Connected Production: Rural Broadband, Commuting Networks, and the Location of Manufacturing Output*

Junho Choi
Columbia University

November 2026

Abstract

Governments now treat high-speed broadband as essential infrastructure, yet the economic effects of wiring rural places remain poorly understood, in part because connectivity in one place can change economic activity in the places it is linked to. I study the Federal Communications Commission's Rural Digital Opportunity Fund (RDOF), a 2020 reverse auction that awarded $9.2 billion to deploy mostly gigabit-capable service to about 5.2 million unserved locations. Identification uses the timing of deployment within pre-determined recentering cells, following the design-based approach of Borusyak and Hull (2023), and extends it to exposure through pre-period commuting networks. Deployment raises measured speeds and high-speed subscription. A county's own deployment has no detectable effect on its GDP, employment, or manufacturing output. Deployment among its commuting partners, by contrast, raises its real manufacturing GDP (coefficient 2.22, standard error 0.74): a one standard deviation increase in connected exposure, within county and year, raises manufacturing GDP by about 1 percent, and the implied gain in 2024 is about $9.7 billion, or 0.45 percent of baseline. Total GDP and manufacturing employment show no detectable change, and about 72 percent of the gain is in non-labor value added. Announced clean-manufacturing investment does not explain the result. In the main specification, a one standard deviation increase in connected exposure also lowers residential electricity sales by about 0.3 percent and raises NO_2 by about 0.7 percent, likely from vehicle emissions rather than manufacturing, while industrial electricity sales and data center activations do not rise. Broadband in the periphery appears to change where industrial output occurs more than how much economic activity takes place overall.

JEL: L96, O33, R12, R11, Q41, Q56. **Keywords:** broadband, spatial spillovers, recentered instruments, manufacturing, commuting networks, electricity use.

*Draft prepared for the NBER digital economics conference. Preliminary and incomplete; please do not cite. I thank Doug Almond, Rob Metcalfe, Harrison Hong, Matthew Gomez, Jeffrey Shrader, José Scheinkman, and participants of the Sustainable Development Colloquium and the Financial Economics Colloquium at Columbia University for helpful comments.

1 Introduction

High-speed internet access has moved from a convenience to something closer to a public utility. Federal and state governments have committed tens of billions of dollars to close rural connectivity gaps, on the premise that broadband will raise local employment, attract firms (Kolko, 2012; Whitacre et al., 2014; Rupasingha et al., 2024), and, through remote work and online substitution, lower the energy and emissions intensity of economic activity (Lange et al., 2020).

An older tradition in the economics of innovation suggests a different expectation. Information and communication technology (ICT) is a general purpose technology whose returns depend on complementary investments in physical capital, organization, and skills (Bresnahan and Trajtenberg, 1995; Brynjolfsson and Hitt, 2000; Brynjolfsson et al., 2021). In manufacturing specifically, network connectivity is embedded in computer-controlled machinery, inspection equipment, and process control (Bartel et al., 2007), and the returns to advanced internet use have been concentrated in places that already had the relevant industrial and skill base (Forman et al., 2012). Under this view, connectivity complements physical production rather than replacing it, and its returns need not appear where the cable is laid. If rural broadband mainly integrates the periphery with the places its residents and firms already trade and commute with, the effects of a rural program will show up in those linked places. A secondary question this paper examines is whether connectivity and electricity act as complements in manufacturing: if networked machinery raises the return to energy-using capital, broadband could raise electricity use instead of lowering it.

This paper asks where the economic effects of rural broadband appear, and in particular whether connecting rural areas changes industrial output in the places linked to them. The setting is the Rural Digital Opportunity Fund Phase I (RDOF), a reverse auction the Federal Communications Commission (FCC) concluded in November 2020. RDOF awarded $9.2 billion over ten years to 180 winning bidders to serve roughly 5.2 million unserved locations; the FCC reported that 99.7 percent of these locations were assigned speeds of at least 100/20 Mbps and more than 85 percent were assigned gigabit service (Federal Communications Commission, 2020). Recipients must reach all assigned locations within six years under interim milestones, which produces staggered deployment between 2021 and the present.

Evaluating such a program raises two problems at once. The first is that deployment is not random. Providers choose which obligated areas to build first, and these choices may track local demand, terrain, or proximity to existing networks. The second is that each county cannot be considered in isolation. A county's residents commute to other counties, and its firms buy from and sell to them. If broadband in one county changes activity in the counties linked to it, a comparison of treated and untreated counties confounds the direct effect with the spillover, and the stable unit treatment value assumption fails.

I address both problems with a single design. Following Borusyak and Hull (2023), I treat the timing of deployment within narrowly defined recentering cells as the source of identifying variation. For each county and month, I compute the share of the county's broadband serviceable locations covered by RDOF deployment, and subtract its expected value over permutations of de-

ployment paths within cells. The resulting recentered instrument isolates the part of deployment that is unexpected given the program's structure. Because connected exposure, the commuting-weighted sum of other counties' own exposure, is a linear function of the same deployment paths, the same recentering applies to it directly. This yields three instruments: own exposure, the average of residence-side and workplace-side connected exposure, and their difference. The design estimates own and connected effects jointly, so each is identified net of the other.

Three findings emerge from the main specification. Own and connected exposure are not two groups of counties: a county with its own deployment can also have commuting partners with deployment, and the design separates the two effects within the same county.

First, RDOF deployment improves connectivity, both through a county's own deployment and through its partners'. Own exposure raises measured download and upload speeds in both M-Lab and Ookla speed-test data and raises the share of households with a high-speed subscription in the American Community Survey. Connected exposure also raises high-speed subscription and measured speeds. This matters for interpretation: connected exposure captures not only commuting-network spillovers but also shared infrastructure, such as middle-mile fiber and provider follow-on investment, in places adjacent to RDOF builds. Descriptively, speeds rose throughout Mississippi and Arkansas during the build-out: from 2019 to 2024, the median block group's fourth-quarter Ookla download speed rose from 88 to 290 Mbps and from 92 to 285 Mbps in commuting-linked non-RDOF block groups, and from 42 to 260 Mbps and from 38 to 231 Mbps in RDOF block groups (Appendix Figures 9 and 10). Nearly all non-RDOF block groups in these states are commuting-linked, so these numbers illustrate the setting and do not show a premium for commuting-linked places.

Second, the own effect on aggregate activity is close to zero. Own exposure has no detectable effect on real GDP, private-industry GDP, payroll employment, resident employment, or County Business Patterns employment. The confidence intervals rule out effects that would be economically large relative to the observed variation in exposure.

Third, the connected effect on real manufacturing GDP is strong. The coefficient on average connected exposure is 2.22 (standard error 0.74) on log real manufacturing GDP, with a Kleibergen-Paap first-stage F of 138. A one standard deviation increase in connected exposure, within county and year, corresponds to about a 1 percent increase in manufacturing GDP. Total GDP, private GDP, construction GDP, and retail GDP do not respond, and the own-exposure coefficient on manufacturing is small and insignificant at the 5 percent level. Applying the estimates county by county implies an increase in 2024 manufacturing GDP of about $9.7 billion, or 0.45 percent of baseline. QCEW manufacturing employment does not respond (0.07, standard error 0.55), and its confidence interval excludes the GDP coefficient. Splitting BEA manufacturing GDP directly into employee compensation and non-labor value added shows that 72 percent of the dollar-scaled gain is in the non-labor part, which is the only component that responds significantly in logs (2.70, standard error 0.83) and is also the part BEA estimates least directly. Taken together, the own effect, which would capture channels such as a county's own consumption and use of the new service, is weak, while the connected effect, which I measure through commuting partners but

which could reflect other links between places, is strong.

Because these connected effects are new and their mechanism is not yet pinned down, a large part of the paper examines what could produce them besides broadband itself. The manufacturing estimate is robust to manufacturing-share-by-year controls, Census division-by-year fixed effects, and dropping any one state, and announced clean-manufacturing investment does not explain it. It remains positive, though less precisely estimated, with state-by-year fixed effects (1.61) and commuting zone-by-year fixed effects (1.35), which absorb about 26 and 69 percent of the variance of the connected instrument. I discuss three threats in detail: a regional manufacturing investment wave that coincided with the build-out, measurement features of how the Bureau of Economic Analysis (BEA) allocates county manufacturing GDP, and changes in the coverage of commercial data sources with connectivity.

Related literature. This paper contributes to three literatures. First, it adds to work on the economic returns to broadband and ICT. Earlier U.S. studies examine the expansion of DSL and cable broadband in the 2000s (Kolko, 2012; Atasoy, 2013; Whitacre et al., 2014; Dettling et al., 2018), and international studies exploit the staggered arrival of fast internet (Akerman et al., 2015; Hjort and Poulsen, 2019) or broadband enablement of local telephone exchanges (DeStefano et al., 2018). Much of this work estimates effects in the places that receive access. I study a recent, largely gigabit-tier program and show that its measurable output effects run through a county's connected exposure rather than its own exposure. Forman et al. (2012) find that the wage returns to advanced internet investment were concentrated in counties with pre-existing advantages; here too the returns are not evenly shared, although the proportional manufacturing response to connected exposure is largest where manufacturing is a small share of the economy (Table 9). DeStefano et al. (2018) find that broadband-enabled communication ICT raised firm scale (sales and employment) rather than productivity. The margin here differs: with connected exposure, manufacturing value added rises without a detectable change in headcount, and most of the gain is non-labor. The settings also differ, since that paper studies firms in newly enabled areas and this one separates own exposure from exposure through commuting partners.

Second, the paper contributes to the design-based literature on spillovers and network exposure. Borusyak and Hull (2023) show how to recenter treatments that are functions of exogenous shocks and non-random exposure, including market-access-type treatments that depend on a network. I apply that framework to a staggered infrastructure program with commuting-weighted exposure and show that own and connected effects can be separately identified from the same deployment paths. The results illustrate that evaluations which ignore network exposure can miss most of a program's measurable effect. The commuting linkages I use are motivated by quantitative spatial models in which local shocks propagate through commuting (Monte et al., 2018), and the setting relates to work on how infrastructure that links peripheral and core regions reallocates industrial activity (Faber, 2014).

Third, the paper speaks to the environmental economics of digitalization. A literature asks whether ICT reduces energy demand on net (Lange et al., 2020), and decomposes the environmen-

tal consequences of economic change into scale, composition, and technique effects (Copeland and Taylor, 2004). Evidence on the energy-ICT relationship is mixed. In French services, Collard et al. (2005) find that electricity intensity rises with computers and software but falls with communication equipment, and at the aggregate level energy appears to be a strong short-run complement to capital and labor (Hassler et al., 2021). By documenting that rural broadband raises manufacturing output through connected exposure, I provide evidence on the composition and location channels through which digital infrastructure could alter the spatial distribution of energy use and emissions. Section 8 reports first evidence on electricity and air quality; a fuller treatment of the environmental consequences is the subject of companion work.

The rest of the paper proceeds as follows. Section 2 describes RDOF. Section 3 describes the data. Section 4 presents the identification strategy. Section 5 presents the main results, Section 6 examines manufacturing in detail, Section 7 discusses threats to identification, Section 8 reports electricity, air quality, and outages, Section 9 offers a brief interpretive framework, and Section 10 concludes.

2 Institutional Background

The auction. RDOF Phase I (FCC Auction 904) was a multi-round descending clock auction for universal service support. Areas eligible for bidding were census blocks the FCC classified as lacking 25/3 Mbps service, aggregated into census block groups (CBGs) for bidding. Bidders committed to serve an area at a given performance tier and latency in exchange for a support amount, and bids for higher speeds and lower latency were favored through tier and latency weights. Bidding closed on November 25, 2020. The auction assigned $9.2 billion over ten years to 180 bidders for about 5.22 million locations, against reserve prices of roughly $26 billion (Federal Communications Commission, 2020). Winners included cable operators, electric cooperatives and consortia of cooperatives, incumbent telephone companies, fixed wireless providers, and satellite providers.

Authorization and defaults. Winning bids became binding obligations only after a long-form review, and authorizations were granted in waves beginning in 2021. A substantial share of winning bids were defaulted or rejected during or after review, and about 3.3 million of the 5.2 million assigned locations remain obligated after defaults. I define obligations at the authorization vintage and exclude defaulted and denied areas from obligated counts; the staggered authorization dates define *authorization cohorts*, which enter the specification through cohort-specific time controls.

Deployment obligations. Support recipients must complete deployment to 40 percent of their required locations by the end of the third year, to 60 percent by the end of the fourth, to 80 percent by the end of the fifth, and to all required locations by the end of the sixth (47 CFR § 54.802(c)(1)). Compliance binds at the recipient-by-state level, which leaves providers discretion over the order in which obligated CBGs within a state are built. This within-obligation discretion over timing

is the variation the design exploits. Recipients report deployed locations to the Universal Service Administrative Company (USAC) through the High Cost Universal Broadband (HUBB) portal, which provides the deployment data used here, and certify that they have met each milestone by March 1 of the following year (47 CFR § 54.316(b)(5)).

Why RDOF is useful for identification. Two features of the program help. First, eligibility and award amounts were set by a national auction based on pre-auction coverage maps, so the set of obligated areas and their obligation sizes are pre-determined relative to post-2020 local conditions. Second, the auction was not designed around environmental or industrial outcomes, so there is no mechanical link between eligibility and the outcomes I study. The remaining concern is the timing of deployment within obligations, which the recentering in Section 4 addresses directly.

3 Data

Deployment and exposure. Obligated and deployed locations come from USAC HUBB and FCC auction records, recorded by census block and aggregated to the county by month. The denominator for all exposure measures is the count of broadband serviceable locations (BSLs) in the FCC Broadband Data Collection Fabric. Because no pre-auction Fabric vintage exists and the FCC operates an RDOF-specific location adjustment process, the BSL denominator is not strictly pre-determined; I discuss this in Section 7. The counts come from the FCC Broadband Data Collection layer published on Esri's ArcGIS Living Atlas (the December 2024 service, with an item page dated December 2025), and access to the Fabric itself is being negotiated.

Commuting networks. Commuting weights come from the LEHD Origin-Destination Employment Statistics (LODES8, 2019, primary jobs), aggregated from block pairs to county-to-county flows and held fixed.

Outcomes. The main outcome for this paper is county real GDP by industry from the BEA (CAGDP9, chained 2017 dollars; current-dollar GDP from CAGDP2), annually for 2018 to 2024. The manufacturing series is suppressed for some county-years; the manufacturing sample has 19,043 county-years against 21,349 for total GDP. Labor market outcomes come from the Quarterly Census of Employment and Wages (QCEW), Local Area Unemployment Statistics (LAUS), Quarterly Workforce Indicators (QWI), and County Business Patterns (CBP). Establishment counts and sizes come from Data Axle historical business files (2019 to 2024). Connectivity outcomes come from M-Lab and Ookla speed tests and from ACS household internet subscription measures. Electricity sales by customer class come from EIA Form 861 at the utility-by-state level (Section 8). Air quality comes from TROPOMI NO_2 columns and ACAG $PM_{2.5}$, with wildfire smoke days from NOAA Hazard Mapping System plumes. Outages come from Eagle-I, data center activations from the Aterio data center inventory (Aterio, 2026), and announced clean-manufacturing investment from the Clean Investment Monitor (Rhodium Group and MIT CEEPR, 2026).

Table 1: Distribution of Exposure Measures

	Mean	SD (raw)	SD (within county)	p99	Share nonzero
Own exposure, S^{own}	0.0063	0.034	0.026	0.169	0.17
Connected, average, S^{A}	0.0019	0.0062	0.0044	0.032	0.54
Connected, difference, S^{D}	−0.0004	0.0026	0.0021	0.0043	0.54
Recentered Z^{own}	0.0003	0.024	0.019	0.083	0.38
Recentered Z^{A}	0.0002	0.0031	0.0024	0.014	0.57

Notes: County-year panel, manufacturing GDP sample, 2018 to 2024. Exposure is the BSL-normalized share of locations covered by RDOF deployment. $S^{A} = (S^{\text{res}} + S^{\text{work}})/2$ and $S^{D} = S^{\text{res}} - S^{\text{work}}$. Statistics include the zero-exposure years 2018 to 2020. Within SDs net out county and year effects; with the full controls of the main specification, the within SD of S^{A} is 0.0039.

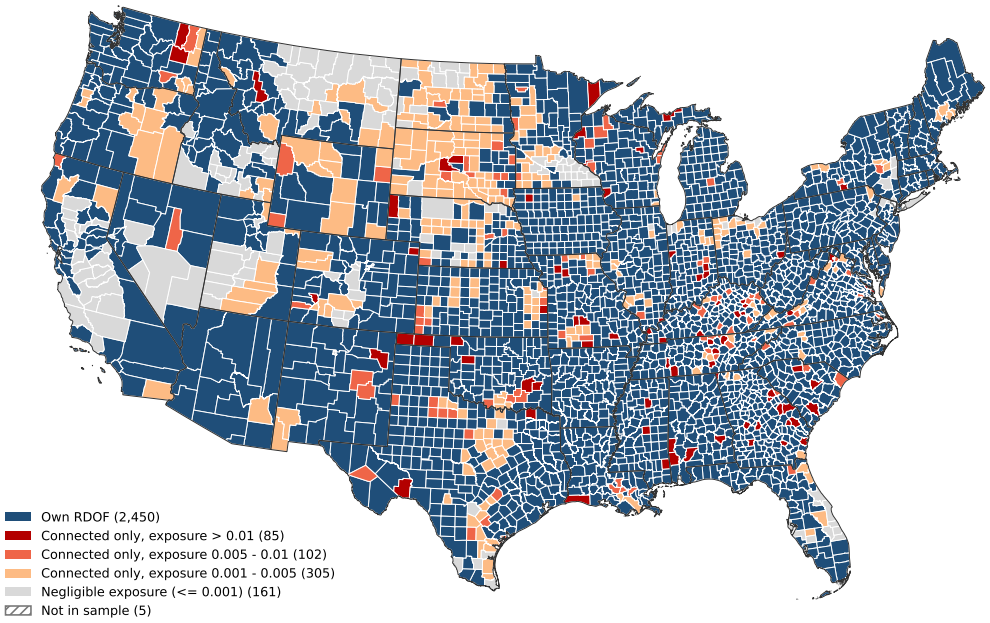
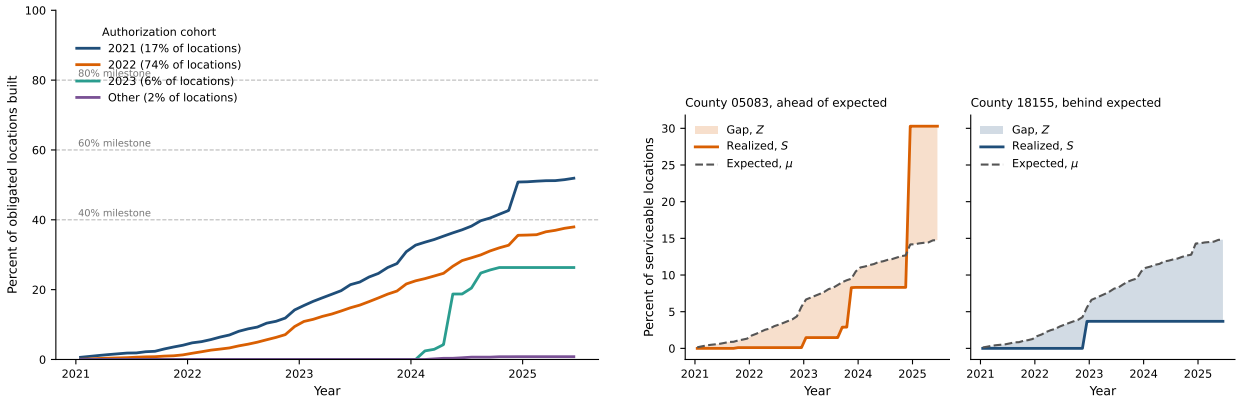


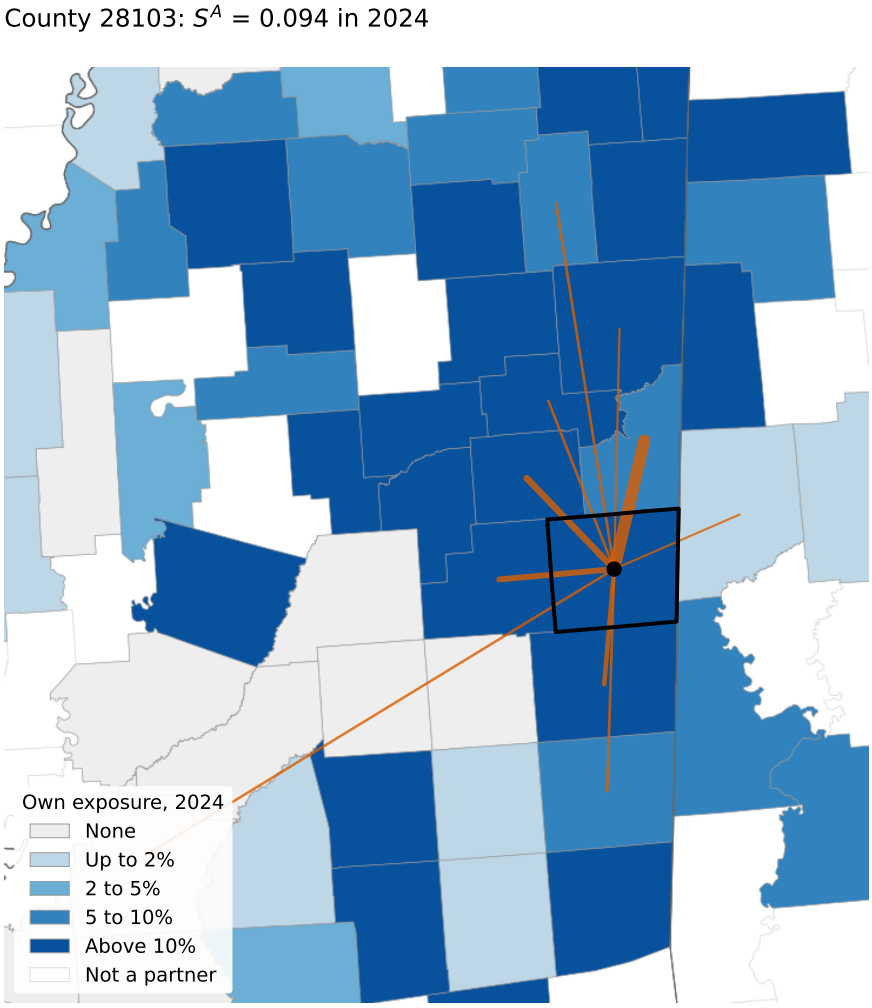
Figure 1: Counties with own RDOF obligations and connected exposure

Notes: Own RDOF counties have obligated locations. Connected-only counties have none, and are classified by their largest connected exposure in any month.

Sample and timing. Exposure is measured monthly from 2021 onward. For annual outcomes, I aggregate exposure to the county-year and include 2018 to 2020 as pre-period years at zero exposure. Table 1 reports the distribution of the exposure measures, and Figures 1 to 3 show where and when deployment and exposure occur. Connected exposure varies much less than own exposure: the raw standard deviation of average connected exposure is 0.0062, against 0.034 for own exposure, and its standard deviation within county and year is about 0.0044. All coefficients below are per unit of exposure (a move from zero to full coverage), so I translate them into effects of realistic changes throughout.



(a) **(b)**



(c)

Figure 2: The RDOF rollout and the identifying variation

Notes: (a) Share of obligated locations built, by authorization cohort; dashed lines mark milestone levels, not dates. (b) Realized and expected own exposure for two counties with similar expected paths, one ahead of expectation and one behind, chosen by rule for illustration; the shaded gap is the recentered instrument. (c) One Mississippi county (outlined) and its commuting partners, shaded by 2024 own exposure, with lines for its ten largest commuting weights. The county's connected exposure is the sum over partners of the commuting weight times the partner's own exposure.

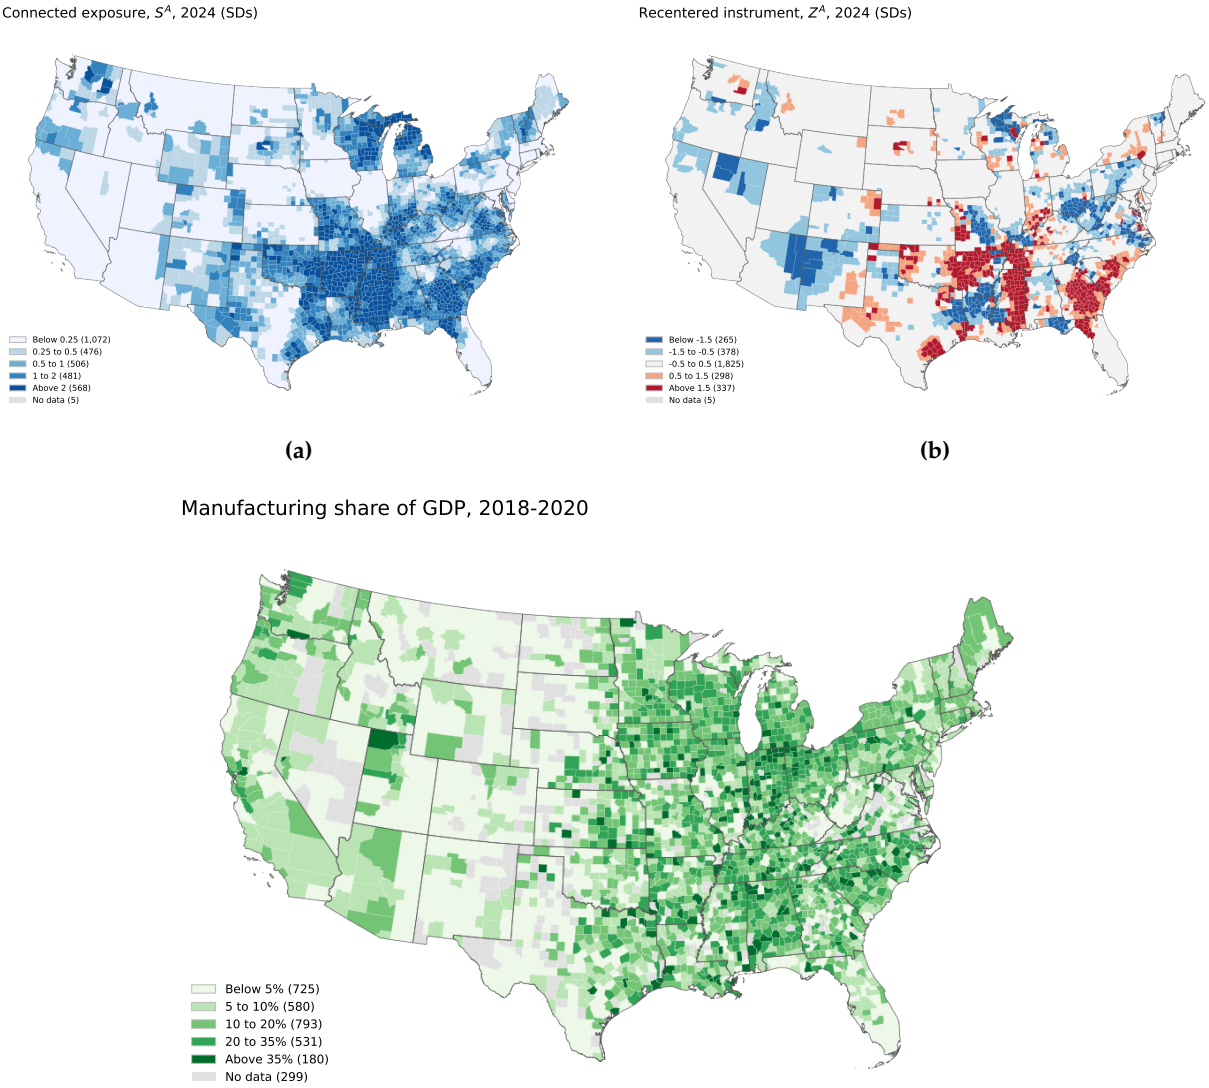


Figure 3: Connected exposure, the recentered instrument, and baseline manufacturing

Notes: Means over 2024. S^A and Z^A are each divided by their own standard deviation across county-years in an annual panel like the GDP regressions (2018 to 2024, zero in 2018 to 2020), without demeaning. A county at 2 on panel (a) has an implied manufacturing effect of about 2×1.38 log points. Manufacturing share is BEA current-dollar manufacturing GDP over total GDP, 2018 to 2020.

4 Identification Strategy

4.1 Own exposure

Index census blocks by k, counties by i and j, and months by t. Let n_k be the number of RDOF-obligated locations in block k, N_i the number of BSLs in county i, and $g_{kt} \in [0, 1]$ the share of block k's obligated locations deployed by month t, weakly increasing in t. County i's own exposure is

the share of its serviceable locations covered by RDOF deployment:

$$S_{it}^{\text{own}} = \frac{1}{N_i} \sum_{k \in i} n_k \, g_{kt}. \tag{1}$$

The deployment path $g_k = (g_{kt})_t$ is the shock. Obligation sizes n_k, BSL counts N_i, the commuting network, and the cell assignment described below are treated as fixed and collected in w.

Recentering. Each obligated block belongs to a recentering cell $c(k)$, defined by four pre-determined characteristics: the FCC performance tier; the capacity of the parent company (a size bin for small and medium firms, the firm itself for large ones); the within-state quartile of reserve price per location; and December 2019 connectivity from Form 477. Baseline cells are national: they do not condition on state. The identifying assumption, stated formally below, is that within a cell the timing of deployment is exchangeable across obligated locations, conditional on w.

Under this assumption, the expected deployment rate of block k is the cell's location-weighted pooled rate, $\bar{g}_{ct} = \sum_{k \in c} n_k g_{kt} / \sum_{k \in c} n_k$, computed over all live blocks in the cell including k itself (the permutation construction, PERM). A block that defaults after authorization leaves the pool from its default month, keeps its realized deployment in S^{own}, and takes its own frozen rate as its expected rate, so it contributes nothing further to the instrument. Expected own exposure is

$$\mu_{it}^{\text{own}} = \mathbb{E}[S_{it}^{\text{own}} \mid w] = \frac{1}{N_i} \sum_{k \in i} n_k \, \bar{g}_{c(k),t}, \tag{2}$$

and the recentered instrument is

$$Z_{it}^{\text{own}} = S_{it}^{\text{own}} - \mu_{it}^{\text{own}}. \tag{3}$$

The recentering removes the part of exposure that is predictable from the program's structure. A county with many obligated locations in fast-building cells has high expected exposure regardless of which particular blocks were built first; if such counties differ systematically in their outcome trends, S^{own} is confounded but Z^{own} is not. The county fixed effect does not solve this, because expected exposure varies over time differentially across counties. Borusyak and Hull (2023) show that controlling for or subtracting this expected value is what makes a treatment built from exogenous shocks and non-random exposure usable as an instrument. By construction, the location-weighted deviation $\sum_{k \in c} n_k (g_{kt} - \bar{g}_{ct})$ over live blocks is zero within each cell and month.

4.2 Connected exposure

Let W^{res} be the county-by-county matrix whose (i, j) entry is the share of county i's employed residents who work in county j, and let W^{work} be the matrix whose (i, j) entry is the share of workers employed in county i who live in county j. Both are computed from 2019 LODES flows with own-county workers in the denominator. The diagonal is then set to zero, rows are not

renormalized, and entries below 0.001 are dropped. Each row therefore sums to about the county's out-commuting share rather than to one, so connected exposure and its instrument both scale with how open a county's labor market is. Connected exposures are

$$S_{it}^{\text{res}} = \sum_{j \neq i} W_{ij}^{\text{res}} S_{jt}^{\text{own}}, \qquad S_{it}^{\text{work}} = \sum_{j \neq i} W_{ij}^{\text{work}} S_{jt}^{\text{own}}. \tag{4}$$

S^{res} is high when the places a county's residents work in are being built out; S^{work} is high when the places a county's workers live in are being built out.

Connected exposure is a treatment of the general network form considered in Borusyak and Hull (2023): it depends on the deployment paths of many blocks outside county i. Because it is linear in those paths given the fixed network, its expected value under the same permutation design is $\mu_{it}^{m} = \sum_j W_{ij}^{m} \mu_{jt}^{\text{own}}$ for $m \in \{\text{res}, \text{work}\}$, and the recentered instrument is $Z_{it}^{m} = \sum_j W_{ij}^{m} Z_{jt}^{\text{own}}$. No additional assumptions are required beyond those for own exposure and the pre-determination of W.

Reparameterization. Because S^{res} and S^{work} are highly correlated, separate coefficients on them are imprecise and their sampling errors are strongly negatively correlated. I therefore estimate on

$$S_{it}^{A} = \tfrac{1}{2}\left(S_{it}^{\text{res}} + S_{it}^{\text{work}}\right), \qquad S_{it}^{D} = S_{it}^{\text{res}} - S_{it}^{\text{work}}. \tag{5}$$

This is an exact reparameterization. Writing $\delta_r S^{\text{res}} + \delta_w S^{\text{work}} = \kappa S^{A} + \tau S^{D}$ gives $\kappa = \delta_r + \delta_w$ and $\tau = \frac{1}{2}(\delta_r - \delta_w)$, and conversely $\delta_r = \kappa/2 + \tau$ and $\delta_w = \kappa/2 - \tau$. The coefficient κ on the *average* is the *sum* of the two side-specific effects: it is the effect of raising both connected exposures by one unit. The shared variation in the two exposures loads on S^{A} and is estimated precisely; the small asymmetric component loads on S^{D}, which has a raw standard deviation of only 0.0026 and is correspondingly imprecise. The own-exposure coefficient, fitted values, and residuals are identical under either parameterization.

4.3 Estimating equation

The main specification is

$$y_{it} \;=\; \beta\, S_{it}^{\text{own}} + \kappa\, S_{it}^{A} + \tau\, S_{it}^{D} + X_{it}'\gamma + \alpha_i + \lambda_t + \varepsilon_{it}, \tag{6}$$

estimated by two-stage least squares with $(S^{\text{own}}, S^{A}, S^{D})$ instrumented by $(Z^{\text{own}}, Z^{A}, Z^{D})$. Table 2 lists every element of the main specification. Robustness checks change one element at a time and are reported separately (Table 10).

Why instrument rather than control. An alternative is to include S^{own}, S^{A}, and S^{D} directly and control for their expected values. I instrument with the recentered measures because it keeps the

Table 2: Main Specification

Element	Main specification
Endogenous regressors	S^{own}, S^{A}, S^{D}
Instruments	Z^{own}, Z^{A}, Z^{D} (PERM recentering within cells)
Exposure denominator	Broadband serviceable locations (FCC Fabric)
Recentering cells	Tier, parent-company capacity, within-state reserve-price quartile, Dec. 2019 connectivity (national cells)
Controls X_{it}	Weather (temperature, precipitation, heating and cooling degree days); obligation intensity by authorization cohort $\times$ time
Fixed effects	County and time (year for annual outcomes, month for monthly outcomes); utility-state and year for electricity
Sample	All counties (connected exposure is defined for counties without own obligations)
Clustering	County (about 2,900 clusters); utility-state for electricity
Weak-IV diagnostic	Kleibergen-Paap rk F for the joint model

Notes: The county is the level at which repeated observations share exposure. Recipient-by-state is not the level of assignment under the design, in which timing is exchangeable across locations within cells (Abadie et al., 2023). Neither level captures dependence across counties that share commuting partners.

estimand in units of realized exposure while using only the unexpected component for identification. Since $Z = S - \mu$, the first stage is strong by construction wherever realized exposure varies around its expectation; the joint Kleibergen-Paap F statistics are between 117 and 215 for the county outcomes in Tables 3 and 4, and between 99 and 127 for the utility-level electricity regressions. Table 10 also reports the OLS specification that includes S^{own}, S^{A}, and S^{D} directly and controls linearly for their expected values μ. That estimate is 3.42, against 2.22 when instrumenting with the recentered measures, and the difference is statistically significant ($t = 3.10$). With state-by-year fixed effects the two are closer (1.94 and 1.61), so the gap appears to come from between-state variation, the same margin where the pre-period imbalance appears (Section 7).

Identifying assumption. Let ε denote the vector of structural errors. The key condition is that deployment timing is exchangeable across obligated locations within cells, conditional on w and ε: any reassignment of build dates across locations in the same cell is equally likely. Under this condition, $\mathbb{E}[Z_{it} \mid w, \varepsilon] = 0$ for each recentered instrument, which implies the exclusion restriction needed for (6). The condition is about the *order* of construction within cells, not about which areas were eligible or how large obligations were. It would fail if, within a cell, providers systematically built first in blocks whose counties (or whose commuting partners) were on different outcome trajectories, for example blocks near growing industrial employers.

What the coefficients mean. β is the effect of a county's own deployment holding connected exposure fixed; κ is the effect of deployment among a county's commuting partners holding own exposure fixed. Neither is a total program effect. Under a national rollout, a county's own and connected exposure rise together, and the combination $\beta + \kappa$ summarizes whether the program's effect in a county is offset by reallocation from its partners. For manufacturing GDP, $\hat{\beta} + \hat{\kappa} = 2.05$

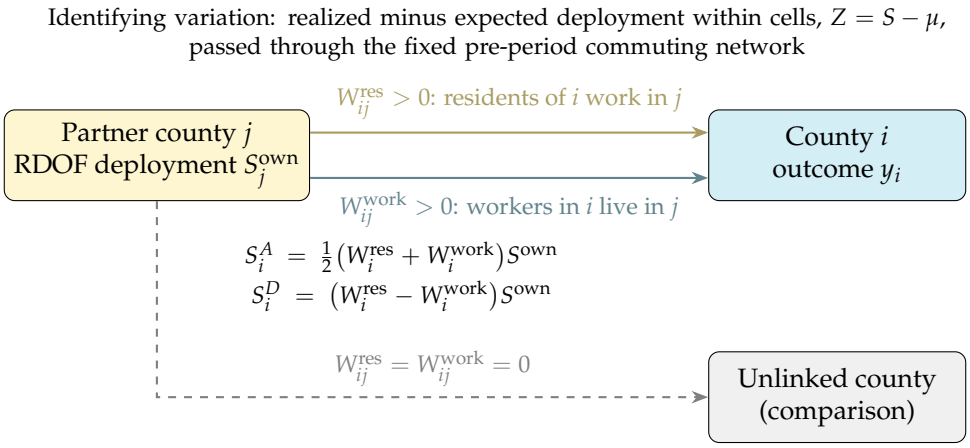


Figure 4: Own and connected exposure

Notes: Schematic. Own exposure is the share of a county's serviceable locations covered by RDOF deployment. Connected exposure is the commuting-weighted sum of other counties' own exposure, using 2019 LODES flows with the diagonal set to zero. County i can also have its own deployment. Identification compares counties whose partners were built earlier than expected with counties whose partners were built later than expected, within the same recentering cells.

(standard error 0.71).

Inference. Standard errors are clustered by county. Following Abadie et al. (2023), the choice turns on how treatment is assigned: under the design, timing is exchangeable across locations within cells, not assigned by recipient and state, and the county is the level at which repeated observations share exposure. Counties that share commuting partners also share connected exposure, so county clustering may understate uncertainty for the connected coefficients. An Anderson-Rubin joint test, which is robust to weak instruments, rejects that all three manufacturing effects are zero ($p = 0.019$). All stars use county-clustered standard errors and normal critical values; randomization inference and spatially robust standard errors are planned.

5 Main Results

5.1 RDOF deployment improves connectivity

Table 3 shows that deployment changes connectivity. Own exposure raises log download speeds by 0.95 in M-Lab and 0.75 in Ookla data per unit of exposure, and log upload speeds by 1.44 and 1.99. It raises the ACS high-speed subscription share by 0.062 per unit. These are first-order validations that the deployment data capture real changes in service.

Deployment among commuting partners also predicts better broadband: its coefficient on the high-speed subscription share is 1.18 (standard error 0.14), and its coefficients on measured speeds are positive and significant at the 5 percent level, except Ookla upload at the 10 percent level. Scaled by a one raw standard deviation change in each exposure, the connected effect on subscriptions is about three and a half times the own effect, while for measured speeds the connected

Table 3: Deployment and Connectivity (Main Specification)

	M-Lab		Ookla		ACS share	
	Log down (1)	Log up (2)	Log down (3)	Log up (4)	High-speed (5)	Any access (6)
Own, S^{own}	0.948***	1.437***	0.746***	1.990***	0.062***	0.009
	(0.220)	(0.341)	(0.171)	(0.302)	(0.019)	(0.012)
Connected avg., S^A	2.682**	6.388***	3.003**	4.071*	1.175***	0.449***
	(1.270)	(2.065)	(1.278)	(2.188)	(0.145)	(0.086)
Connected diff., S^D	0.905	−0.584	3.079	4.557	−0.072	−0.432***
	(2.302)	(3.756)	(2.111)	(3.389)	(0.258)	(0.161)
Observations	116,145	116,145	148,047	148,047	12,403	12,403
KP first-stage F	214.81	214.81	161.35	161.35	129.16	129.16

Notes: Each column is a separate 2SLS regression of the outcome on own exposure, average connected exposure, and the connected difference, instrumented by their recentered counterparts, with the main specification in Table 2. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ (normal critical values). Coefficients are per unit of exposure. M-Lab is county-month. Ookla is quarterly, with each quarter's value repeated across its months. The ACS share is county-year.

effect is 40 to 80 percent of the own effect. Neither measure is clean. Speed-test outcomes depend on who chooses to run tests, and the ACS county rates are PUMA estimates allocated to counties, so counties in a multi-county PUMA share values, which can load mechanically on connected exposure. I therefore interpret the magnitude of the connected subscription estimate with caution.[1]

This result shapes how the connected effects should be read. Connected exposure does not isolate a pure commuting spillover holding the county's own connectivity fixed. It captures the full reduced-form consequence of partners' RDOF deployment, which includes changes in the county's own service through shared middle-mile infrastructure and provider investment in adjacent areas. Separating the two channels requires controlling for non-RDOF changes in availability, which I discuss in Section 7.

5.2 Aggregate activity: a weak own effect

Table 4 reports effects on aggregate output and employment. Own exposure has no detectable effect on any measure. The own coefficients on log real GDP and log private-industry GDP are -0.017 and -0.019 with standard errors near 0.05. Scaled by one raw standard deviation of own exposure (0.034), the 95 percent confidence interval for real GDP runs from about -0.4 to $+0.3$ percent. QCEW payroll employment, LAUS resident employment, and CBP employment show the same pattern.

Connected exposure also has no detectable effect on total county GDP or payroll employment. Resident employment falls in counties whose partners are treated (LAUS, $\hat{\kappa} = -0.247$, standard error 0.105). The connected difference S^D is significant for CBP employment and marginally signifi-

[1] Future work will add test-count controls for the speed-test outcomes and re-estimate subscriptions on single-county PUMAs or a finer source.

Table 4: Aggregate Output and Employment (Main Specification)

	BEA real GDP		QCEW		LAUS	CBP
	Total	Private	Log emp.	Log wage	Log emp. (res.)	Log emp.
	(1)	(2)	(3)	(4)	(5)	(6)
Own, S^{own}	−0.017	−0.019	−0.018	0.015	−0.003	0.004
	(0.048)	(0.053)	(0.018)	(0.012)	(0.013)	(0.060)
Connected avg., S^A	0.027	0.097	−0.075	−0.081	−0.247**	−0.180
	(0.244)	(0.278)	(0.151)	(0.115)	(0.105)	(0.366)
Connected diff., S^D	−0.328	−0.340	−0.043	0.525*	−0.257	−1.867**
	(0.593)	(0.650)	(0.308)	(0.269)	(0.201)	(0.774)
Observations	21,349	21,349	167,310	167,310	148,464	18,584
KP first-stage F	116.76	116.76	119.92	119.92	161.27	146.00
Frequency	Annual	Annual	Monthly	Monthly	Monthly	Annual

Notes: See Table 3. BEA GDP in chained 2017 dollars, 2018 to 2024, with 2018 to 2020 at zero exposure. QCEW wage is the average weekly wage. QCEW employment is monthly; wages are quarterly, repeated across the quarter's months.

cant for QCEW wages; given the weak identification of S^D and the number of outcomes examined, I do not interpret these. Two sector results outside this table also stand out: QCEW construction employment and CBP accommodation and food employment rise with connected exposure (2.15, standard error 0.95, and 2.16, standard error 0.77). Table 6 scales them, and Section 9 returns to them.

6 The Connected Effect on Manufacturing Output

6.1 Main estimate

Figure 5 and Table 7 present the central result. Across the BEA industry aggregates, manufacturing is the only series that responds to connected exposure. The coefficient on S^A is 2.222 (standard error 0.740, $t = 3.00$). The own-exposure coefficient is −0.172 (0.103), and the connected difference is 1.648 (1.266). The joint first-stage F is 138.

Magnitude. Because connected exposure varies over a narrow range, the per-unit coefficient overstates what the program delivered. Table 5 translates the estimate, and Table 6 does the same for other outcomes. A one standard deviation increase in S^A within county and year (about 0.0044), which is the variation the regression uses, raises manufacturing GDP by about 0.98 percent; a one raw standard deviation increase (0.0062) raises it by about 1.4 percent. For an upper bound, if a county's partners were fully built at an average obligation share of about 3 percent of locations, S^A would rise by about 0.03 times the county's out-commuting share, because the commuting weights sum to that share and not to one. The mean share is about 0.50, so S^A would rise by about 0.015 and manufacturing GDP by about 3.4 percent; this is outside the range of variation in the data and is not a prediction.

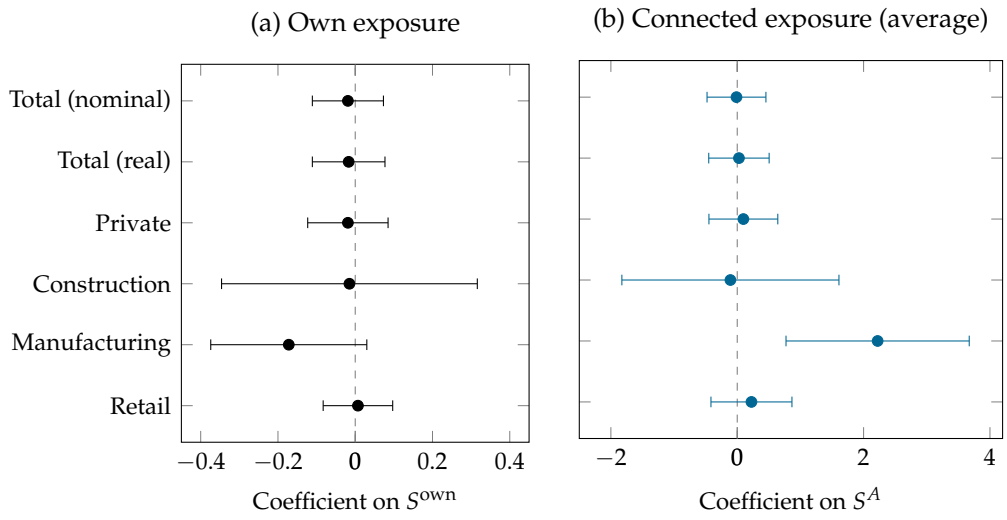


Figure 5: County GDP by industry: own and connected effects

Notes: Point estimates and 95 percent confidence intervals (clustered by county) from the main specification, one regression per outcome; each regression includes S^{own}, S^A, and S^D. Outcomes are log BEA county GDP, chained 2017 dollars except "Total (nominal)." Coefficients are per unit of exposure; note the different horizontal scales. Raw standard deviations are 0.034 for S^{own} and 0.0062 for S^A.

Table 5: Magnitudes: Manufacturing GDP

A. Percent change for realistic changes in connected exposure			
Change in exposure	ΔS^A	Δ ln mfg. GDP	% change
One SD of S^A within county and year	0.0044	0.0098	0.98
One raw SD of S^A	0.0062	0.0138	1.39
Full partner build, 3% obligation (ceiling)	0.015	0.0333	3.38
B. Implied 2024 change, summed over 2,802 counties			
Version	Total	% of baseline	5th–95th pct.
Own + S^A model	\$9.72B	0.45	\$3.32B to \$16.17B
S^A term only	\$12.07B	0.56	\$4.99B to \$19.13B
Own term only (reference)	−\$2.27B	−0.11	−\$4.60B to \$0.06B

Notes: Panel A uses $\hat{\kappa} = 2.222$ from Table 7, column (1); percent change is $100\,(e^{\hat{\kappa}\Delta S^A} - 1)$. Panel B applies the own + S^A model ($\hat{\beta} = -0.167$, $\hat{\kappa} = 1.923$) county by county to 2024 exposures and baseline real manufacturing GDP (\$2.145 trillion). Rows are not additive because of the exponential transformation.

To express the estimate in dollars, I apply the estimated coefficients to each county's 2024 exposure and baseline manufacturing GDP and sum across counties. Using the specification with own exposure and S^A, the implied increase in 2024 manufacturing GDP across 2,802 counties is \$9.7 billion, or 0.45 percent of the \$2.15 trillion baseline, with a 5th to 95th percentile range of \$3.3 to \$16.2 billion. I use this specification because it does not project the imprecise S^D term onto county baselines, and the range comes from draws of the coefficients with their county-clustered covariance. Setting the own-exposure term to zero gives \$12.1 billion. For scale, RDOF authorized support is on the order of \$0.6 billion per year. These are partial-equilibrium, average-effect translations and do not account for offsetting changes elsewhere.

Table 6: Connected Effects per Standard Deviation of S^A

Outcome	$\hat{\kappa}$ per unit	Per raw SD	Per within SD
Log real manufacturing GDP	2.22	1.38 (0.46)	0.98 (0.33)
QCEW log manufacturing employment	0.07	0.07 (0.50)	0.04 (0.28)
QCEW log construction employment	2.15	1.78 (0.79)	1.00 (0.44)
CBP log accommodation and food employment	2.16	0.85 (0.31)	0.65 (0.23)
Log NO_2	1.52	1.26 (0.16)	0.73 (0.09)
Residential log electricity sales	−0.78	−0.60 (0.22)	−0.32 (0.12)
Industrial log electricity sales	−1.13	−0.87 (1.55)	−0.45 (0.81)
Ookla log download speed	3.00	2.44 (1.04)	1.42 (0.60)
ACS high-speed share (percentage points)	1.18	0.93 (0.11)	0.53 (0.06)
Own exposure, for comparison: log manufacturing GDP	−0.17	−0.58 (0.35)	−0.45 (0.27)

Notes: Effects are $100 \times \hat{\kappa} \times$SD, in log points (approximately percent) unless noted, with standard errors in parentheses, from the main specification. Raw SD is the outcome sample's standard deviation of the exposure. Within SD nets out unit and time effects and comes from the matching weather-only runs. For manufacturing GDP, the within SD under the full controls is 0.0039, which would give 0.86 percent.

6.2 Which side of the commuting link?

Column (2) of Table 7 reports the same model in the side-specific parameterization. The residence-side coefficient is $\hat{\delta}_r = 2.759$ (standard error 1.46) and the workplace-side coefficient is $\hat{\delta}_w = -0.537$ (1.16). Taken at face value, the gain loads on deployment in the counties where a county's residents work rather than in the counties where its workers live. I do not interpret the split. The two side-specific estimates have a sampling correlation of about -0.87, the difference S^D is weakly identified, and neither side-specific coefficient is significant at the 5 percent level. The data identify the effect of connected exposure as a whole, not its direction. In particular, the evidence does not establish that the gains occur in "destination hubs" that absorb commuters from rural areas with new broadband; that reading would predict a positive δ_w.

6.3 Output and headcount

Table 8 examines the margins through which manufacturing GDP could rise. QCEW manufacturing employment does not respond to connected exposure ($\hat{\kappa} = 0.07$, standard error 0.55), and its confidence interval excludes the GDP coefficient. QWI manufacturing employment gives a similar point estimate with a wider interval (0.25, standard error 0.97). Manufacturing wages rise weakly in QCEW (0.51, standard error 0.27), and QWI earnings give the same point estimate (0.51, standard error 0.44). Employment and wages together sum to about 0.6, well short of 2.22, so most of the GDP response is not in labor compensation. The QCEW specifications begin in 2021, without the pre-period used for GDP, but splitting BEA's own manufacturing series over the full 2018 to 2024 panel tells the same story: 72 percent of the gain is in non-labor GDP (Table 11). That makes the non-labor allocation discussed in Section 7 central to interpreting the result.

Data Axle establishment records are harder to interpret. Manufacturing establishment counts fall with connected exposure (-0.82, standard error 0.46), significant only at the 10 percent level,

Table 7: Manufacturing GDP (Main Specification and Reparameterizations)

	Log real manufacturing GDP		
	Main: own, S^A, S^D	Side-specific	Own, S^A only
	(1)	(2)	(3)
Own, S^{own}	−0.172*	−0.172*	−0.167
	(0.103)	(0.103)	(0.104)
Connected avg., S^A (κ)	2.222***		1.923***
	(0.740)		(0.672)
Connected diff., S^D (τ)	1.648		
	(1.266)		
Residence side, S^{res} (δ_r)		2.759*	
		(1.465)	
Workplace side, S^{work} (δ_w)		−0.537	
		(1.155)	
$\hat{\beta} + \hat{\kappa}$	2.050	2.050	1.756
	[0.711]	[0.711]	[0.640]
Observations	19,043	19,043	19,043
KP first-stage F	138.21	138.21	517.35

Notes: BEA real manufacturing GDP (CAGDP9), chained 2017 dollars, 2018 to 2024. Column (2) is an exact reparameterization of column (1): $\delta_r = \kappa/2 + \tau$, $\delta_w = \kappa/2 - \tau$. Column (3) drops S^D. Square brackets: standard error of $\hat{\beta} + \hat{\kappa}$ from the joint covariance. Main specification as in Table 2.

while QCEW establishment counts do not move (0.14, standard error 0.36). Data Axle's coverage and industry classification change over the sample, and its listings may themselves respond to connectivity (Section 7), so I place more weight on QCEW.

Industrial electricity sales do not rise detectably with connected exposure (Section 8). The per-customer coefficient is positive (5.33, standard error 3.85) but insignificant. Because the two regressions share a sample, these coefficients imply a coefficient of −6.47 on log industrial customers, so the per-customer result appears to be driven by the customer count and not by usage. It is also not robust to including the 2018 to 2020 pre-period (coefficient 2.27, $t = 0.65$). I do not treat it as evidence of more intensive electricity use.

Across three-digit manufacturing subsectors in QWI (Appendix Table 13), connected exposure has large positive coefficients for nonmetallic minerals (327), primary metals (331), and miscellaneous manufacturing (339), and negative coefficients for furniture (337) and chemicals (325). With 21 subsectors and thin samples in several, a handful of nominal rejections in both directions is expected by chance, so the subsector results are exploratory.

6.4 Where the effect appears

Table 9 estimates the main specification separately by group. Identification comes mostly from the South, where the standard deviation of S^A is 0.0084, against 0.0012 in the West and 0.0014 in the Northeast; first-stage F statistics in those two regions are below 3, so I do not report them. The

Table 8: Manufacturing Margins (Main Specification)

	Own S^{own}	Conn. avg. S^{A}	Conn. diff. S^{D}	Obs.	KP F
Output					
BEA log real mfg. GDP	−0.172*	2.222***	1.648	19,043	138.21
	(0.103)	(0.740)	(1.266)		
Labor (QCEW, 2021 onward)					
Log mfg. employment	−0.056	0.074	−0.803	126,882	73.42
	(0.069)	(0.549)	(0.987)		
Log mfg. average weekly wage	−0.041	0.515*	0.764*	126,882	73.42
	(0.031)	(0.268)	(0.458)		
Log mfg. establishments	−0.027	0.141	0.233	126,882	73.42
	(0.055)	(0.365)	(0.615)		
Labor (QWI, quarterly, 2015 onward)					
Log mfg. employment	0.124	0.252	−0.233	343,564	114.50
	(0.114)	(0.971)	(1.472)		
Log mfg. earnings	−0.014	0.511	0.936	361,901	121.38
	(0.052)	(0.437)	(0.868)		
Establishments (Data Axle, 2019 to 2024)					
Log mfg. establishments	−0.025	−0.819*	−0.622	18,606	119.77
	(0.061)	(0.456)	(0.875)		
Log mean establishment size	−0.095	−1.700	−1.258	18,310	120.79
	(0.130)	(1.124)	(1.918)		
Industrial electricity (EIA-861, utility-state, 2021 to 2024)					
Log sales	0.204	−1.132	2.229	4,663	99.38
	(0.501)	(2.015)	(3.807)		
Log sales per customer	0.611	5.334	8.318	4,663	99.38
	(0.748)	(3.847)	(5.659)		

Notes: Each row is a separate 2SLS regression with the main specification (Table 2). Electricity regressions use utility-by-state and year fixed effects, with exposure aggregated to utility service territories, and cluster by utility-state. Other rows cluster by county. QCEW rows begin in 2021, with no pre-period. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

manufacturing coefficient has the same sign and a similar size in the South (1.79, standard error 0.76) and the Midwest (2.75, standard error 2.05), though the Midwest estimate is imprecise. Log NO_2 and residential electricity sales, by contrast, change sign between those regions.

The manufacturing response is largest where manufacturing is a small share of the economy: 4.06 and 4.05 in the bottom and middle terciles of baseline manufacturing share, against 1.10 in the top tercile (test of equality, $p = 0.096$). Because the outcome is in logs, these are proportional effects on smaller bases. The pattern does not suggest that connectivity works mainly by deepening an existing industrial base. It could also reflect noisier measurement of small manufacturing sectors in BEA's county estimates. Log NO_2 does not vary with manufacturing share ($p = 0.95$), so it does not appear to be driven by manufacturing.

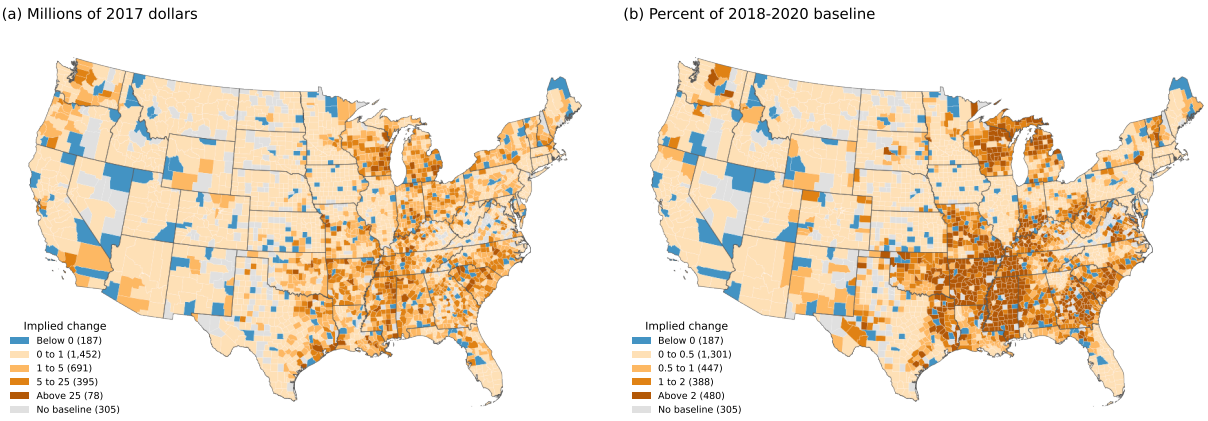


Figure 6: Where the implied manufacturing gains are

Notes: Implied 2024 change in real manufacturing GDP from the own + S^A model, $(e^{\hat{\beta}S_i^{\text{own}}+\hat{\kappa}S_i^A}-1)\times \text{GDP}_{i,\text{pre}}^{\text{mfg}}$, in dollars and as a percent of the 2018 to 2020 baseline. These are translations of the estimated coefficients, not separately estimated county effects. The top ten states by implied gain hold 68 percent of the gain and 34 percent of baseline manufacturing GDP; Mississippi holds 9.3 percent of the gain on 0.8 percent of baseline.

Table 9: Connected Effect by Group ($\hat{\kappa}$)

	Mfg. GDP	Log NO_2	Residential sales	Comm. + ind. sales
Census region				
South	1.790**	0.971***	−0.565**	−2.235
	(0.764) [99]	(0.115) [101]	(0.256) [102]	(2.497) [94]
Midwest	2.750	−1.988	1.216	4.184
	(2.052) [21]	(1.281) [37]	(1.158) [17]	(3.115) [17]
Baseline manufacturing share				
Bottom tercile	4.063*	1.658***		
	(2.088) [39]	(0.321) [56]		
Middle tercile	4.051***	1.684***		
	(1.421) [73]	(0.601) [43]		
Top tercile	1.099	1.468***		
	(0.671) [115]	(0.315) [74]		

Notes: Each cell is the main specification on one group alone. Standard errors in parentheses, clustered by county (utility-state for electricity); own-sample first-stage F in brackets. Log sales in the electricity columns. Northeast and West are omitted because their first-stage F statistics are below 3. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7 Threats to Identification and Robustness

7.1 Robustness of the manufacturing estimate

Table 10 and Figure 7 report robustness checks for the connected coefficient on manufacturing GDP. Each row changes one element of the main specification.

Adding manufacturing-share-by-year controls, which allow counties with different baseline manufacturing intensity to follow different aggregate paths, raises the estimate to 2.51 ($t = 3.37$). Census division-by-year fixed effects, which absorb any shock common to a division, leave it at 2.17 ($t = 2.76$). Dropping any one state leaves the point estimate between 1.88 and 2.63, and drop-

Table 10: Robustness: Connected Effect on Log Real Manufacturing GDP

Row	Change from main specification	$\hat{\kappa}$	SE	t
0	None (main specification)	2.222	0.740	3.00
1	Add manufacturing share × year	2.51	0.745	3.37
2	Census division × year FE	2.17	0.788	2.76
3	State × year FE	1.61	0.879	1.83
4	Commuting zone × year FE	1.35	1.096	1.23
5	Drop S^D (own + S^A only)	1.923	0.672	2.86
6	Control linearly for μ (OLS) instead of 2SLS	3.42	0.966	3.54
7	Drop ten largest manufacturing states (N = 13,025)	2.24	0.831	2.70
8	Recenter within state and cohort (weather only; 2.10 without)	1.53	0.896	1.71
9	Pre-2021 clean-manufacturing capex × year	2.23	0.740	3.01
10	Cumulative announced clean-manufacturing capex	2.22	0.740	3.00
11	Drop counties with any clean-manufacturing facility	2.39	0.750	3.19
12	Add wildfire smoke days (N = 16,313)	1.71	0.612	2.79
13	South only (N = 8,552)	1.79	0.764	2.34
Placebo: 2019 relative to 2020, reduced form, 2015 to 2024				
P1	County and year FE	−1.45	0.710	−2.04
P2	State × year FE	−0.01	0.766	−0.02

Notes: Each row changes one element of the main specification. All rows cluster by county. $N = 19{,}043$ unless noted; row 11 has $N = 17{,}090$. The standard error for row 7 is implied by its t-statistic. Placebo coefficients are on the final-year instrument, on the reduced-form scale of Figure 8.

ping the ten largest manufacturing states gives 2.24 ($t = 2.70$). Finer geographic controls remove much of the variation that identifies a connected effect, because connected exposure is spatially smooth by construction: state-by-year fixed effects absorb about 26 percent of the variance of the connected instrument and commuting zone-by-year fixed effects about 69 percent, and counties in the same commuting zone are often partners of one another. The estimate stays positive and is less precise under them, at 1.61 ($t = 1.83$) and 1.35 ($t = 1.23$). Recentering within state and authorization cohort, which changes the cells and not the fixed effects, gives 1.53 ($t = 1.71$), against 2.10 in the same weather-only specification. The confidence intervals under these specifications include the main estimate.

The placebo rows and Figure 8 report the event study in the reduced form. Relative to 2020, the 2019 coefficient is -1.45 ($t = -2.04$) under county and year fixed effects, reflecting a dip in 2018 and 2019, and -0.01 ($t = -0.02$) under state-by-year fixed effects. The 2024 coefficient is nearly the same in both (2.13 and 2.06, with $t = 1.95$ and 1.76). The joint test of the five pre-period coefficients does not reject in either case ($p = 0.30$ and 0.65). The effect builds as deployment proceeds, and the post-deployment trajectory is nearly the same under both sets of fixed effects: the 2022 to 2024 coefficients are 1.16, 1.81, and 2.13 under county and year fixed effects, and 0.93, 1.80, and 2.06 under state-by-year fixed effects. What state-by-year fixed effects remove is the 2019 dip, and the lower pooled coefficient under them is consistent with that dip no longer counting toward the comparison of later years with earlier ones.

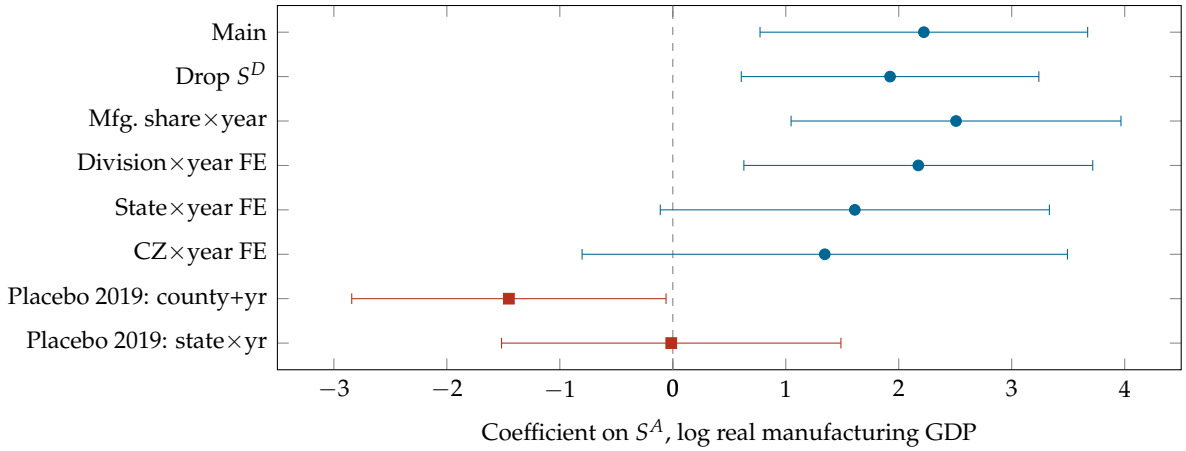


Figure 7: Robustness of the connected manufacturing effect

Notes: Point estimates and 95 percent confidence intervals from Table 10. Blue circles: 2SLS post-period estimates on S^A. Red squares: reduced-form placebo coefficients on Z^A for 2019 outcomes relative to 2020, as in Figure 8, which are on a different scale.

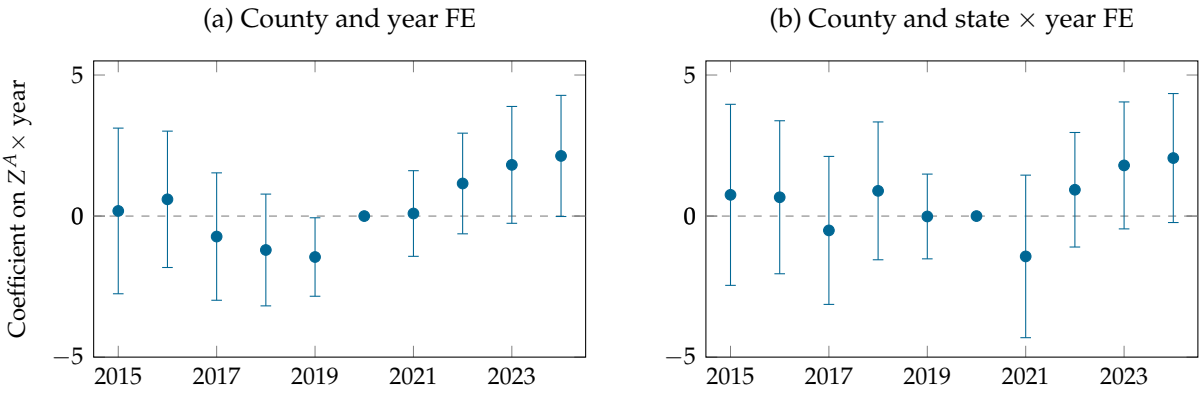


Figure 8: Dynamics of the connected manufacturing effect

Notes: Reduced-form coefficients on the recentered connected instrument interacted with year, relative to 2020, in log real manufacturing GDP, 2015 to 2024, with own and difference terms included. Bars are 95 percent confidence intervals, clustered by county. Joint test of the five pre-period coefficients: $p = 0.30$ in panel (a) and $p = 0.65$ in panel (b).

7.2 A coinciding regional manufacturing shock

The leading alternative explanation is that connected exposure correlates with a regional manufacturing investment wave in the same years, notably the large battery and electric vehicle plant investments in the rural Southeast and Midwest after 2022. Such a shock would raise manufacturing GDP and construction activity in the same regions where RDOF deployment was proceeding. Other outcomes illustrate the risk. Log outage customer-hours have no obvious reason to respond to partners' broadband, yet they load strongly on connected exposure under county and month fixed effects (7.65, $t = 4.07$) and the estimate disappears with commuting zone-by-month fixed effects (-1.19, $t = -0.39$). Log NO_2 shows the same pattern (1.44, $t = 7.30$, falling to 0.22, $t = 1.50$, both without smoke controls; with them the main estimate is 1.52).

Clean-technology plants do not explain the estimate. Using facility-level announcements from the Clean Investment Monitor, adding pre-2021 announced capital spending interacted with year leaves the coefficient at 2.23 ($t = 3.01$), controlling for cumulative announced spending leaves it at 2.22 ($t = 3.00$), and dropping all counties with any facility gives 2.39 ($t = 3.19$). Connected exposure also does not predict new announcements: its coefficient on any announcement in a year is -0.07 ($t = -0.52$), and on log cumulative announced spending it is -2.80 ($t = -1.87$). Two limits remain. The Clean Investment Monitor covers clean technology only, so it does not speak to semiconductors or other plant investment, and it records announced totals at the announcement date, not spending over time. Also, QCEW construction employment and wages rise with connected exposure (2.15, $t = 2.25$, and 1.09, $t = 3.23$), although BEA construction GDP does not. Those estimates are not changed by the clean-investment controls, and they could reflect network construction or other building in counties with connected exposure.[2]

7.3 How BEA measures county manufacturing GDP

County GDP is not directly measured. For most industries, BEA constructs county GDP from county compensation and proprietors' income plus an allocated share of the state's remaining income components (gross operating surplus and taxes less subsidies). For goods-producing industries, the allocation of this non-labor component uses Economic Census value added and payroll in benchmark years and is interpolated or extrapolated between benchmarks with sales from a commercial establishment database (Bureau of Economic Analysis, 2020). Real county GDP is current-dollar GDP deflated by national industry price indexes, so it contains no local price variation.

This construction matters in two ways. First, manufacturing GDP can rise without manufacturing earnings rising if the non-labor component moves, so the combination of rising GDP with flat employment and modest wage gains is not impossible by construction. Second, if the commercial data used to extrapolate between census years change their coverage or modeled sales with connectivity, BEA's county manufacturing GDP could inherit that change. The description here follows BEA's 2020 primer; the methodology for the 2018 to 2024 vintage may differ, including in how it incorporates the 2022 Economic Census.

Table 11 splits manufacturing GDP directly, using BEA compensation of employees (CAINC6N, by place of work) and current-dollar manufacturing GDP (CAGDP2), with non-labor GDP as their difference. On county-years with both series, log compensation does not respond detectably (0.97, standard error 0.74), log non-labor GDP does (2.70, standard error 0.83), and compensation's share of GDP falls (-0.53, standard error 0.27). In dollars scaled by each county's 2018 to 2020 manufacturing GDP, the response of 2.16 splits exactly into 0.60 of compensation and 1.57 of non-labor GDP, so about 72 percent of the gain is non-labor. This agrees with the QCEW results. It leaves two readings open: returns to capital may truly be rising, or BEA's allocation of the non-labor component may be moving with its commercial data inputs. These data cannot separate the two,

[2]Future work will add non-clean plant investment, such as CHIPS Act awards.

Table 11: Labor and Non-Labor Parts of BEA Manufacturing GDP, and Suppression

	Own S^{own}	Conn. avg. S^{A}	Conn. diff. S^{D}	Obs.	KP F
Logs, county-years with both series (current dollars)					
Log manufacturing GDP	−0.120	1.672***	0.746	18,705	133.3
	(0.098)	(0.629)	(1.120)		
Log compensation	−0.112	0.968	−0.307	18,705	133.3
	(0.118)	(0.737)	(1.329)		
Log non-labor GDP	−0.087	2.699***	2.390	18,672	135.5
	(0.120)	(0.835)	(1.547)		
Compensation share of GDP	0.014	−0.534**	−0.710	18,705	133.3
	(0.046)	(0.271)	(0.448)		
Dollars / county's mean 2018 to 2020 manufacturing GDP					
Manufacturing GDP	−0.163	2.161**	0.395	18,570	135.6
	(0.181)	(1.029)	(2.032)		
Compensation	−0.066	0.595	−0.948	18,570	135.6
	(0.120)	(0.665)	(1.620)		
Non-labor GDP	−0.098	1.566**	1.343	18,570	135.6
	(0.132)	(0.772)	(1.370)		
Suppression indicators, county-years with total GDP					
Manufacturing GDP suppressed	0.117	−0.871	0.367	21,349	116.8
	(0.093)	(0.743)	(1.358)		
Compensation suppressed	0.094	−0.568	0.252	21,349	116.8
	(0.082)	(0.671)	(1.231)		

Notes: Main specification, 2018 to 2024, county clusters. Non-labor GDP is manufacturing GDP (CAGDP2) minus compensation of employees (CAINC6N). Non-labor GDP can be negative: 33 such county-years are lost from the log row and kept in the scaled rows, which add up exactly. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

so the remaining test is whether commercial-database sales move with exposure.

The nominal coefficient in Table 11 (1.67) is below the real coefficient in the main specification (2.22). If real GDP were nominal GDP deflated by a single national manufacturing price index, the two would match on the same sample. The gap may reflect the slightly smaller sample here or deflation at a finer industry level, which estimating both on the same county-years will separate.

Suppression. The manufacturing sample is about 11 percent smaller than the total GDP sample because of disclosure suppression. If suppression is related to exposure, the estimate could reflect selection. On the 21,349 county-years with total GDP, connected exposure does not predict suppression of manufacturing GDP (−0.87, standard error 0.74) or of manufacturing compensation (−0.57, standard error 0.67).

7.4 Other threats

Endogenous denominator. No Fabric vintage predates the auction, and the FCC's RDOF-specific location adjustment process means the BSL count can respond administratively to the program. Net changes across Fabric versions are under 1 percent but gross additions and deletions are about

4 percent and are concentrated in rural and tribal areas. An eligible-locations denominator, which is fixed at the auction, is a planned robustness check.

Pre-determined network. The commuting matrices are built from pre-period LODES flows and held fixed, so post-period changes in commuting that broadband itself may cause do not enter the exposure measures. LODES omits self-employment, much agricultural work, and non-work trips; connected exposure therefore captures linkages through formal wage employment.

Shared infrastructure. As discussed in Section 5, connected exposure predicts improvements in a county's own connectivity. Controlling for non-RDOF changes in availability from the FCC Broadband Data Collection would indicate how much of the manufacturing effect runs through the county's own service rather than through its partners; this is planned for future work.

Multiple testing. The project has estimated several hundred outcomes, many of them sector or subsector breakdowns. The manufacturing GDP result was not selected from among these: BEA industry aggregates were estimated as a set (Figure 5), and manufacturing is the only one that responds ($p = 0.0027$, or $p = 0.016$ after a Bonferroni correction across the six BEA aggregates). Subsector results are presented as exploratory.

8 Electricity, Air Quality, and Outages

If connectivity and electricity are complements in manufacturing, electricity use should rise where manufacturing output does. Table 12 reports what the available data show. Exposure for electricity outcomes is less clean than for county outcomes, as described below, and the air quality and outage results point to regional confounding under the main fixed effects. I therefore treat this section as suggestive.

Electricity. EIA-861 sales have no county field. Exposure for a utility-state-year is a population-weighted mean over the counties that utility lists in its service territory. Three features of this mapping matter. The list can change from year to year, so exposure can move for reasons unrelated to deployment; the population weight appears to be the whole county's, not the population that utility serves; and utilities with no listed counties drop out, including TVA's direct retail sales in Alabama, Mississippi, and Tennessee. To test the first, I hold each utility's county list and weights at their 2019 values.

Industrial sales do not change detectably, and the result does not depend on the list (-1.13, standard error 2.02, and -1.27, standard error 2.02, with lists fixed). Commercial sales and combined commercial and industrial sales are insignificant, and their point estimates move substantially when lists are fixed, which suggests that changes in the lists affect them. The combined measure addresses inconsistent classification across the two classes, and it also shows no detectable change. These intervals are wide, so they cannot rule out an electricity response of the same size

as the manufacturing GDP response. Residential sales fall (-0.78, standard error 0.29) and are stable with fixed lists (-0.75, standard error 0.29). Sales per residential customer do not change detectably (-0.31, standard error 0.27). The residential decline is present in the South and not in the Midwest, and I have not estimated it under state-by-year fixed effects.

Air quality. Log NO_2 rises with connected exposure in the main specification (1.52, standard error 0.20, with smoke controls). Without smoke controls, where the main estimate is 1.44, it falls to 0.48 (standard error 0.16) with state-by-month fixed effects and to 0.22 (standard error 0.14) with commuting zone-by-month fixed effects, and it comes entirely from the South (Table 9). NO_2 comes mainly from combustion, and the estimate does not vary with baseline manufacturing share, which points to vehicle emissions more than manufacturing, although I have not yet tested this with traffic data. $PM_{2.5}$ at the block group level shows no detectable effect once smoke is controlled for (-0.31, standard error 0.23). I do not use the county $PM_{2.5}$ file, which covers about three quarters of counties and disagrees with the block group estimate.

Outages as a negative control. There is no obvious reason for commuting partners' broadband to raise power outages, yet log customer-hours out loads strongly on connected exposure in the main specification (7.65, standard error 1.88). The estimate falls to 3.24 (standard error 2.09) with state-by-month fixed effects and to -1.19 (standard error 3.06) with commuting zone-by-month fixed effects. This indicates that connected exposure under county and time fixed effects can pick up regional patterns unrelated to broadband.

Data centers. Data center activations might seem a natural place for broadband to matter. Own exposure has no detectable effect. Connected exposure is negative in the main specification (log activations -0.036, standard error 0.014; log megawatts -0.18, standard error 0.07), but both vanish with commuting zone-by-month fixed effects, the same pattern as NO_2 and outages. A null is what one would expect. RDOF builds last-mile access to homes and small businesses, while data centers need constant connectivity over redundant long-haul routes and large amounts of power, which last-mile deployment does not provide.

9 Interpretation

The results so far establish a reduced-form fact: RDOF deployment in a county's commuting partners is associated with higher manufacturing value added in that county, with no detectable change in total output or manufacturing employment. This section offers a simple way to think about why connectivity could raise manufacturing output and what that would imply for energy use. It is not estimated.

Consider a plant that combines connectivity-augmented capital ZK and electricity E in a bundle Φ, with output subject to decreasing returns because other factors, such as labor and the plant

Table 12: Electricity, Air Quality, Outages, and Data Centers ($\hat{\kappa}$ on S^A)

	Main	State × time	CZ × time	Fixed 2019 lists
Log NO_2	1.515***	0.477***	0.217	
	(0.197)	(0.158)	(0.144)	
Log $PM_{2.5}$, block group†	−0.306		−0.245	
	(0.228)		(0.192)	
Log outage customer-hours	7.654***	3.238	−1.189	
	(1.882)	(2.091)	(3.065)	
Log data center activations	−0.036***		0.000	
	(0.014)		(0.025)	
Log data center megawatts	−0.181**		−0.053	
	(0.071)		(0.131)	
Log electricity sales				
Residential	−0.781***			−0.749***
	(0.286)			(0.286)
Commercial	−2.130			−0.355
	(2.096)			(1.250)
Industrial	−1.132			−1.267
	(2.015)			(2.015)
Commercial + industrial	−1.128			0.545
	(2.007)			(1.088)

Notes: Main specification, with smoke controls for NO_2 and $PM_{2.5}$ in the Main column. County clusters; utility-state clusters for electricity. State and commuting zone columns for NO_2 and $PM_{2.5}$ do not include smoke controls. †Recipient-by-state clusters, not yet re-run with county clusters. Blank cells were not estimated. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

itself, are fixed:

$$Y = \Phi^{\mu}, \quad 0 < \mu < 1, \qquad \Phi = \left[b\,(ZK)^{\frac{\zeta-1}{\zeta}} + (1-b)\,E^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}}, \tag{7}$$

where Z indexes connectivity-enabled capital productivity (networked controls, sensors, enterprise software, remote coordination) and ζ is the elasticity of substitution between effective capital and electricity. With input and output prices fixed, profit maximization implies

$$\frac{d\ln Y}{d\ln Z} = s_K \frac{\mu}{1-\mu}, \qquad \frac{d\ln E}{d\ln Z} = s_K\Big(\frac{1}{1-\mu} - \zeta\Big), \qquad \frac{d\ln(E/Y)}{d\ln Z} = s_K\,(1-\zeta), \tag{8}$$

where s_K is the cost share of effective capital in Φ. Connectivity raises electricity use unless capital and electricity are very strong substitutes, because the plant scales up. Electricity intensity rises only if the two are complements ($\zeta < 1$). The ratio of the two responses, $(d\ln E)/(d\ln Y) = [1-\zeta(1-\mu)]/\mu$, depends only on ζ and μ, so an estimated electricity response relative to the output response would reveal which case applies.

The data do not yet discriminate among these cases. Output rises without a detectable employment response, and the effect builds over time, as a slow-moving complementarity would. But the industrial electricity estimate (−1.13, 95 percent interval roughly −5.1 to 2.8) includes both no response and a response equal to the output response, and utility-level sales are a blunt measure of

plant-level use. The largest output responses are also not in manufacturing-heavy counties (Table 9), which is not what deepening of an existing base would predict. Existing evidence is mixed: energy appears to be a strong short-run complement to capital at the aggregate level (Hassler et al., 2021), in French services computers and software raised electricity intensity while communication equipment lowered it (Collard et al., 2005), and in UK firms broadband-enabled communication ICT raised firm size but not productivity (DeStefano et al., 2018). What the framework clarifies is what a test needs: output and electricity responses measured for the same plants or places.

The framework is silent on why the response runs through connected exposure rather than own exposure. Three explanations are consistent with the evidence and are not yet separated: (i) shared infrastructure, in which RDOF construction upgraded middle-mile capacity and spurred provider investment in linked counties, raising their own connectivity (Table 3); (ii) labor and supplier integration, in which rural workers and suppliers with better broadband become more productive inputs to manufacturers in linked counties; and (iii) construction staging, in which network construction temporarily raised activity in the places from which it was staged. QCEW construction employment and wages and CBP accommodation and food employment all rise in linked counties, which fits staging, as do the QWI employment gains in nonmetallic minerals and primary metals, although BEA construction GDP does not rise and the subsector results are exploratory. The non-labor result fits the first two explanations if better connectivity raises returns on existing plant without more headcount, and it fits none of them if it reflects BEA's allocation of non-labor GDP. The last predicts that effects track the pace of partner deployment and fade after construction; the first two predict persistence. Distinguishing the flow of deployment from its stock is the natural test.

Community opposition may also shape where large electricity users locate, which would matter for reading the data center results and for any local energy response. Nothing in this paper tests that, and I mention it only as a question for future work.

10 Conclusion

Rural broadband programs are evaluated, implicitly or explicitly, by what happens in the places that get wired. This paper shows that for the largest recent U.S. program, the measurable output effects so far run through a county's links to other places more than through its own deployment. Using the timing of RDOF deployment within recentering cells and pre-period commuting networks, I find that the own effects on GDP and employment are close to zero, while connected exposure raises manufacturing value added without a detectable rise in manufacturing employment. Deployment among commuting partners also improves a county's broadband.

The result has two implications. For program evaluation, designs that compare treated and untreated places will miss effects that run through networks, and may misattribute them to the comparison group. For the environmental economics of digitalization, the location and composition of the response matter: if broadband raises industrial output through connected exposure rather than substituting digital activity for physical activity where the infrastructure is built, its

consequences for energy use and emissions depend on where those counties are and what they produce.

Several next steps would sharpen these results. On inference, I plan randomization inference and spatially robust standard errors that account for dependence across counties sharing commuting partners. On mechanisms, I plan to separate broadband connectivity from commuting links, by comparing partners that border a county with partners farther away and by controlling for non-RDOF changes in a county's own availability; to test whether BEA's non-labor allocation, through which most of the manufacturing gain runs, moves with its commercial data inputs; to control for manufacturing investment beyond clean technology; and to test the vehicle-emissions reading of the NO_2 result with traffic data. On energy, a sharper measure of industrial electricity use would allow a direct test of the complementarity hypothesis, and power plant activations and generation from EPA Continuous Emissions Monitoring System data will extend the analysis to the supply side. Direct access to the Fabric will allow a better-documented exposure denominator. Finally, these estimates can feed a fuller accounting of the program's benefits and costs, using the marginal value of public funds (Hendren and Sprung-Keyser, 2020): output gains, the program's fiscal cost, and external costs and benefits such as changes in emissions and electricity use.

References

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2023). When Should You Adjust Standard Errors for Clustering? *Quarterly Journal of Economics* 138(1), 1–35.

Akerman, A., I. Gaarder, and M. Mogstad (2015). The Skill Complementarity of Broadband Internet. *Quarterly Journal of Economics* 130(4), 1781–1824.

Atasoy, H. (2013). The Effects of Broadband Internet Expansion on Labor Market Outcomes. *ILR Review* 66(2), 315–345.

Aterio (2026). Data Center Inventory, snapshot of March 4, 2026. Proprietary dataset. aterio.io.

Bartel, A., C. Ichniowski, and K. Shaw (2007). How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills. *Quarterly Journal of Economics* 122(4), 1721–1758.

Bureau of Economic Analysis (2020). A Primer on Local Area Gross Domestic Product Methodology. *Survey of Current Business*, March 2020.

Borusyak, K., and P. Hull (2023). Nonrandom Exposure to Exogenous Shocks. *Econometrica* 91(6), 2155–2185.

Bresnahan, T. F., and M. Trajtenberg (1995). General Purpose Technologies: "Engines of Growth"? *Journal of Econometrics* 65(1), 83–108.

Brynjolfsson, E., and L. M. Hitt (2000). Beyond Computation: Information Technology, Organizational Transformation and Business Performance. *Journal of Economic Perspectives* 14(4), 23–48.

Brynjolfsson, E., D. Rock, and C. Syverson (2021). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics* 13(1), 333–372.

Rhodium Group and MIT CEEPR (2026). Clean Investment Monitor, manufacturing download, 2026 Q3 release. cleaninvestmentmonitor.org.

Collard, F., P. Fève, and F. Portier (2005). Electricity Consumption and ICT in the French Service Sector. *Energy Economics* 27(3), 541–550.

Copeland, B. R., and M. S. Taylor (2004). Trade, Growth, and the Environment. *Journal of Economic Literature* 42(1), 7–71.

DeStefano, T., R. Kneller, and J. Timmis (2018). Broadband Infrastructure, ICT Use and Firm Performance: Evidence for UK Firms. *Journal of Economic Behavior & Organization* 155, 110–139.

Dettling, L. J., S. Goodman, and J. Smith (2018). Every Little Bit Counts: The Impact of High-Speed Internet on the Labor Force Participation of Married Women. *Review of Economics and Statistics* 100(2), 260–273.

Faber, B. (2014). Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System. *Review of Economic Studies* 81(3), 1046–1070.

Federal Communications Commission (2020). Rural Digital Opportunity Fund Phase I Auction (Auction 904) Closes; Winning Bidders Announced; FCC Form 683 Due January 29, 2021. Public Notice DA 20-1422, December 7, 2020.

Forman, C., A. Goldfarb, and S. Greenstein (2012). The Internet and Local Wages: A Puzzle. *American Economic Review* 102(1), 556–575.

Hassler, J., P. Krusell, and C. Olovsson (2021). Directed Technical Change as a Response to Natural Resource Scarcity. *Journal of Political Economy* 129(11), 3039–3072.

Hendren, N., and B. Sprung-Keyser (2020). A Unified Welfare Analysis of Government Policies. *Quarterly Journal of Economics* 135(3), 1209–1318.

Hjort, J., and J. Poulsen (2019). The Arrival of Fast Internet and Employment in Africa. *American Economic Review* 109(3), 1032–1079.

Kolko, J. (2012). Broadband and Local Growth. *Journal of Urban Economics* 71(1), 100–113.

Lange, S., J. Pohl, and T. Santarius (2020). Digitalization and Energy Consumption. Does ICT Reduce Energy Demand? *Ecological Economics* 176, 106760.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018). Commuting, Migration, and Local Employment Elasticities. *American Economic Review* 108(12), 3855–3890.

Rupasingha, A., J. Pender, and R. Williams (2024). Broadband and Rural Business: Evidence from the Broadband Initiatives Program. *Economic Inquiry* 62(2), 705–728.

Whitacre, B., R. Gallardo, and S. Strover (2014). Broadband's Contribution to Economic Growth in Rural Areas: Moving towards a Causal Relationship. *Telecommunications Policy* 38(11), 1011–1023.

A Appendix Tables and Figures

Table 13: QWI Manufacturing Employment by Three-Digit NAICS (Main Specification)

NAICS	Subsector	Own, S^{own}	Conn. avg., S^{A}	Conn. diff., S^{D}	Obs.	KP F
311	Food	0.119 (0.427)	0.387 (3.964)	1.305 (5.050)	197,139	122.89
312	Beverage and tobacco	−0.974* (0.519)	−1.344 (8.160)	−11.43 (15.73)	97,773	27.74
313	Textile mills	−4.830* (2.560)	−15.14* (8.751)	−16.36 (13.03)	27,308	19.27
314	Textile products	2.382 (1.495)	−7.553* (3.941)	5.388 (11.60)	59,542	14.23
315	Apparel	−1.320 (2.905)	26.84 (29.02)	10.57 (36.03)	31,763	23.56
316	Leather	−4.027 (7.778)	41.67 (72.59)	−0.848 (62.01)	13,331	1.52†
321	Wood products	−0.184 (0.218)	1.540 (1.820)	5.326** (2.544)	170,534	71.92
322	Paper	0.707 (1.048)	3.819 (5.174)	1.004 (6.075)	55,296	43.59
323	Printing	−0.523 (0.655)	−0.057 (3.914)	1.924 (4.900)	142,617	75.63
324	Petroleum and coal	−4.292 (2.771)	5.084 (5.286)	−14.15 (15.82)	26,524	25.20
325	Chemicals	−0.432 (0.515)	−6.821* (3.927)	−5.039 (5.707)	124,824	18.04
326	Plastics and rubber	0.010 (0.535)	0.427 (2.426)	−4.043 (4.212)	111,765	62.57
327	Nonmetallic minerals	−0.335 (0.405)	7.524** (3.302)	4.851 (4.480)	152,430	89.09
331	Primary metals	0.151 (0.656)	9.194** (4.473)	−10.64 (6.873)	59,068	37.19
332	Fabricated metals	0.175 (0.250)	−0.853 (1.550)	0.580 (2.893)	226,191	75.53
333	Machinery	0.385 (0.246)	−0.699 (1.776)	0.548 (3.614)	172,487	106.21
334	Computer and electronic	0.906 (1.288)	−6.464 (7.040)	−20.81* (11.95)	87,342	24.70
335	Electrical equipment	0.184 (2.017)	−13.54 (16.55)	−19.56 (16.13)	68,211	28.26
336	Transportation equipment	−0.198 (0.416)	0.142 (2.538)	2.025 (3.989)	124,029	70.61
337	Furniture	1.063** (0.533)	−8.723*** (3.320)	−9.595* (5.799)	128,521	73.09
339	Miscellaneous	0.193 (0.515)	11.86*** (3.833)	3.033 (5.753)	145,445	92.77

Notes: QWI log employment, 2015Q1 to 2025Q1; quarterly values repeat across months, with month fixed effects. Main specification. County-clustered standard errors in parentheses. † Kleibergen-Paap $F < 10$. All 21 manufacturing subsectors estimated are reported. Exploratory: with 21 subsectors, several nominal rejections in each direction are expected by chance.

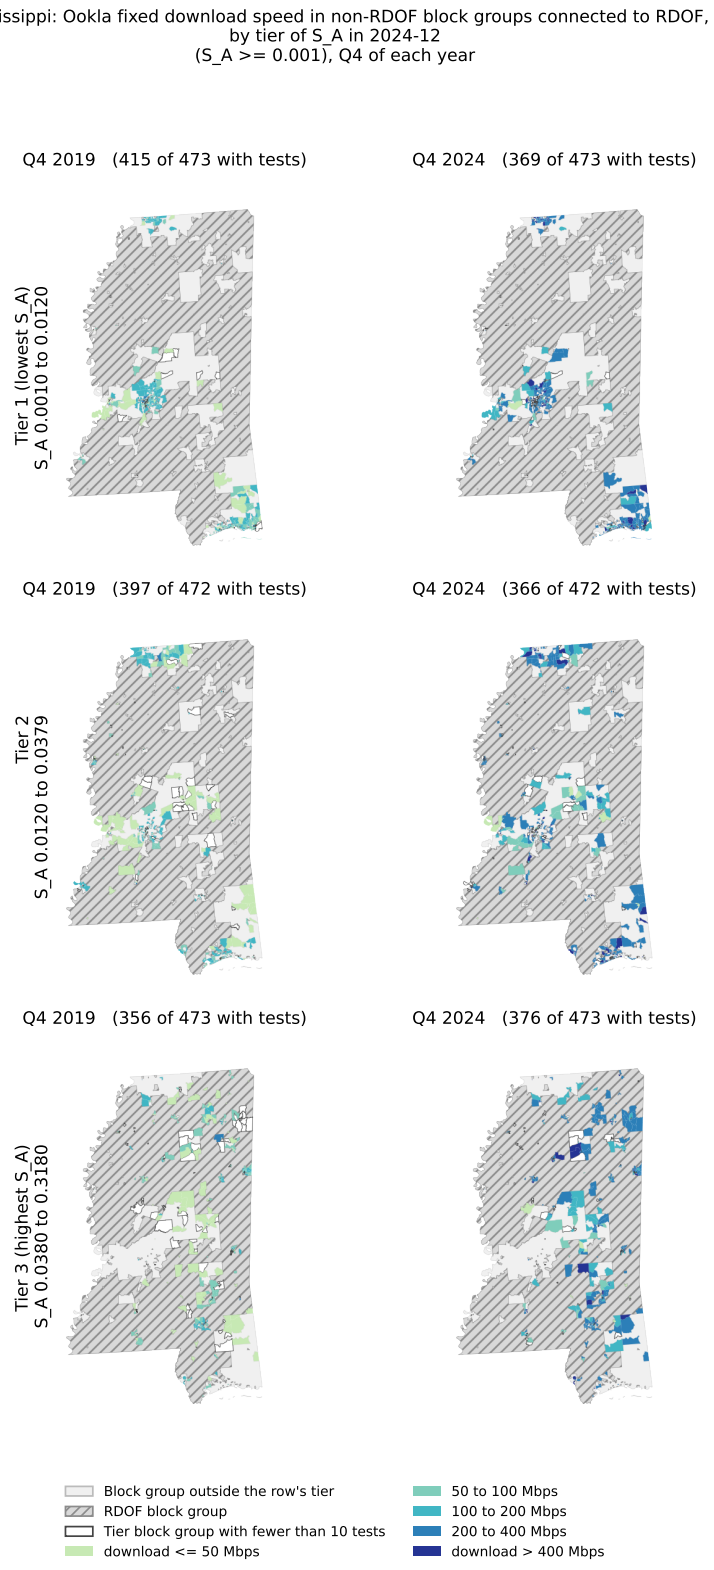


Figure 9: Speeds in commuting-linked non-RDOF block groups, Mississippi

Notes: Descriptive. Fourth-quarter Ookla fixed download speed in 2019 and 2024, in block groups by tier of connected exposure. Each block group's value is the test-weighted mean of the Ookla tiles whose centers fall inside it. Speeds rose in all groups, so no premium for commuting-linked places is implied.

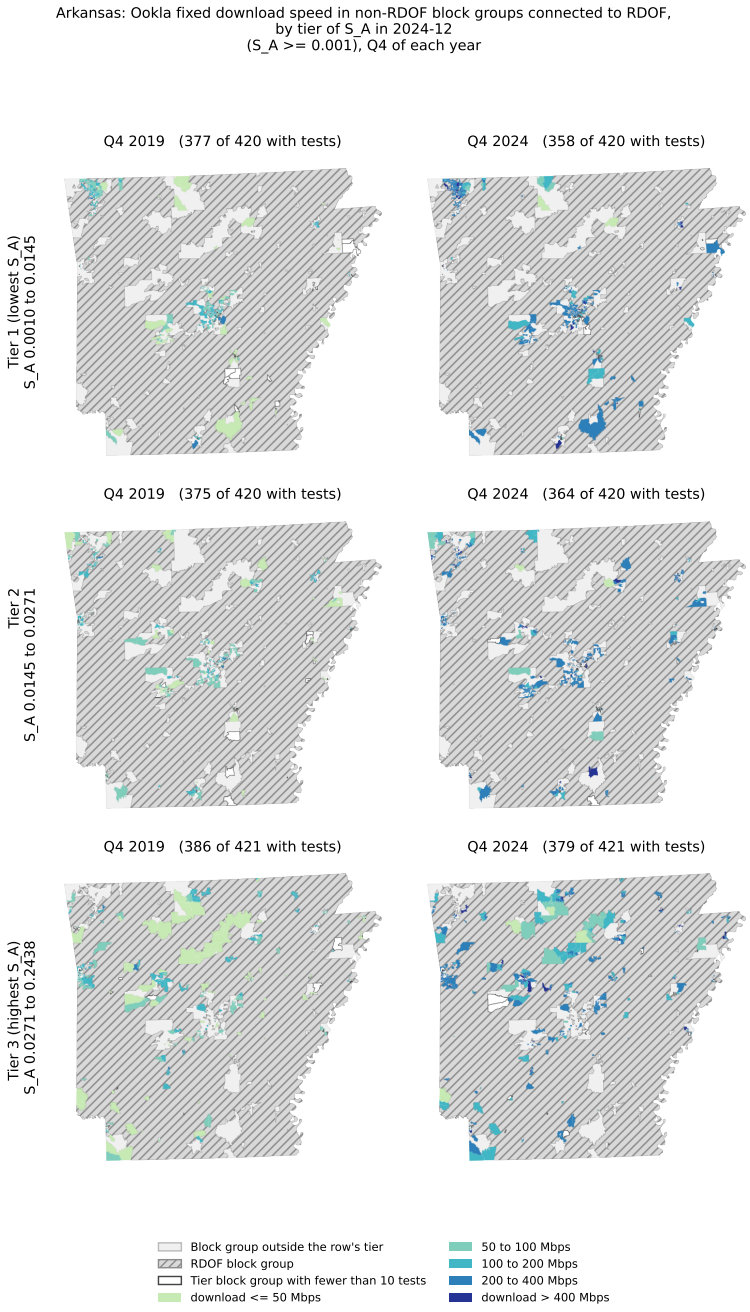


Figure 10: Speeds in commuting-linked non-RDOF block groups, Arkansas

Notes: See Figure 9.