculations.

[49] `lm2_referee_checks.py`, written after a referee pass on 5 September 2026 and importing the v2.2 sample; output `lm2_referee_checks.json`. Carries the interaction with state fixed effects, the Wald test of treated minus Group B, the non-Inside three-group medians, the fab and battery metro cut, the cut of metros whose sites are all dated and all 2025 or later, aggregate membership, weighted least squares, leave-one-out, and randomization inference (5,000 draws, seed 1). The fab and battery metros are named from public knowledge of the TSMC, Intel, Micron and Panasonic sites; no project list was retrieved. Author’s calculations.

[50] `lm2_v2_1_results.json`, Group A. “Operators publishing” is the count of the four operators with at least one published site in the metro.

[51] 36 distinct states among the 89 Inside locals. Author’s count.

[52] `lm2_identification_results.json`. Table 3 rows correspond to runs R3, R4, R4 (second coefficient), R1 and R2 in that file. Estimated with `lm2_identification.py`, which imports the sample from `lm2_v2_1.py`. Author’s calculations.

[53] `lm2_identification_results.json`, R1: coefficient on `inside` +9.97, SE 4.28, $p = 0.020$, which in an interaction model is the Inside premium among untreated locals. The premium among treated locals is that coefficient plus the interaction, 29.65, SE 10.45, $p = 0.005$ on a linear-combination test of the same fitted model; that test is run by `lm2_referee_checks.py` and reported as `inside_among_treated` in `lm2_referee_checks.json`. Author’s calculations.

[54] `lm2_identification_results.json`, event. Balanced panel: every Inside local with a non-missing membership figure in all eight filing years. $89 \times 8 = 712$.

[55] Greenstone, Hornbeck and Moretti, WP 13833, p. 31: “Importantly, there is not evidence of differential trends in the years before the MDP’s opening and statistical tests confirm this visual impression.”

[56] `lm2_v2_1_results.json`, timing table. Cumulative medians from a 2018 base on the balanced panel.

[57] `lm2_identification_results.json`, robust. Median regression is least absolute deviations without clustering; its interval is the default from `statsmodels.quantreg`.

[58] `lm2_v2_1_results.json`, not-Inside table.

[59] Table 3, row 4.

[60] The 23 non-Inside locals in treated metros, with growth: Grove City OH +48.18, Collegeville PA +46.38, Gilbert AZ +44.09, Southfield MI +41.26, Batavia IL +36.91, Indianapolis 1393 +30.92, Minneapolis 160 +29.95, Forestville MD +27.24, Smyrna GA +22.16, Newtown Square PA +17.66, Kansas City 53 +13.35, Tinley Park IL +12.89, Aurora IL +9.94, Des Moines 55 +5.63, Downers Grove 15 -4.54, Phoenix 266 -5.73, Las Vegas 396 -6.46, Burnsville MN -6.73, Highland Springs 50 -14.98, Phoenix 387 -16.24, Largo MD -21.30, Cedar Rapids 204 -22.23, Downers Grove 21 -51.54.

[61] Ines Helm, “National Industry Trade Shocks, Local Labour Markets, and Agglomeration Spillovers,” *Review of Economic Studies* 87, no. 3 (2020). Working-paper version dated 10 January 2019, abstract: “I find considerable employment spillovers from other tradable industries’ trade shocks and even stronger effects within the same broad sector. Spillovers are larger for industries employing similar workers.” Saved at `lm2_lit/Helm_2019_trade_shocks_agglomeration.txt`; tables beyond Table 1 are not in the saved file.

[62] U.S. Bureau of Labor Statistics, Quarterly Census of Employment and Wages, Open Data Access industry slices, annual averages, private ownership: `data.bls.gov/cew/data/api/{2018,2024,2025}/a/ind` retrieved 5 September 2026; server last-modified dates 6 September 2019 (2018), 2 September 2025 (2024) and 21 August 2026 (2025). Rows for the 481 counties and 73 statistical areas holding a sample Inside local are saved verbatim at `qcew_src/QCEW_23821_238212_extract_2026-09-05.txt`. NAICS 238210 does not exist as a QCEW file; 23821 is the five-digit total and 238212 the nonresidential six-digit code.

[63] `lm2_qcew_benchmark.py`, reading the extract at the previous note, `lm2_local_metro_omb.csv`, `lm2_qcew_targets.json` and the v2.2 sample; output `lm2_qcew_benchmark.csv` and `.json`. Employment is summed over the counties of each statistical area disclosed in both years; a local is used if the 2018 sum is at least 100. Standard errors clustered by state. Author’s calculations.

[64] `CBRE_Market_Inventory.csv`, H2 2025 rows, joined to Group A by the market crosswalk in `lm2_v2_1.py`.

[65] Pearson r over twelve locals, author’s calculation, `lm2_v2_1_results.json`.

[66] CBRE, *North America Data Center Trends H2 2025*, Figure 1 heading and rate bullet; *H1 2026*, Figure 1 heading; *U.S. Data Center Trends Report H1 2018*, Figure 13 “Market Definitions.” Saved as `cbre_src/CBRE_NADC_Trends_H2_2025.txt`, `_H1_2026.txt` and `_H1_2018.txt`, retrieved 5 September 2026. The 2025 and 2026 files are the overview pages, whose figure tables are chart embeds and render as images rather than text, so the per-market values in `CBRE_Market_Inventory.csv` are not in the saved text; the per-market chapter pages were not retrieved. The H2 2025 figures were read from the published page on 5 September 2026 and the inventory and vacancy columns verified for all sixteen markets; see the provenance note at the end of the paper. The 2018 file is from a Texas A&M Real Estate Research Center mirror.

[67] CBRE editions held: H1 2018, H2 2025, H1 2026. Augusta does not appear as a market in any of the three. Editions between H1 2018 and H2 2025 were not examined.

[68] OLMS reporting history for file 012-627, held at `OLMS_LM2_IBEW_Local26_012-627_ReportingHistory_2000-` retrieved 4 September 2026.

[69] IBEW Local 26 FY2025 Form LM-2, RPT_ID 924002, period 1 July 2024 to 30 June 2025, Schedule 13, from the 2025 `ar_membership` extract and the filed PDF at note 5.

[70] Schedule 13 apprentice category by fiscal year from the 2018 through 2025 extracts, with category labels matched across spelling variants. `LM2_IBEW_Local26_Membership_Series_2026-09-04.md`. Author’s calculation.

[71] Slaiman, 6 July 2026, [01:00:24]–[01:00:48], as transcribed: “the 14 million hours we used to work 8 years ago. My local union works 28 million hours last year. We’re going to work north of 33 million hours. We have raised uh we went from uh 9,000 work uh members to 17,000 in that time period.” `Labor_Report_Source_Register_v0_1.md` C2. Advocate-sourced claim.

[72] Reporting history at note 68: most recent report FY2025, received 26 September 2025; no FY2026 report listed as of 4 September 2026.

[73] Slaiman, Loudoun County Planning Commission, Spring Valley hearing, 23 July 2026, § 3.4.49 beginning [04:50:47], quotation at [04:51:15]–[04:52:50], as transcribed: “My local union doesn’t really need this - we have 97 [percent] of the work in the data center[s].” `Loudoun_Extract_Case_Hearings_v1_0.md` § 3.4.49; `Labor_Report_Source_Register_v0_1.md` C1. Advocate-sourced claim.

[74] Individual checks of 4 September 2026 as recorded in `LM2_Findings_CONSOLIDATED_v2_2.md` § 6, with the pages examined listed there.

[75] Fiscal-year-end month from `lm_data`, 2025 extract, across the 154 sample locals. Author’s count.

[76] Patrick Kline and Enrico Moretti, “Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority,” *Quarterly Journal of Economics* 129, no. 1 (2014). Working-paper version, NBER WP 19293, abstract: “Gains in manufacturing employment, by contrast, continued to intensify well after federal transfers had lapsed” and “agglomeration gains in the TVA region were offset by losses in the rest of the country.” Saved at `lm2_lit/Kline_2014_TVA_big_push.txt`; tables beyond Table I are not in the saved file.

[77] Timothy J. 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; published in *Journal of Economic Perspectives* 34, no. 3 (2020). Working-paper text saved at `lm2_lit/Bartik_2019_PlaceBasedWP.txt`: “Migration shocks to local labor supply induce shocks to local labor demand of similar or larger size” and “Migrants demand housing, affecting construction.”

[78] Edward C. Kokkelenberg and Donna R. Sockell, “Union Membership in the United States, 1973–1981,” *Industrial and Labor Relations Review* 38, no. 4 (1985): 497–543. Abstract only, `lm2_lit/Kokkelenberg_1985_union_membership_ABSTRACT_ONLY.txt`: estimates “by occupation, industry, selected states and SMSAs.” The claim that it is the last such published series is the author’s and is bounded to the sources consulted for this paper.

[79] Hunt Allcott and Daniel Keniston, “Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America,” *Review of Economic Studies* 85, no. 2 (2018). Working-paper version, NBER WP 20508 revised June 2017, abstract: “Local wages rise during oil and gas booms, but manufacturing is not crowded out ... Tradable manufacturing subsectors do contract during resource booms, but their productivity is unaffected, so there is no evidence of foregone local learning-by-doing effects.” Saved at `lm2_lit/Allcott_2017_dutch_disease_agglomeration.txt`; tables are not in the saved file.

[80] Vicki Mayer and Julia Velkova, “This site is a dead end? Employment uncertainties and labor in data centers,” *The Information Society* 39, no. 2 (2023): 112–122, interviews and observation at trade expos and at a data center in Finland. Saved at `lm2_lit/Mayer_2023_DeadEnd.txt`. The characterisation of the field as qualitative and operations-side is bounded to the works retrieved for this paper, fifteen files in `lm2_lit/`.

Provenance and open items

Every citation in this paper rests on a document held on disk, with one class of exception: the operator site pages behind `Hyperscaler_Campus_Locations.csv`, at notes [36], [37] and [74], were read and transcribed into that file rather than saved as page text, so the file is the author’s transcription and the pages are not in the deposit. The four citations that once rested on general knowledge, notes [13], [14], [24] and [30], were retrieved on 5 September 2026 and are saved in `cites_src/`. The OLMS extracts, the IBEW map text items, the LM-2 instructions, the QCEW slices, the CBRE overview texts, the literature and the Census Geocoder responses in full are saved under the directories the notes name; `lm2_deposit_build.py` assembles them and `_zenodo_lm2/README.md` states the order in which the scripts reproduce every table.

Four limits on that apparatus are stated here rather than in a note. The published-version page numbers for Greenstone, Feyrer, Kline and Moretti, Allcott and Keniston, Helm, Ellison and Glaeser, and Rosenthal and Strange are not in the saved working-paper texts, and those works are cited to the versions held. The CBRE report is a copyrighted industry publication and cannot be deposited. Its H2 2025 figures were read from the published page on 5 September 2026 and checked against `CBRE_Market_Inventory.csv`, twice for the two market tables: inventory and vacancy match for all sixteen markets on both an image read and an independent transcription, and the eight primary inventories sum to the 9,432 the page states. Under construction is published as a bar chart with the value printed above each bar; all sixteen match the file market by market, and the eight primary values sum to the 5,994.4 megawatts the page’s text states. Fifteen of those sixteen were read from images held at `cbre_src/verification_2026-09-05/`; Austin-San Antonio’s label falls outside the crop and was read from the live chart. Those images and the transcription beside them are excluded from the deposit because the report is copyrighted. The rendered-map images used to validate four state rosters were not kept. And the classification history in section 3.3 is recorded in `LM2_Findings_CONSOLIDATED_v2_2.md`, which is in the deposit.