Northwestern | RESEARCH AND INNOVATION FOR SOCIAL AND ECONOMIC INCLUSION LAB

POLICY BRIEF · SERIES NO. 2026-04 · JULY 2026

The Demand Side of Digital Infrastructure

Michelle Yin, PhD, Director and Professor · **Diego Guerrero**, Lead Economist | School of Education and Social Policy · RISEI Lab · Northwestern University

Causal evidence from a $9 billion natural experiment on who gains from broadband investment — and what it means for BEAD, the $21 billion in restructuring savings, and the AI-era buildout

OVERVIEW

Who is connected determines what a connection is worth. America is in the middle of the largest public and private buildout of digital infrastructure in its history, and nearly the entire policy debate is about supply — how much to build, how fast, and where. Evidence from a $9 billion natural experiment shows that the returns are set on the other side of the wire, by the populations served, the constraints they face, and the institutions around them. Getting that side wrong wastes money and can widen the very gaps the investment was meant to close.

$42.45B	**$21B**	**+4.7 pp**	**$20,800**
BEAD deployment funds now moving to construction	unallocated BEAD restructuring savings awaiting NTIA decision	participation gain for high-search-cost workers; ≈0 for the general population	federal cost per job under CAF II, vs. $92,000 per job-year for ARRA

Five Things Policymakers Need to Know

1. **Infrastructure returns are heterogeneous, not uniform.**
The same wire produced no measurable participation effect for the general population and a gain of up to 4.7 percentage points for workers facing high job-search costs. Averages conceal the entire policy story.

2. **Broadband is a labor-market matching technology.**
The employment channel is not remote work but the collapse of search and information frictions. Benefits reach across the whole occupational distribution, including construction, food service, and production.

3. **Deployment without adoption creates two-tier access.**
When infrastructure arrived, better-off households upgraded while cost-constrained households took the cheapest tier and owned fewer devices. Coverage maps improved while the quality gap widened.

4. **Returns depend on local conditions.**
Effects concentrate where job seekers and vacancies coexist — in middle-income and urban markets with moderate to high unemployment — and vanish in the thinnest and poorest markets absent complementary investment.

5. **The policy window is open now.**
BEAD construction begins this summer under rules that stripped adoption, affordability, and workforce provisions, and roughly $21 billion in program savings awaits an allocation decision. That decision is where this evidence bites.

SECTION 01 · THE QUESTION THE DEBATE KEEPS SKIPPING

Public infrastructure economics has a blind spot, and it is the person at the end of the line

Three infrastructure debates are running in Washington at once. The $42.45 billion BEAD program has cleared its final approvals and begins construction this summer. Roughly $700 billion in private capital is flowing into AI data centers this year alone, provoking a bipartisan backlash over power costs, land, and who benefits. And Congress is weighing everything from data-center moratoria to federal preemption of state siting rules. All three debates share the same frame, which treats infrastructure as a supply problem measured in gigabits, gigawatts, coverage maps, and cost per passing.

That frame has a long pedigree. From rural electrification to the interstate highways, federal infrastructure policy has allocated dollars by where the gaps are. It is intuitive, auditable, and it fits on a map. But it embeds an assumption that our evidence shows to be wrong: that a connection is worth roughly the same to everyone who gets one. **The economic value of digital infrastructure depends on what constraint it relaxes for the specific people it reaches**, and those constraints vary enormously across the population. For a household that already searches, applies, and works online, a new wire is a convenience. For a worker whom the labor market has priced out of searching at all, it is the difference between being in the economy and out of it.

This brief presents causal evidence on that heterogeneity from the Connect America Fund Phase II (CAF II), the largest federal broadband investment before BEAD, and draws out what it means for three live policy decisions: how states spend their BEAD awards, how NTIA allocates roughly $21 billion in restructuring savings, and how the country should weigh the equity consequences of the broader AI-era buildout.

SECTION 02 · THE NATURAL EXPERIMENT

Geology decided who got connected first, and that lets us measure what connection did

Between 2015 and 2021, CAF II disbursed roughly $9 billion to extend broadband into high-cost areas. Crucially for research, eligibility and timing were set not by politics or applications but by an engineering formula, the Connect America Cost Model, built on terrain slope, soil composition, depth to bedrock, and loop length. These variables predict the physical cost of laying network plant and have no direct relationship to local labor-market trajectories. Two otherwise similar counties got broadband at different times because trenching was cheaper in one. That quasi-random variation, combined with 71 million person-year observations from the American Community Survey across 3,144 counties, lets us estimate what broadband deployment actually did, both overall and for the populations facing the highest barriers.

FIGURE 1

The same infrastructure, radically different returns

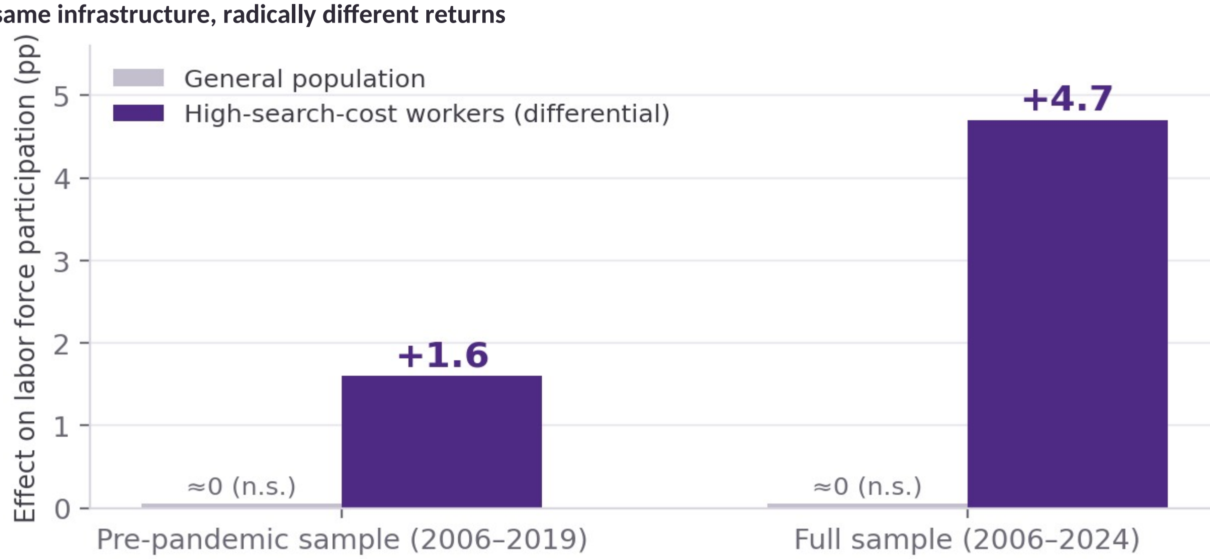

Differential effect of CAF II deployment on labor force participation (Post × CAF × Disability), relative to the general-population effect, which is statistically indistinguishable from zero in both samples. From a 37 percent baseline, the differentials represent 4.3 and 12.7 percent relative improvements. County-clustered standard errors are 0.2–0.3 percentage points; estimates are significant at the 1 percent level. Data: Yin and Guerrero (2026), Tables 2 and A8.

We focus on workers with disabilities as the identifying population not because the findings stop there, but because their elevated search costs are documented, measurable, and unambiguous. Transportation barriers shrink the geographic reach of in-person search, job postings almost never disclose accessibility or accommodation practices, and disclosing a disability cuts callback rates by 26 percent. If broadband works by dissolving search frictions, it should show up first and largest in this group — and it does.

Two benchmarks put the magnitude in context. Acemoglu and Angrist (2001) estimate that the Americans with Disabilities Act reduced employment of workers with disabilities by 0.8 percentage points in the years after enactment. A single infrastructure program moved participation by two to six times that magnitude in the opposite direction without mandating anything of employers. The effects also compound rather than spike: event-study estimates show flat pre-trends, a small effect at deployment, and steady growth through six to seven years while networks build out and households accumulate digital search capital. **Infrastructure returns are a stock that grows, not a stimulus that fades.**

SECTION 03 · THE MECHANISM

The gains come from easier search rather than remote work, and that changes who benefits

Why the mechanism matters is not academic. If broadband raises employment by enabling telework, benefits accrue narrowly to workers in remote-feasible occupations, who are disproportionately college educated and already within reach of jobs; policy should then target telework-amenable populations. If broadband instead works by making it cheap to find suitable employers — to see postings, screen workplaces, submit applications, and complete first interviews — then benefits extend across the entire occupational distribution and reach precisely the workers standard policy misses. Three tests separate the hypotheses, and all three point the same way.

FIGURE 2

Three tests reject remote work as the channel

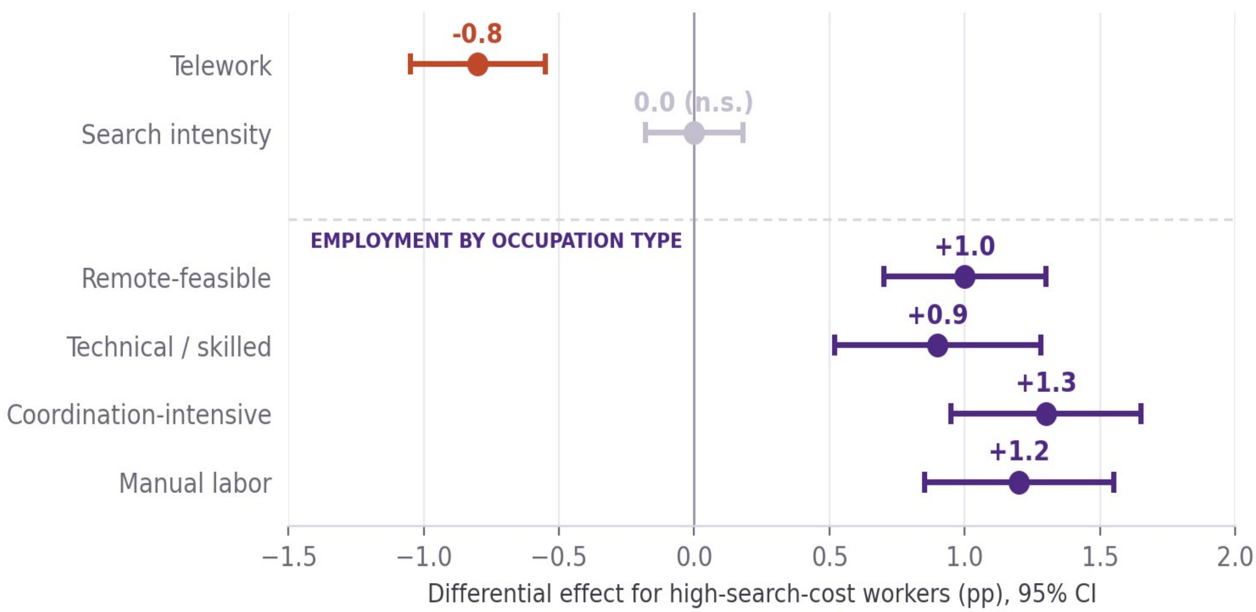

Differential effects (Post × Disability) with 95 percent confidence intervals. If remote work drove the gains, the telework differential would be positive and employment gains would concentrate in remote-feasible occupations. Instead telework fell relative to the general population, gains are uniform across all four occupation categories (Dingel and Neiman 2020), and search intensity did not move. Commutes fell only for workers with mobility-related impairments, which points to people finding jobs closer to home rather than working from home. Data: Yin and Guerrero (2026), Table 4.

Two auxiliary results sharpen the interpretation. Migration fell, meaning workers were not relocating toward opportunity but matching with employers already near them whom they previously could not identify. The response also came overwhelmingly from labor force entry rather than from the unemployed searching harder: the participation differential is roughly three times the conditional employment differential, and induced workers moved disproportionately into full-time positions, with 1.5 percentage points more full-time work and 1.8 percent more hours. The binding constraint was never willingness or capability but the cost of finding out where one could work.

This generalizes beyond disability

> The identifying population is workers with disabilities because their search costs are documented and exogenous to the geology that timed deployment. But the mechanism — information about suitable employers as the binding constraint — applies to workers re-entering after incarceration, immigrants without local professional networks, displaced workers in unfamiliar sectors, single parents screening for scheduling flexibility, and older workers screening for age-inclusive employers. Together these groups are a substantial share of the non-employed working-age population. We frame these extensions as testable predictions, but the policy logic is the same: wherever the constraint is informational, infrastructure that cuts information costs is employment policy.

SECTION 04 · THE TWO-TIER ADOPTION PROBLEM

Households came to the new network, but not to the same one

One finding should most unsettle a supply-side view of infrastructure. Deployment raised high-speed broadband subscriptions by about 1.2 percentage points for the general population. For households including a worker with a disability, the high-speed response was 0.9 points lower, which nets out to essentially zero. Those households did respond — by adopting entry-tier, low-speed service, precisely the opposite of the general population, which was migrating off it. The same trench and the same fiber produced a two-tier outcome sorted by household budget, compounded by a device gap, since these households were also markedly less likely to own a computer at baseline.

FIGURE 3

The two-tier response: deployment sorted households by budget

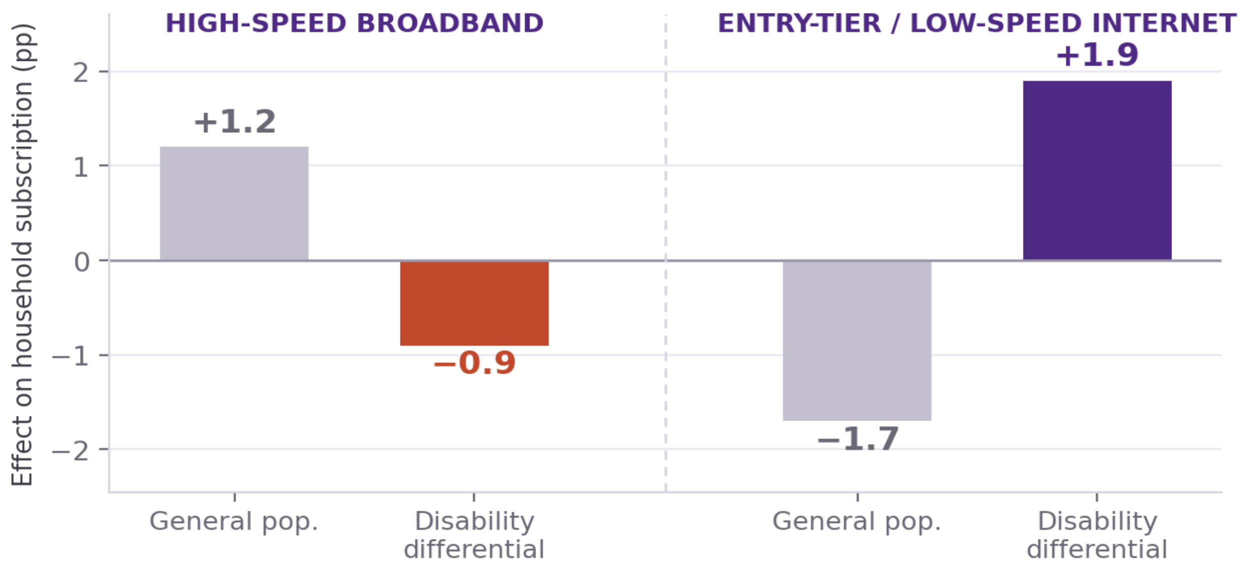

Effects of CAF II deployment on household internet subscription by service tier. The general population upgraded toward high-speed service and abandoned entry-tier connections, while households including a worker with a disability did the reverse. Entry-tier figures average effects across pre-program county speed terciles. The negative high-speed differential concentrates where faster tiers already existed, which reflects a price response rather than an infrastructure one. Data: Yin and Guerrero (2026), Tables 2, 3, and A3.

Two implications follow for policy. The first is that the labor market gains documented above ran through *any* connectivity rather than premium speed, so even modest adoption support has high returns for constrained households. The second, and the more troubling one for the current buildout, is that infrastructure spending that ignores affordability and devices mechanically widens the quality-of-access gap even as it narrows the coverage gap. The June 2025 BEAD restructuring removed the program's affordability and low-cost-plan requirements and its adoption and workforce provisions in the name of cost savings. Our evidence says those were not peripheral requirements. They were the part of the program that determines whether the marginal dollar reaches the households where its return is highest.

SECTION 05 · THE GEOGRAPHY OF RETURNS

A wire needs a labor market around it

Search-friction reduction creates jobs only where there are matches waiting to be made, with workers seeking and vacancies unfilled. The heterogeneity across local labor markets follows exactly that logic, and it complicates the politics of both rural-first broadband and place-based industrial policy.

FIGURE 4

Returns concentrate where matches can form

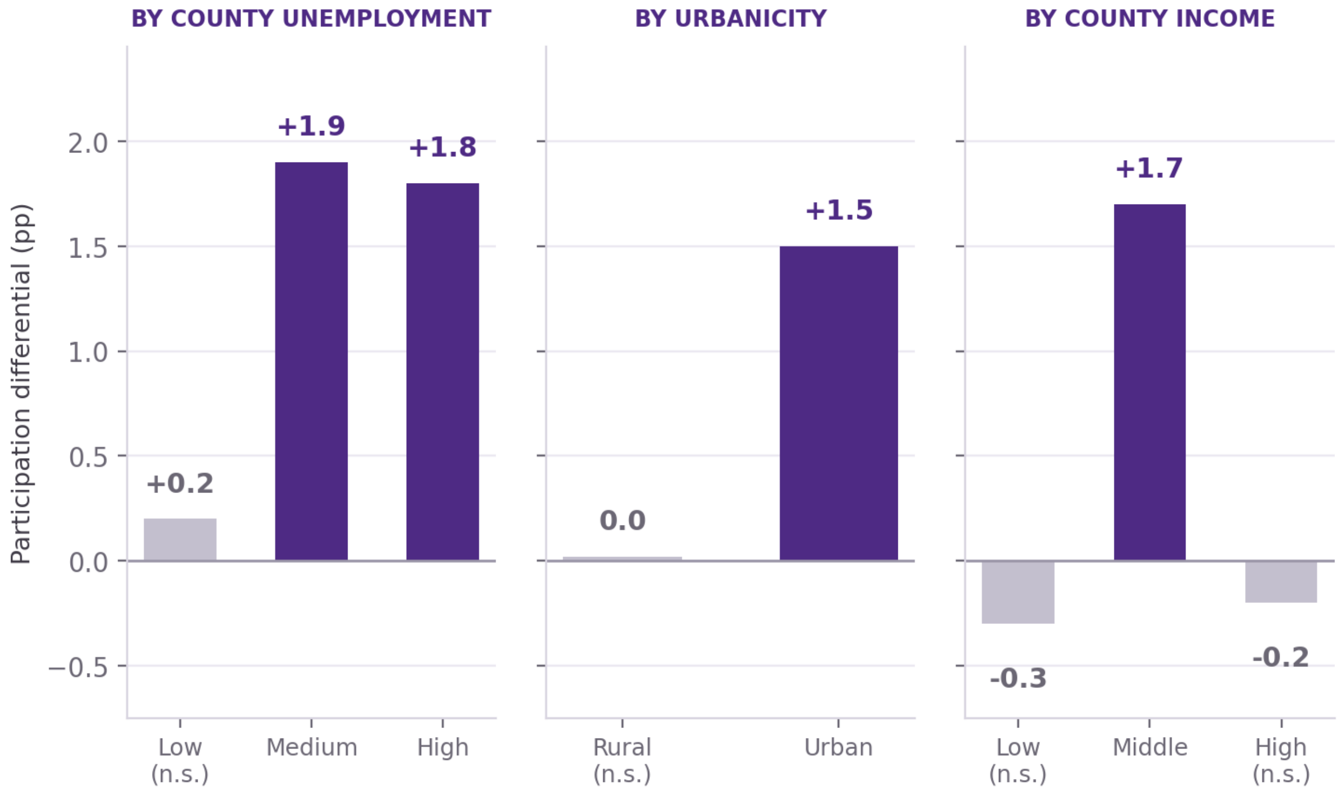

Differential participation effects within subsamples. Effects are large and precise in medium- and high-unemployment counties, urban labor markets, and middle-income counties — the places that combine connectivity gaps, employer density, and slack. They are indistinguishable from zero in tight markets, thin rural markets, and both tails of the income distribution. Gray bars are not statistically significant. Data: Yin and Guerrero (2026), Table 6.

This is not an argument against rural investment: the equity case for universal service stands on its own, and both CAF II and BEAD rightly prioritize unserved places. It is an argument that employment returns are a separate objective with a separate geography, and that programs claiming job-creation benefits should be honest about where those benefits will and will not materialize. In low-income and thin rural markets, connectivity is necessary but nowhere near sufficient. Without complementary investment in training, matching institutions, and local employer density, the wire reaches a labor market too small or too constrained to respond.

SECTION 06 · WHAT A DOLLAR BUYS UNDER TWO ALLOCATION RULES

The cheapest jobs program the federal government has run lately was a broadband program

The aggregate numbers are striking even before targeting. Applying the estimated employment effects to the roughly 120 million working-age adults in treated counties, CAF II generated approximately 432,000 jobs at about $20,800 in federal outlay per job (192,000 jobs at $46,900 under the conservative pre-pandemic estimates). The Council of Economic Advisers benchmark for the 2009 Recovery Act was roughly $92,000 per job-year. Because the calculation counts only wages and excludes consumer surplus, firm productivity, telehealth and education benefits, and the twenty-year service life of the asset, it is a floor.

FIGURE 5

Federal cost per job for CAF II against the stimulus benchmark

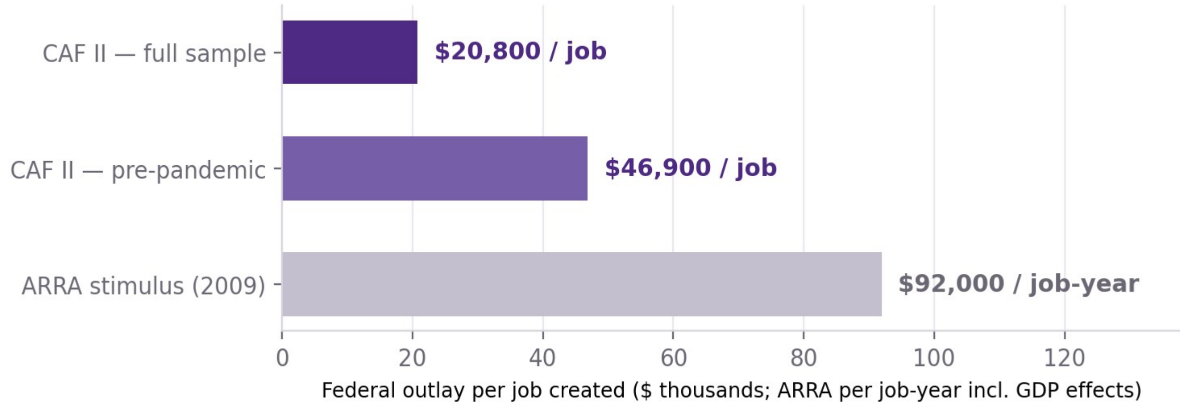

The CAF II cost per job applies estimated employment effects from Tables 2 and A8 to the treated working-age population against the $9 billion outlay. Direct comparison is imperfect because ARRA measured job-years including multiplier effects, but the order-of-magnitude gap survives any reasonable adjustment. The figure excludes non-labor-market benefits and is a lower bound. Data: Yin and Guerrero (2026), Table A15, and CEA (2009).

Targeting changes the arithmetic further. We simulate allocating a BEAD-sized budget of $42.45 billion across counties under two rules. The current connectivity-gap rule sends money where the unserved people are; a search-cost-weighted rule sends the same money where the unserved *and search-constrained* people live. Holding per-connection cost fixed, so that the difference reflects demand-side targeting alone, the second rule generates thousands of additional jobs and cuts the cost per job by nearly six percent — all from a change in the allocation formula that costs nothing.

TABLE

Two allocation rules for a $42.45B budget

ALLOCATION RULE	JOBS GENERATED (ANNUAL)	JOB-YEARS (20-YR LIFE)	COST PER JOB-YEAR
Connectivity-gap rule (current approach)	63,250	1,265,000	$33,557
Search-cost-weighted rule	**67,127**	**1,342,540**	**$31,619**
Gain from demand-side targeting	*+3,877 (+6.1%)*	*+77,540*	***−5.8%***

Both rules connect about 15.2 million previously unserved people at a constant $2,800 per connection across 3,033 counties with a positive December 2015 connectivity gap. Jobs are computed from disability-share-weighted employment effects of 0.018 and 0.002. Figures are undiscounted and conservative because event-study effects compound over time. Data: Yin and Guerrero (2026), Table 7.

SECTION 07 · THE POLICY WINDOW IS OPEN RIGHT NOW

The evidence arrives at exactly the moment the money moves

This would be an academic exercise if the allocation questions were settled, but they are not. As of this month, 54 of 56 states and territories have approved BEAD final proposals and construction begins this summer. The June 2025 Restructuring Policy Notice re-scored awards on lowest cost alone and removed the affordability, low-cost-plan, workforce development, and adoption requirements, which together formed the demand-side architecture. That restructuring produced roughly $21 billion in projected savings, and NTIA, after an early-2026 listening session, is now deciding among uses that explicitly include supplemental deployment, middle-mile investment, and digital adoption and workforce programs. Meanwhile the political system is consumed by the supply side of the AI buildout, with its moratoria, siting fights, and ratepayer protests, while the demand-side question of who can use any of this infrastructure goes largely unasked.

The risk in plain terms

A buildout that funds deployment while stripping adoption support will connect the map and miss the households. Our two-tier adoption evidence shows this is not hypothetical. Absent affordability and device support, the households with the highest marginal returns to connectivity systematically end up on the lowest tier of it or off it entirely. As hiring, benefits, education, and government services migrate onto AI-mediated platforms, the penalty for being on the wrong side of that line compounds annually. A buildout designed around supply alone does more than underdeliver — it actively finances divergence.

SECTION 08 · POLICY ACTIONS

Recommendations

1. Direct a substantial share of the ~$21 billion in BEAD savings to demand-side programs — adoption, devices, affordability, and digital workforce capacity — weighted toward high-search-cost communities.

This is the single highest-leverage decision on the table. The restructuring removed demand-side requirements from the deployment grants, and the savings are the vehicle to restore that function where the evidence says returns are largest.

2. Add demand-side weights to any supplemental deployment rounds.

Where remaining funds go to construction, weight allocation by the concentration of search-constrained residents alongside connectivity gaps. Our counterfactual shows this raises jobs per dollar by about 6 percent at zero fiscal cost, and because effects compound, the long-run gain is larger.

3. Require states to track adoption, tier, and device access, not just passings.

Availability data cannot detect two-tier adoption. Program accountability should measure subscription rates, service tier, and device access by income and disability status, so divergence is visible before it hardens.

4. Wire the matching institutions into the buildout.

The gains run through better matches. Vocational rehabilitation, American Job Centers, community colleges, and reentry programs are what convert a connection into a placement, and they should be funded and coordinated as part of infrastructure policy rather than adjacent to it.

5. Be honest about where employment returns will and will not appear.

Universal service is a legitimate equity goal independent of jobs. But job-creation claims should be conditioned on local labor-market thickness, and the thinnest markets should get the complementary investments in training and employer development that make connectivity convert.

6. Apply the demand-side lens to the AI buildout writ large.

The distributional consequences of compute, connectivity, and AI-mediated services are set by who can access and use them. Public policy toward private AI infrastructure — incentives, siting, and benefit agreements — should price in adoption and access, not only construction.

SECTION A · METHODS APPENDIX

Data, Identification, and Robustness

A. Setting and identifying variation

The Connect America Fund Phase II disbursed about $9 billion between 2015 and 2021 to extend broadband to census blocks where existing speeds fell below FCC thresholds and commercial deployment was deemed non-viable. Eligibility was determined mechanically by the Connect America Cost Model (CAM), an FCC engineering model predicting deployment cost from terrain slope, elevation, surface ruggedness, soil composition, depth to bedrock, vegetation density, and loop length. Blocks with modeled cost between $52.50 and $198.60 per location per month received support. Carriers accepted state-level offers, could not cherry-pick blocks, and faced mandatory milestones of 40, 60, 80, and 100 percent of locations by 2017 through 2020, monitored through the HUBB portal. Because CAM inputs predict construction cost but bear no direct relationship to labor-market trajectories, deployment timing is plausibly orthogonal to local economic conditions conditional on county and year fixed effects. Approximately 81 percent of the 3,144 sample counties were treated; the 19 percent never treated serve as the comparison group.

B. Data

Outcomes and demographics come from ACS microdata for 2006 to 2024, covering ages 16 to 64 outside group quarters — about 71 million person-year observations. The six-item ACS disability module, stable since 2008, identifies hearing, vision, cognitive, ambulatory, self-care, and independent-living difficulty and covers about 10 percent of working-age adults. Labor outcomes include participation, employment, hours, full-time status, earnings, telework, search methods among the unemployed, commute time, and occupation grouped following Dingel and Neiman

(2020). Technology outcomes for 2013 to 2024 include internet subscription of any type, high-speed broadband subscription, and computer ownership. Treatment timing comes from HUBB administrative filings geocoded to counties; a county is treated in year t if the first CAF II deployment occurred by the end of t. Geography uses a PUMA-to-county crosswalk on 2020 boundaries with population reweighting, and results replicate on the one-to-one matched subsample at 3.9 against 4.6 percentage points.

C. Estimation

The baseline is a reduced-form difference-in-differences with a triple-difference for the population differential: $Y(ict) = \delta\cdot(Post(ct) \times CAF(c) \times Disability(i)) + \beta_1\cdot Post(ct) + \beta_2\cdot CAF(c) + \beta_3\cdot Disability(i) + X(i)\cdot\gamma + \alpha(c) + \lambda(t) + \varepsilon(ict)$, with controls for age, education, race, sex, and marital status, county and year fixed effects, county-clustered standard errors, and county and person weights. The design is reduced-form on the deployment shock rather than 2SLS on subscription, so coefficients are intent-to-treat with no per-unit-broadband interpretation, and the CAM-based rollout supports parallel trends rather than an exclusion restriction. Dynamics use heterogeneity-robust event-study estimators from Borusyak, Jaravel, and Spiess (2024), de Chaisemartin and D'Haultfœuille (2023), and Sun and Abraham (2021), alongside TWFE, addressing negative-weighting bias under staggered timing (Goodman-Bacon 2021). Pre-treatment coefficients are economically small and statistically indistinguishable from zero across outcomes and estimators.

D. Robustness

A placebo permutation with 500 random reassignments of rollout timing centers at zero, and the true estimates of 0.047 for participation and 0.016 for employment lie in the extreme right tail with a permutation p value below 0.01. Balance tests regress deployment year on 2010–2014 changes in county characteristics and show predominantly small coefficients; residual differences are absorbed by controls and tested through pre-trends. Across specifications the participation differential ranges from 4.0 to 4.7 percentage points, from no controls through county-by-year fixed effects. The estimates are also robust to trimming density extremes (3.4 points), excluding sensory impairments (4.6), and restricting to ambulatory limitations alone (4.3). County-level five-year ACS estimates replicate with attenuated magnitudes of 1.3 points and flat pre-trends. The differential survives controls for SSDI, SSI, and public assistance receipt and for county industry location quotients. The pre-pandemic sample yields 1.6 percentage points, which we treat as the conservative causal benchmark, and gains persisted after 2019 rather than reverting.

E. Cost-benefit and targeting assumptions

Jobs figures apply employment effects of 0.2 percentage points for the general population and 1.8 total for workers with disabilities (0 and 1.6 before the pandemic) to about 120 million working-age adults in treated counties with 10 percent disability prevalence, against the $9 billion outlay. Wage gains use pre-treatment mean earnings of $37,900 and $28,700. The targeting counterfactual allocates $42.45B at a constant $2,800 per connection across 3,033 counties with a positive Dec-2015 connectivity gap, comparing allocation proportional to unserved population against unserved population weighted by disability share; a twenty-year service life converts to job-years, undiscounted. All calculations exclude consumer surplus, firm productivity, and non-labor benefits and hold treatment effects constant despite evidence of compounding — both choices bias toward understatement.

F. Limitations

Mechanism identification relies on differential responses across populations with different baseline search costs, and whether the small general-population effects operate through the same channel remains open. Heterogeneity by local conditions is descriptive, not experimentally identified. The household margin relevant to the differential response is connectivity at any tier, not high-speed broadband specifically. Extensions to other search-constrained populations are testable predictions, not findings of this design.

SUGGESTED CITATION

Yin, M., & Guerrero, D. (2026). *The Demand Side of Digital Infrastructure* (Policy Brief No. 2026-04). RISEI Lab, Northwestern University. The brief summarizes Yin and Guerrero (2026), *Connecting the Disconnected*, a working paper on the heterogeneous returns to broadband infrastructure investment. Additional sources: Acemoglu and Angrist (2001); Ameri et al. (2018); CEA (2009); Dingel and Neiman (2020); NTIA BEAD Restructuring Policy Notice (June 2025) and Progress Dashboard (July 2026).

DISCLAIMER

The contents of this paper were developed under grant Disability Innovation Fund (DIF 84.421E) from the U.S. Department of Education (Department). The Department does not mandate or prescribe practices, models, or other activities described or discussed in this document. The contents of this manuscript may contain examples of, adaptations of, and links to resources created and maintained by another public or private organization. The Department does not control or guarantee the accuracy, relevance, timeliness, or completeness of this outside information. The content of this manuscript does not necessarily represent the policy of the Department. This publication is not intended to represent the views or policy of or be an endorsement of any views expressed or materials provided by any Federal agency.

ABOUT RISEI LAB

Research and Innovation for Social and Economic Inclusion

The RISEI Lab at Northwestern University, led by Michelle Yin, is a multidisciplinary research and innovation hub focused on education, workforce development, and inclusive economic mobility. The lab conducts rigorous, policy-relevant research using large-scale data and causal methods, and designs and evaluates interventions in partnership with federal and state agencies, with a focus on how policy, systems, and emerging technologies shape pathways to education and employment for individuals with disabilities and other underserved populations.

michelle.yin@northwestern.edu · michelleyin.org · riseilab.org

School of Education and Social Policy · Northwestern University · Annenberg Hall, 2120 Campus Drive, Evanston, IL 60208