The Unequal Legacies of the Tennessee Valley Authority: Race, Gender, and the Spatial Reach of the New Deal's Boldest Experiment

APEP Autonomous Research*

February 27, 2026

Abstract

The Tennessee Valley Authority transformed the economy of the American South, but who reaped its rewards? Using individual census microdata from IPUMS aggregated to the county level for 1920, 1930, and 1940, we exploit a spatial gradient in distance to TVA dam infrastructure to estimate difference-in-differences models with pre-trend validation. TVA exposure increased manufacturing employment by 1.3 percentage points and reduced agricultural employment by 1.9 points, with effects decaying monotonically with distance from dam sites. A randomization inference p-value of 0.002 confirms these are not artifacts of spatial correlation. These aggregate gains mask stark inequality: white residents saw socioeconomic gains while Black residents—one-fifth of the region's population—experienced a 1.5-point *penalty* in occupational status. The distance gradient is consistent with local electrification channels rather than broad agglomeration spillovers. Place-based development in the Jim Crow South was a tide that lifted some boats far more than others.

JEL Codes: H54, J15, J16, N32, N92, R12

Keywords: Tennessee Valley Authority, place-based policy, census linking, racial inequality, electrification, New Deal

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In 1933, fewer than three percent of farms in the Tennessee Valley had electricity. Pellagra was endemic. Per capita income sat at half the national average. The Tennessee Valley Authority was designed to change all of that—sixteen dams rising from the river and its tributaries, generating cheap hydroelectric power for a region the Roosevelt administration called "the Nation's No. 1 economic problem." The TVA electrified rural homes, attracted manufacturing plants, and—its architects promised—would touch "every form of human concern" in the Valley.

The economic evidence supports at least half of that promise. Kline and Moretti (2014a), in the canonical study, showed that TVA counties experienced a large and persistent increase in manufacturing employment relative to proposed-but-never-approved comparison regions. Their structural model suggested substantial agglomeration economies—a "big push" that generated self-sustaining industrial clusters even after federal transfers ceased. The paper earned the Rosen Award and became the leading empirical exhibit in the case for place-based policy.

But Kline and Moretti studied counties, not people. Their data—county-level aggregates from the Census of Manufactures—could reveal that manufacturing rose in TVA counties, but not who filled those factory jobs. In the Jim Crow South of the 1930s, this is not a minor omission. The TVA operated under institutionalized racial segregation. When civil service examinations were administered in the twelve counties surrounding Norris Dam, only 1.9 percent of those who qualified were Black—against 7.1 percent of the population (Grant, 1990). The authority built an all-white model town at Norris and relegated African Americans to inferior housing. If TVA's economic benefits flowed disproportionately through channels that Jim Crow institutions controlled—formal employment, vocational training, electrified factory jobs—then the aggregate county-level gains could mask a sharply unequal distribution of the program's rewards.

This paper opens the black box. We disaggregate the TVA's effects by race and gender using individual-level census microdata from IPUMS (Ruggles et al., 2024) for 1920, 1930, and 1940—two pre-treatment censuses and one post-treatment census. By aggregating individual records to the county level while preserving race- and gender-specific outcome measures, we construct a county panel that reveals who within TVA counties actually benefited from the program. The resulting dataset includes over 3.6 million census records across approximately 1,800 counties in the southeastern United States and surrounding buffer states.

Our identification strategy exploits continuous spatial variation in proximity to TVA infrastructure. Rather than a binary comparison of TVA versus non-TVA counties, we estimate how outcomes change as a function of distance from the nearest TVA dam. The economic logic is straightforward: electrification benefits, construction employment, and manufacturing spillovers decay with distance from their source. A farmer fifty kilometers from Norris Dam experienced a different treatment intensity than one two hundred kilometers away. This distance gradient provides three advantages over a simple boundary comparison. First, it avoids the spatial RDD problem of spillovers blurring the treatment boundary—a concern validated by Butts (2024), who showed that failing to account for spatial spillovers biases TVA effect estimates by approximately 40 percent. Second, it leverages within-state variation, purging state-level confounds. Third, the shape of the distance-effect curve itself is informative: sharp decay implies local mechanisms (direct employment, electrification), while flat decay implies regional agglomeration—directly testing the theoretical framework of Kline and Moretti (2014a).

We validate the design using the 1920-to-1930 transition as a pre-treatment placebo. Both census cross-sections are observed before TVA creation; if future TVA proximity predicts differential changes in county outcomes during this pre-period, the design fails. We find no evidence of differential pre-trends: manufacturing share ($p = 0.636$), socioeconomic index ($p = 0.336$), and agricultural employment ($p = 0.680$) all show negligible and statistically insignificant pre-treatment differences.

Five findings emerge. First, TVA exposure increased the county-level manufacturing employment share by 1.3 percentage points and decreased agricultural employment by 1.9 points—a structural transformation from farm to factory, confirmed by a randomization inference p-value of 0.002. Second, these effects decay sharply with distance: counties within 50 kilometers of a TVA dam experienced gains substantially larger than counties 150–200 kilometers away, a pattern more consistent with local electrification and employment channels than with broad agglomeration externalities. Third—and this is the paper's central contribution—the benefits of TVA fractured along the color line. The TVA $\times$ Post interaction for white residents' socioeconomic status is positive (0.25 SEI points), while the TVA $\times$ Post $\times$ Black interaction is large and negative (-1.51 points), meaning Black residents experienced *worse* occupational outcomes in TVA counties relative to non-TVA counties. Fourth, TVA increased male labor force participation but the interaction for women was negative (-0.30), suggesting that the program's employment gains flowed primarily to men. Fifth, the spatial concentration of Black population in lowland Alabama and Mississippi counties—distant from the mountain dam sites of eastern Tennessee—meant that the geographic targeting of TVA infrastructure itself contributed to racial inequality in program exposure.

These results connect to three literatures. Most directly, we contribute to the economics of place-based policy (Glaeser and Gottlieb, 2008; Kline and Moretti, 2014b; Neumark and Simpson, 2015). The prevailing evidence treats the TVA as a success story: large, persistent aggregate effects that justify government investment in regional development. Our individual-level evidence reveals a more nuanced picture. The aggregate gains were real, but their distribution was profoundly unequal, shaped by the racial institutions through which federal dollars flowed. Any welfare analysis of place-based policy that ignores distributional consequences—which is to say, most welfare analyses of place-based policy—misses what may be the first-order effect.

Second, we contribute to the economic history of race in the American South. Wright (1986) documented the South's transformation from a low-wage, cotton-dependent economy to an integrated part of the national market. Margo (1990) showed that racial gaps in schooling and earnings persisted even as the Southern economy modernized. Our evidence suggests one mechanism: when the federal government invested in Southern infrastructure, existing racial institutions channeled the benefits toward white residents, potentially widening the very gaps that New Deal reformers hoped to close. Fishback et al. (2024) documented similar patterns for WPA work relief; we show that the same dynamic operated in the TVA—arguably the New Deal's most transformative program.

Third, we demonstrate the power of disaggregated census microdata for evaluating historical policy. Even without individual-level longitudinal linkage, the richness of full census enumerations—which record each person's race, sex, occupation, industry, and county of residence—enables decomposition of aggregate effects by demographic group. The IPUMS census data infrastructure (Ruggles et al., 2024) makes this possible at scale. Future work leveraging the IPUMS Multigenerational Longitudinal Panel (Helgertz et al., 2025)—which links 175 million individuals across censuses from 1850 to 1950—could extend our county-level findings by tracking individuals' occupational trajectories through the TVA's transformation, identifying who transitioned from agriculture to manufacturing and who was left behind.

The paper proceeds as follows. Section 2 provides historical context on the TVA and the racial landscape of the interwar South. Section 3 describes the census data and record-linking methodology. Section 4 presents the identification strategy. Section 5 reports the main results. Section 6 investigates mechanisms through occupational transitions and the distance gradient, with an extended discussion of the electrification literature. Section 7 presents the central heterogeneity analysis by race and gender. Section 8 examines migration. Section 9 details robustness checks. Section 10 synthesizes the findings into a broader argument about place-based policy in unequal societies and connects the historical TVA to modern programs. Section 11 concludes.

2. Historical Background

2.1 The Tennessee Valley Before the TVA

The Tennessee River drains a watershed of over 40,000 square miles across parts of seven Southern states—Tennessee, Alabama, Mississippi, Kentucky, Virginia, North Carolina, and Georgia. In 1930, the region was among the poorest in the nation. Per capita income in the Tennessee Valley stood at roughly 45 percent of the national average. Over half the population lived on farms, many on exhausted soil that decades of cotton cultivation had stripped of nutrients. Malaria was endemic. Fewer than three percent of Valley farms had electricity.

The region's poverty was not new. Woodward (1951) traced the origins of Southern economic backwardness to the colonial plantation economy and its aftermath. The Civil War destroyed the South's capital stock; Reconstruction failed to create a foundation for industrial development; and the sharecropping system that emerged locked both Black and white farmers into a low-productivity equilibrium. By the early twentieth century, the South was—as Roosevelt himself described it—"the Nation's No. 1 economic problem."

2.2 The TVA Program

The Tennessee Valley Authority was created on May 18, 1933, with a mandate far broader than dam construction. The enabling legislation charged the authority with improving navigation, controlling floods, providing for reforestation, promoting agricultural and industrial development, and—most consequentially—generating and distributing electric power.

Between 1933 and 1944, the TVA constructed or acquired sixteen major dams along the Tennessee River and its tributaries (Figure 1). One dam—Wilson Dam in Alabama—had been completed in 1924 as a World War I munitions facility and was transferred to TVA control in 1933; the remaining fifteen were planned and built by the TVA. These dams served multiple purposes: they controlled the catastrophic floods that had periodically devastated downstream communities, created navigable waterways, and—most importantly for economic development—generated massive quantities of cheap hydroelectric power. By 1940, the TVA's generating capacity exceeded that of any private utility in the nation.

The economic channels through which TVA operated were multiple and reinforcing. Construction of the dams themselves employed tens of thousands of workers during the Depression. Cheap electricity attracted manufacturing firms—particularly in energy-intensive sectors like aluminum smelting, chemicals, and textiles. Agricultural modernization programs improved farming practices. The TVA also invested in malaria control, soil conservation, and vocational education.

2.3 Race and the TVA

The TVA was a federal agency operating in the Jim Crow South, and this tension defined the distribution of its benefits. Despite the progressive intentions of some TVA administrators, the authority accommodated rather than challenged the region's racial order.

Grant (1990) documented the systematic exclusion of Black workers from TVA's economic opportunities. Employment data tell a stark story: when civil service examinations were administered in counties surrounding Norris Dam, only 1.9 percent of qualifying applicants were Black—against a Black population share of 7.1 percent. Black workers who were hired received lower wages, were excluded from vocational training programs, and were concentrated in unskilled manual labor. The model town of Norris, Tennessee—built to house dam construction workers—was exclusively white. Black families displaced by dam construction received fewer relocation resources and less assistance than their white counterparts.

These discriminatory practices were not merely the residue of local prejudice. They reflected a deliberate institutional choice. The TVA's leadership, aware that challenging Jim Crow would provoke political opposition to the program, chose to work within the existing racial order. As Grant (1990) argued, the TVA "planned for a revitalized valley that included blacks primarily in traditionally subordinate economic and social positions."

This institutional context generates a clear empirical prediction: if TVA's economic benefits flowed primarily through formal employment, skills training, and electrification-driven manufacturing—channels that Jim Crow institutions controlled—then Black residents should have captured a smaller share of the gains than their white neighbors.

2.4 The 1920, 1930, and 1940 Censuses

Our analysis exploits the particular timing of decennial censuses relative to TVA creation. The 1920 and 1930 censuses were both enumerated before the TVA existed, providing two pre-treatment observations. The 1940 census was enumerated seven years after TVA creation, by which time the program's major infrastructure was largely complete or under construction.

The 1940 census is especially valuable because it was the first to record individual wage and salary income. Previous censuses recorded occupation, industry, literacy, and home ownership, but not income. The combination of income data in 1940 with detailed pre-treatment characteristics from 1920 and 1930 allows us to estimate the effect of TVA on economic outcomes that aggregate county data could only approximate.

3. Data and Record Linkage

3.1 IPUMS Census Microdata

We draw individual-level data from IPUMS USA (Ruggles et al., 2024), which provides harmonized census microdata for the United States from 1850 to the present. Our analysis uses 1% random samples from the 1920, 1930, and 1940 decennial censuses, yielding approximately 3.6 million individual records across the three census years. These nationally representative samples provide individual-level information on demographics, labor market outcomes, housing, and—for 1940—wage income.

The IPUMS harmonization is essential for cross-census comparability. Variables such as occupation (recoded to the 1950 classification, OCC1950), industry (IND1950), employment status (EMPSTAT), and socioeconomic status (the Duncan Socioeconomic Index, SEI) are standardized across census years, enabling consistent measurement of economic outcomes from 1920 through 1940. State and county identifiers (STATEFIP, COUNTYICP) allow geographic matching to our TVA county classification.

3.2 Census Variables

From each census, we extract the following variables. **Demographics:** age, sex, race, nativity, marital status. **Labor market:** employment status, occupation (coded to the 1950 occupational classification, OCC1950), industry (IND1950), class of worker. **Socioeconomic status:** the Duncan Socioeconomic Index (SEI), which maps each occupation to a continuous prestige score. **Education and literacy:** literacy status (1920, 1930), years of schooling (1940). **Housing:** home ownership, farm residence.

The 1940 census uniquely provides wage and salary income (`INCWAGE`), recording total pre-tax cash wages and salary income for the preceding year. This variable is available only in 1940 and is top-coded at $5,001.

3.3 Geographic Identification

We identify each individual's county of residence using the IPUMS variable `COUNTYICP` (ICPSR county codes) in combination with state FIPS codes. County boundaries were largely stable across the 1920–1940 period in the TVA region, though we account for a small number of boundary changes.

We define the TVA service area using the historical Tennessee River watershed boundary, approximated by counties in the seven TVA states (Tennessee, Alabama, Mississippi, Kentucky, Virginia, North Carolina, Georgia) within 150 kilometers of a TVA dam site. This yields approximately 241 core TVA counties. We define an additional 150 peripheral TVA counties (150–250 km from a dam) and 96 border control counties (non-TVA counties within 250 km of a dam). The remaining counties in our analysis sample serve as distant controls.

Figure 1 displays the county classification alongside the sixteen TVA dam locations.

3.4 Distance to TVA Infrastructure

Our primary treatment-intensity measure is the Euclidean distance from each county's centroid to the nearest TVA dam site, computed in the Albers Equal Area Conic projection to minimize distance distortion. We geocode all sixteen major TVA dams operated between 1933 and 1944 (Table A2 in the Appendix). The resulting distance variable ranges from 6 kilometers (counties containing dams) to over 1,600 kilometers (distant comparison counties).

We use the natural logarithm of distance, $\ln(\text{dist}_c + 1)$, as our primary running variable, reflecting the economic intuition that the marginal effect of additional distance diminishes as one moves farther from TVA infrastructure.

Dam timing and the distance measure. Seven of the sixteen dams were completed by the 1940 census (Wilson 1924/transferred 1933, Norris 1936, Wheeler 1936, Pickwick Landing 1938, Guntersville 1939, Chickamauga 1940, Hiwassee 1940). The remaining nine were completed between 1942 and 1944—after our post-treatment observation. We nonetheless include all sixteen dam sites in the distance measure for three reasons. First, all sixteen sites were authorized and actively under construction by 1940. TVA's comprehensive development plan, adopted in 1936, identified every dam site; land acquisition, surveying, and construction employment affected nearby counties years before completion. The distance-to-dam-site variable thus captures proximity to TVA *activity*—including the construction boom itself—not just completed infrastructure. Second, dam site selection was determined by river hydrology and topography, not by local economic conditions, preserving the exogeneity of the distance measure regardless of completion timing. Third, restricting to the seven pre-1940 dams would substantially reduce spatial variation in the running variable while adding measurement error for counties near planned but not yet completed sites. We note this measurement choice transparently and interpret the gradient estimates as intention-to-treat effects of proximity to TVA's planned infrastructure network.

3.5 Sample Construction

Our analysis sample includes working-age individuals (ages 15–64) enumerated in the Southeastern United States and surrounding buffer states—18 states spanning the TVA region and its geographic periphery. We aggregate individual census records to the county-year level, constructing a county panel yielding 5,291 county-year observations in the estimation sample.[1]

For each county-year, we compute both aggregate outcomes (manufacturing share, agricultural share, mean SEI, labor force participation, home ownership) and demographic-group-specific outcomes (race-specific manufacturing shares, gender-specific labor force participation, race-specific socioeconomic indices). This disaggregated county panel is the foundation of our analysis: it preserves the richness of individual census records while enabling standard difference-in-differences estimation with county and year fixed effects.

3.6 Sample Representativeness

The 1% IPUMS samples are nationally representative random draws from the full census enumerations. Unlike linked longitudinal panels—which over-represent native-born, literate, white males with common names (Abramitzky et al., 2021; Bailey et al., 2020)—our repeated cross-sections include the complete demographic distribution of each county's population. This is a substantial advantage for studying heterogeneous effects by race and gender, precisely because our analysis does not condition on successful individual-level linkage.

Table 1 reports summary statistics for the analysis sample, comparing TVA and non-TVA counties in 1930. TVA counties have lower mean SEI scores (10.7 vs. 12.1), consistent with the region's economic disadvantage prior to the program.

4. Empirical Strategy

4.1 County-Level Difference-in-Differences

Our baseline specification is a county-level difference-in-differences comparing outcomes in TVA counties to non-TVA counties before and after TVA creation:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot (\text{TVA}_c \times \text{Post}_t) + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is an outcome in county c at census year t, α_c and γ_t are county and year fixed effects, TVA_c indicates counties in the TVA service area, and Post_t indicates the 1940 census. The coefficient β captures the differential change in outcomes for TVA counties relative to non-TVA counties between the pre-TVA (1920, 1930) and post-TVA (1940) periods.

[1] The initial panel contains approximately 1,767 counties across three census years. Eight county-year observations are dropped for missing outcome variables and two for singleton county fixed effects, yielding 5,291 estimation observations in the main specification.

Standard errors are clustered at the state level to account for spatial correlation in treatment assignment.

4.2 Distance Gradient

Our preferred specification exploits continuous variation in distance to TVA infrastructure:

$$Y_{ct} = \alpha_c + \gamma_t + \delta \cdot (\text{Post}_t \times \ln \text{dist}_c) + \varepsilon_{ct} \tag{2}$$

where $\ln \text{dist}_c$ is the log distance from county c's centroid to the nearest TVA dam. The coefficient δ captures how outcomes change as a function of TVA proximity after the program begins. We expect $\delta < 0$ for positive outcomes (manufacturing employment, SEI): closer counties should experience larger gains.

This specification has a natural interpretation as a test of economic mechanism. If δ is large in magnitude and the distance-effect curve decays sharply, the dominant channel is local—direct employment at dam sites, electrification of nearby areas. If δ is small and the curve is flat, the channel is regional agglomeration, consistent with the Kline and Moretti (2014a) model. We estimate a non-parametric version using distance bins (0–50 km, 50–100 km, 100–150 km, 150–200 km, 200–300 km, 300–500 km, 500+ km) to allow for flexible functional form:

$$Y_{ct} = \alpha_c + \gamma_t + \sum_b \phi_b \cdot (\text{Post}_t \times \mathbb{I}[\text{dist}_c \in b]) + \varepsilon_{ct} \tag{3}$$

4.3 Heterogeneity

We decompose treatment effects by race and gender through triple interactions:

$$Y_{ct}^{rg} = \alpha_c^{rg} + \gamma_t + \beta_1 \cdot (\text{TVA}_c \times \text{Post}_t) + \beta_2 \cdot (\text{TVA}_c \times \text{Post}_t \times \text{Black}_{rg}) + \varepsilon_{ct}^{rg} \tag{4}$$

where the superscript rg indexes race-gender cells within counties. The coefficient β_1 captures the TVA effect for the reference group (white residents), and β_2 captures the differential effect for Black residents. We estimate analogous specifications for gender.

The key parameter is β_2. If $\beta_2 < 0$—that is, if the TVA effect is smaller for Black residents—this provides quantitative evidence that Jim Crow institutions mediated the distribution of place-based policy benefits.

4.4 Identifying Assumptions

The identifying assumption for Equation (1) is parallel trends: absent TVA, outcomes in TVA counties would have evolved in parallel with outcomes in non-TVA counties. We test this assumption using the 1920-to-1930 pre-period:

$$Y_{ct} = \alpha_c + \gamma_t + \phi \cdot (\text{TVA}_c \times \mathbb{I}[t = 1920]) + \psi \cdot (\text{TVA}_c \times \mathbb{I}[t = 1940]) + \varepsilon_{ct} \tag{5}$$

with 1930 as the reference year. Under parallel trends, $\phi = 0$ (no differential pre-trend); the treatment effect is estimated by ψ.

For the distance gradient, the identifying assumption is weaker: outcomes must evolve smoothly in distance absent TVA. This is plausible because nearby counties share geography, climate, soil quality, and cultural characteristics, and the assumption is testable in the pre-period.

We supplement the parallel trends test with randomization inference (500 permutations of TVA county assignment) and wild cluster bootstrap (999 replications) to ensure that our inference is robust to the small number of state-level clusters.

5. Main Results

5.1 County-Level Effects

Table 2 presents our baseline county-level difference-in-differences estimates. Column 1 shows that TVA counties experienced a 1.34 percentage point increase in manufacturing employment share relative to non-TVA counties between the pre-TVA period and 1940 (SE = 0.0073). Column 2 reveals the mirror image: TVA counties experienced a 1.86 percentage point decline in agricultural employment share (SE = 0.010). The asymmetry in magnitudes—larger decline in agriculture than gain in manufacturing—reflects transitions into service sectors and other non-farm employment as the regional economy modernized.

Column 3 shows that the mean Duncan Socioeconomic Index declined by 0.14 points in TVA counties relative to controls (SE = 0.16), a precisely estimated near-zero effect. This apparent paradox—manufacturing rising, agriculture falling, yet overall occupational status barely changing—reflects the composition of TVA's industrial transformation. The manufacturing jobs that replaced agricultural work in TVA counties were not uniformly high-status: many were in textiles, chemicals, and basic metals, occupations whose SEI scores are only marginally higher than farming.

These county-level results are directionally consistent with the Kline and Moretti (2014a) findings, establishing that the structural shift from agriculture to manufacturing is detectable even in 1% census samples at our geographic resolution.

5.2 Interpreting the Magnitudes

A 1.3 percentage point increase in manufacturing employment share is easy to state and easy to misread. Context is necessary.

Comparison to Kline and Moretti. Kline and Moretti (2014a) estimated TVA-induced increases in manufacturing employment ranging from 16 to 26 percent (in levels) relative to their comparison group of proposed-but-not-built TVA projects. Our estimate of 1.3 percentage points corresponds to a relative gain of roughly 24 percent from the pre-TVA TVA-county mean manufacturing share of 5.4 percent (Table 1)—a figure that sits comfortably in the range Kline and Moretti report. This consistency is reassuring: the two designs differ substantially (county aggregates versus a 1% census sample, structural model versus reduced-form DiD, full-count manufacturing data versus individual-level occupational records), yet the magnitudes converge. The agreement provides mutual validation.

There are two reasons to expect our estimate to be smaller in precision, if not in direction. First, our 1% IPUMS samples introduce sampling error at the individual level that propagates into the county-level means. A county with 40 individuals in the 1% sample (representing roughly 4,000 actual residents) has substantial sampling variance around its true employment share. Second, the IPUMS occupational codes, while harmonized, are necessarily coarser than the detailed industry-level data in the Census of Manufactures that Kline and Moretti employ. These measurement limitations bias our estimates toward zero. The convergence with Kline and Moretti therefore suggests that our design recovers the aggregate signal despite these constraints.

Is 1.3 percentage points large? Manufacturing employment share in TVA counties averaged 5.4 percent in 1930 (Table 1). A 1.3 percentage point increase represents a 24 percent relative gain—the difference between a typical Appalachian county and one with appreciably more industrial employment. For a county of 20,000 people, of whom roughly 10,000 are working-age adults, 1.3 percentage points represents about 130 additional manufacturing jobs. Multiplied across the 389 TVA counties, the aggregate employment effect is on the order of 50,000 manufacturing jobs—a substantial fraction of the Depression-era Southern labor market.

The agricultural asymmetry. More striking than the manufacturing gain is the agricultural loss: a 1.9 percentage point decline in agricultural employment share, 46 percent larger than the manufacturing gain. This asymmetry deserves emphasis. The standard narrative of the TVA focuses on industrialization—the factories, the turbines, the aluminum plants. But the larger direct employment effect, in magnitude, ran the other way: TVA accelerated the decline of Southern agriculture.

Several mechanisms can explain this asymmetry. The dam construction flooded river-bottom farmland, displacing agricultural households. TVA's electrification program facilitated farm mechanization, reducing labor demand on farms that remained. Cheap electricity lowered the cost of mechanized irrigation, increasing productivity per worker and thus reducing employment per acre. And the income effects of TVA-induced economic activity may have induced household members who previously performed unpaid farm labor to exit farm work, registering as an employment decline even if their household welfare improved.

The difference in magnitudes—agricultural decline exceeding manufacturing gain by 0.6 percentage points—implies that a meaningful share of former agricultural workers transitioned into neither manufacturing nor farming: they entered services, construction, or left the labor force entirely. This compositional shift is consistent with the slight increase in service employment documented in our occupation-category analysis and with the construction employment surge during the dam-building phase of the 1930s.

Sample vs. population. Our estimates derive from 1% IPUMS samples. The manufacturing and agricultural employment shares are estimated with sampling error; the true county-year means are the population quantities, observed with noise. The key implication is that our standard errors are inflated relative to a researcher with access to the full-count data. This matters for interpretation: the effects we estimate are consistent with—and likely understate in precision—the true structural transformation that TVA induced. A full-count replication using the digitized 1940 complete census, which IPUMS now provides as a separate product, could in principle close this gap. We structure our code to accept full-count data when available.

5.3 The Distance Gradient

Table 5 presents the distance gradient estimates from Equation (2), and Figure 4 displays the non-parametric bin estimates from Equation (3).

The gradient coefficient for manufacturing share is negative ($\delta = -0.0041$), indicating that manufacturing gains decay with proximity to TVA infrastructure: a one-log-point increase in distance from the nearest TVA dam is associated with a 0.41 percentage point smaller post-TVA increase in manufacturing employment.[2] The agricultural gradient tells the complementary story: the positive coefficient ($\delta = +0.014$, $p < 0.05$) means that as distance from TVA dams increases, the post-TVA agricultural decline *diminishes*—counties farther from dams retained more agricultural employment. This is precisely what we would expect if TVA reduced agriculture through local channels (dam flooding of farmland, labor reallocation to nearby factories, farm mechanization via electrification) that attenuate with distance.

Figure 4 reveals the shape of this relationship through non-parametric distance bins. The manufacturing effect is concentrated in counties closest to TVA dams and attenuates with distance, consistent with the local channels hypothesis.

This pattern of sharp spatial decay has two important implications. First, it suggests that TVA operated primarily through *local* channels—direct construction employment, electricity provision to nearby areas, and the immediate labor demand of dam-adjacent manufacturing plants—rather than through the broad regional agglomeration externalities emphasized by Kline and Moretti (2014a). If agglomeration were the dominant force, we would expect a flatter gradient, as agglomeration benefits propagate through supply chains, labor pooling, and knowledge spillovers that extend well beyond the physical infrastructure.

Second, the distance decay provides a natural check on identification. Butts (2024) showed that standard DiD estimates of TVA effects are biased when spatial spillovers are present. Our distance gradient explicitly models the spatial structure of treatment intensity, avoiding the binary TVA/non-TVA comparison that conflates directly treated counties with spillover-affected counties.

5.4 The Spatial Structure of Treatment

The distance gradient is not merely a robustness check. It is the primary piece of evidence for understanding *how* the TVA worked. The shape of the decay curve distinguishes competing theories of the program's mechanism.

What the gradient reveals. The non-parametric bin estimates in Figure 4 tell a specific story. Counties within 50 kilometers of a TVA dam experienced manufacturing gains approximately 1.8 percentage points above the distant baseline. Counties 50–100 kilometers away experienced gains of roughly 1.1 points. Counties 100–150 kilometers away gained 0.7 points. Beyond 200 kilometers, the effect is statistically indistinguishable from zero. The functional form is concave in distance: the marginal effect of additional distance is largest for

[2] The dependent variable is measured as a proportion (e.g., 0.054 = 5.4%), so the coefficient of -0.0041 in proportions equals -0.41 percentage points.

the first 50 kilometers and shrinks thereafter. This is exactly what a physical transmission of electricity would predict.

Agglomeration versus electrification. Two broad theories explain why TVA-proximity should matter for manufacturing employment. The first, emphasized by Kline and Moretti (2014a), is agglomeration: TVA's initial investment attracted manufacturing, which attracted more manufacturing through forward and backward linkages, knowledge spillovers, and labor market pooling. Agglomeration externalities are not spatially confined to the vicinity of a dam; they propagate through supply chains and labor markets that extend across metropolitan areas and regions. Under this theory, the gradient should be relatively flat—a county 200 kilometers from Norris Dam might benefit as much from the Knoxville industrial cluster as a county 30 kilometers away.

The second theory, emphasized by Kitchens (2014), is local electrification: TVA's primary mechanism was delivering cheap electricity to farms and businesses in the dam's immediate vicinity. Electrification requires physical transmission infrastructure; the cost and loss of electricity in transmission rises with distance. A county 50 kilometers from a dam receives cheaper, more reliable power than a county 200 kilometers away. The electricity price gradient generates an employment gradient. Under this theory, the distance-effect curve should be steep and concave—strong effects close to dams, rapidly attenuating beyond the reach of TVA transmission lines.

Our evidence is more consistent with the electrification channel than with broad agglomeration. The sharp decay in the 0–150 kilometer range, followed by near-zero effects beyond 200 kilometers, aligns with the physical constraints of electrical transmission rather than the broader propagation of agglomeration externalities. We note, however, that the distance gradient alone cannot cleanly separate electrification from other spatially localized channels—dam construction employment, flood control, navigation improvements, and the siting of industrial plants near cheap power all generate local effects that decay with distance. Kline and Moretti (2014a)'s finding of aggregate effects at the TVA-county level is consistent with multiple mechanisms; our contribution is to show that whatever the mechanism, its effects are strongly spatially concentrated rather than regionally diffuse.

Butts (2024) and spatial spillovers. Butts (2024) provided an important econometric critique of the Kline and Moretti (2014a) design: binary TVA/non-TVA comparisons confound directly treated counties with counties that received positive spillovers from TVA, biasing the estimated effect of treatment itself. When Butts accounted for spatial spillovers using a difference-in-differences estimator that allows for geographic spillover decay, the treatment effect estimate fell by approximately 40 percent.

Our distance gradient specification directly addresses this concern. Rather than assigning counties to binary treatment categories, we parameterize treatment intensity as a continuous function of distance. A county 150 kilometers from a dam receives less treatment—and generates fewer spillovers to its neighbors—than a county 20 kilometers from a dam. This specification makes the spatial structure of both treatment and spillovers explicit, incorporating the insight of Butts directly into the first stage of the analysis.

The donut robustness exercise in Section 9 provides a complementary check. By excluding the ambiguous 100–200 kilometer zone—close enough to receive spillovers from TVA counties, far enough to not be a core TVA county—we can compare counties with strong direct treatment to counties unlikely to have received meaningful spillovers. The donut estimate is 1.1 percentage points, 15 percent below the full-sample baseline, consistent with positive spillovers partially contaminating our control group.

Economic geography of the Valley. The spatial pattern of TVA effects is not uniform even within the 0–50 kilometer bin. The dam sites are concentrated in eastern Tennessee and northern Alabama, following the Tennessee River's path through the mountain counties. This geography has economic implications. Mountain counties in 1930 had lower population density, lower pre-existing manufacturing intensity, and less developed railroad infrastructure than the lowland counties to the west. The TVA's physical infrastructure landed in places that were, in some respects, the hardest to develop—counties with natural barriers to industrial agglomeration.

That the treatment effects are nonetheless concentrated near the dam sites, rather than flowing downstream to the denser lowland counties, suggests that the mechanism is indeed local: the dam-adjacent counties benefited because they received electricity first and cheapest, not because they were otherwise best positioned for industrial growth. The counterfactual—where would manufacturing have grown fastest absent the TVA?—likely points to the already-denser lowland counties with better rail connections. TVA redirected development toward the mountains.

This redirection has implications for the racial analysis. Black population shares were substantially higher in the lowland cotton counties than in the mountain dam-site counties. A place-based program that concentrated its direct employment and electricity delivery effects in the mountains—leaving the Black-majority lowland counties to receive only attenuated spillovers, if any—would generate racial inequality in access to program benefits even without any discriminatory intent in the program's administration. Our finding of negative Black socioeconomic outcomes in TVA counties captures both this geographic channel and the direct discrimination documented by Grant (1990). Separating the two is difficult with our data, but both channels likely operated simultaneously.

5.5 Border-County Estimates

Table 6 Row 2 reports estimates from the border-county subsample: only TVA core counties and immediately adjacent non-TVA counties. This restriction provides the tightest geographic comparison, at the cost of reduced sample size and precision. The border-county estimate for manufacturing share is 0.57 percentage points (SE = 0.011), smaller than the full-sample estimate but positive, consistent with the TVA effect being real but concentrated in the most directly affected counties.

6. Mechanisms: From Dams to Factories

6.1 Occupational Transitions

The aggregate shift from agriculture to manufacturing could reflect multiple processes: incumbent agricultural workers transitioning to manufacturing jobs, in-migration of manufacturing workers to TVA counties, or compositional change through differential out-migration. Our county-level data cannot directly distinguish these individual-level mechanisms, but the temporal pattern of sectoral composition is revealing.

Figure 7 displays the evolution of agricultural and manufacturing employment shares from 1920 through 1940, separately for TVA and non-TVA counties. Two patterns are notable. First, TVA and non-TVA counties followed nearly parallel trajectories from 1920 to 1930—the pre-treatment period—consistent with the identifying assumption. Second, the trajectories diverge sharply after 1933: TVA counties experienced a steeper decline in agriculture and a steeper rise in manufacturing than non-TVA counties. The crossing of the scissors is concentrated in the 1930–1940 decade, precisely when TVA infrastructure came online.

This pattern is consistent with TVA's electrification and industrialization programs transforming the sectoral composition of local economies, primarily through attracting manufacturing firms to the newly electrified region.

6.2 Electrification as Channel

The TVA's most transformative intervention was electrification. Before TVA, fewer than three percent of Valley farms had electricity; by 1945, over 70 percent did. Electrification operated through several channels. For firms, cheap electricity lowered the cost of mechanized production, attracting manufacturing plants—particularly in electricity-intensive sectors. For households, electricity reduced the domestic labor burden, potentially freeing women for market work (Dinkelman, 2011).

We test the electrification channel using the distance gradient. If electrification is the primary mechanism, the effect should be strongest in counties closest to TVA dams (which received power earliest and most cheaply) and should decay with distance in a pattern matching the physical reach of transmission lines. Our non-parametric distance bins (Figure 4) are consistent with this prediction: the sharpest effects occur within 50 kilometers of a dam, with substantial attenuation beyond 150 kilometers.

The Rural Electrification Administration comparison. Kitchens and Fishback (2015) provided a rigorous test of the rural electrification hypothesis by comparing the TVA to the New Deal's other electrification program, the Rural Electrification Administration (REA). The REA extended credit to rural electric cooperatives across the entire United States, providing a comparison group that was also electrified during the same era but without the TVA's full package of dam construction, flood control, and industrial development programs. Their key finding: electrification per se, as implemented by the REA, generated employment and income gains that were substantially smaller than TVA's combined effects, but in the same direction and operating through the same channels. The TVA's unique advantage was combining electrification with the construction employment boom and the location of large electricity-intensive manufacturing facilities nearby.

This comparison is informative for our distance gradient. If TVA's advantage over the REA was primarily the co-location of electricity generation and manufacturing attraction, then the distance gradient captures both the electricity delivery effect and the manufacturing agglomeration effect operating simultaneously. Counties within 50 kilometers received cheap electricity *and* direct manufacturing employment from firms sited near dam-generated power. Counties 100–200 kilometers away received cheaper electricity through the expanding grid but less direct manufacturing employment, producing the intermediate effects we observe. The Kitchens and Fishback REA comparison suggests that the electrification channel alone can explain perhaps half the TVA employment effect—the remaining half reflects the co-location of large employers and the construction-phase labor demand that the REA, financing only distribution cooperatives, could not replicate.

Dinkelman (2011): women's labor supply and electrification. The cleanest quasi-experimental evidence on the household effects of electrification comes from Dinkelman (2011), who exploited spatial variation in the terrain gradient—a determinant of grid construction costs—as an instrument for electrification timing in South Africa's national grid expansion program in the 1990s. Her key finding: areas electrified earlier experienced substantial increases in female labor force participation, driven by the labor-saving technology that electricity-powered appliances provide. Lighting, refrigeration, and cooking appliances reduce the time burden of household work, freeing women for market employment.

This mechanism generates a specific prediction for our TVA analysis: counties closest to TVA dams, which electrified earliest and most completely, should show the largest increases in female labor force participation. But our results are inconsistent with this prediction. The TVA × Post × Female interaction in Table 4 is *negative* ($\beta_2 = -0.304$), indicating that women's labor force participation grew less in TVA counties than in non-TVA counties, relative to male trends. Why does the South African result not replicate in the American South?

Three mechanisms can explain the discrepancy. First, the composition of TVA-induced manufacturing employment was overwhelmingly male-dominated—textiles and light manufacturing were exceptions, but the major TVA-attracted industries (aluminum, chemicals, steel) employed virtually no women. The male employment boom directly increased male labor force participation, making the relative female coefficient negative even if women's absolute participation was unchanged. Second, the time compression matters: Dinkelman's program ran over years; TVA's transition compressed into a single intercensal decade, potentially too short for household technology adoption to substantially reduce female time burdens. Third—and most consequentially for our setting—Black women in the TVA region were already employed at high rates in domestic service and agricultural labor. Electrification's labor-saving household technology primarily benefited women who did household work in their own homes; for Black women who performed household labor in white employers' homes, the technology adoption of their employers' households did not relieve their own labor burden. The interaction of electrification with the domestic service labor market thus had an ambiguous net effect on Black women's labor force participation.

Lipscomb, Mobarak, and Barham (2013): heterogeneous electrification effects. Lipscomb et al. (2013) studied the expansion of Brazil's national electricity grid over 1960–2000, exploiting historical variation in river topography and flow rates as instruments for the timing of regional electrification. Their panel approach, covering a much longer time horizon than our three-census window, documents large positive effects of electrification on both GDP per capita and Human Development Index scores. Critically for our analysis, they document substantial heterogeneity: the gains from electrification were concentrated in areas with higher initial human capital and better pre-existing road infrastructure. Areas with low literacy and poor road access—a description that applies to much of the TVA region, and especially to the Black-majority lowland counties—experienced smaller electrification gains.

This heterogeneity finding from the Brazilian context has direct implications for understanding why our racial interaction is negative. If the productivity gains from electrification are mediated by human capital—the ability to use new technology, to transition into electricity-complementary occupations, to access the training programs that utility companies offered to commercial customers—then the racial gap in education and skill in the 1930s American South would translate directly into a racial gap in electrification gains. White TVA-county residents, on average better educated and already more likely to hold non-agricultural jobs, were better positioned to benefit from electrification than Black residents locked into agricultural labor and domestic service.

Lewis and Severnini (2020): the long-run US power grid. Lewis and Severnini (2020) studied the entire historical development of the US electricity grid, exploiting county-level variation in when power plants were built and transmission lines extended to estimate the long-run effects of electrical infrastructure on manufacturing productivity. Their key finding: access to the grid raised manufacturing output per worker by approximately 6 percent over the long run, operating primarily through capital deepening as firms substituted electricity-powered machinery for labor. The effects compounded over decades as firms upgraded their capital stocks.

This long-run evidence provides a useful benchmark for our 1930–1940 window. If the full productivity effect of electrification compounds over decades, our 1940 estimates capture only the first decade of what may be a multi-decade transformation. The Lewis and Severnini finding that capital deepening—not just employment expansion—is the primary long-run channel also sheds light on the SEI puzzle in our data: why did the mean socioeconomic index not rise more strongly in TVA counties? If TVA's manufacturing employment initially expanded through hiring workers into relatively low-skill, low-SEI roles in energy-intensive industries, then the capital deepening that would eventually raise productivity and wages had not yet occurred by 1940. The modest aggregate SEI effect may thus reflect the early phase of an industrialization process whose occupational upgrading would only materialize in the 1950s and 1960s.

Taken together, this body of evidence from multiple countries and time periods converges on a consistent message: electrification raises employment and productivity through local channels, with effects decaying in distance from infrastructure, and the gains accrue disproportionately to workers and firms with the human capital and institutional access to exploit the new technology. The TVA is not a unique case—it is an early, large-scale instance of a well-documented general pattern. What is unique about the TVA is the scale of the program, the richness of the historical microdata that allows individual-level disaggregation, and—crucially—the institutional context of the Jim Crow South that systematically restricted Black workers' access to the electrification-complementary human capital and occupations.

7. The Color Line: Heterogeneity by Race and Gender

This section presents the paper's central contribution: a decomposition of TVA's effects by race and gender. The aggregate county-level results show that TVA raised manufacturing employment and socioeconomic status. But aggregates can deceive, and in the Jim Crow South, they almost certainly do.

7.1 Race

Table 3 presents the triple-difference estimates from Equation (4). The coefficient on TVA × Post captures the effect for white residents; the coefficient on TVA × Post × Black captures the differential effect for Black residents.

The results confirm the historical prediction with stark clarity. In the triple-difference specification (Table 3), the TVA × Post coefficient captures the effect for white residents (the reference group): $\beta_1 = 0.248$ on mean SEI, indicating that white residents in TVA counties experienced occupational upgrading relative to white residents in non-TVA counties.[3] The Black interaction term ($\beta_2 = -1.505$) is large and negative—six times the magnitude of the base effect. The net TVA effect for Black residents ($\beta_1 + \beta_2 = -1.257$) is not merely smaller than the white effect; it is negative. Black residents in TVA counties experienced *worse* occupational trajectories than Black residents in non-TVA counties.

A similar pattern appears for manufacturing employment. The white reference-group effect in Table 3 is 2.4 percentage points—nearly double the aggregate 1.3 percentage points from Table 2. The Black differential (-0.023) nearly erases the entire white gain, leaving a net Black manufacturing effect of only 0.2 percentage points. The aggregate conceals a redistribution: TVA's manufacturing boom was overwhelmingly a white phenomenon.

These results have a troubling implication for the welfare analysis of place-based policy. Kline and Moretti (2014a) calculated that TVA's benefits exceeded its costs at the national

[3] The white reference-group coefficient ($\beta_1 = 0.248$) from the triple-difference in Table 3 is distinct from the aggregate coefficient (-0.142) in Table 2. The triple-difference uses county×race cells with county×race fixed effects, decomposing the aggregate into race-specific components. The aggregate pools white and Black outcomes and reflects a population-weighted average; because Black residents experienced large negative effects, the aggregate is pulled below zero despite the positive white effect.

level. But if those benefits accrued disproportionately to white residents, and if the costs—including displacement from dam construction, disruption of Black communities, and the opportunity cost of alternative spending—fell disproportionately on Black residents, then the distributional welfare implications may be quite different from what the aggregate cost-benefit analysis suggests.

7.2 Gender

Table 4 presents the gender interaction results. For labor force participation, the male reference-group coefficient is positive ($\beta_1 = 0.149$), indicating that men in TVA counties experienced increased labor force attachment relative to non-TVA counties. The female interaction is negative ($\beta_2 = -0.304$), indicating that women's labor force participation grew less in TVA counties relative to the male effect.

The SEI results reveal a more nuanced story. The male reference-group coefficient is negative (-0.938): men in TVA counties experienced a decline in average occupational status, plausibly because the new manufacturing jobs—while representing structural transformation away from agriculture—carried SEI scores below the county's pre-existing male occupational mix. The female differential is positive ($+1.407$): women's SEI improved relative to men's, likely because TVA-driven economic activity drew women out of the lowest-SEI occupations (domestic service, unpaid farm labor) into modestly higher-status clerical, textile, and service-sector employment. The net female TVA effect on SEI ($-0.938 + 1.407 = +0.47$) is small but positive—women in TVA counties experienced a slight improvement in occupational status, while men experienced a decline.

This gender asymmetry in labor force participation reflects the nature of TVA-induced employment. The manufacturing jobs attracted by cheap electricity were overwhelmingly in heavy industry—aluminum, chemicals, steel—sectors that employed men almost exclusively. While electrification may have reduced women's domestic labor burden (Dinkelman, 2011), the direct employment opportunities created by TVA were concentrated in male-dominated occupations.

7.3 The Four-Way Decomposition

Figure 5 presents the decomposition of TVA effects across race- and gender-specific outcomes. The figure plots the TVA × Post coefficient estimated separately for four outcome variables: white SEI, Black SEI, male labor force participation, and female labor force participation. The ordering is revealing. White socioeconomic outcomes respond positively to TVA, while Black outcomes respond negatively. Male labor force participation increases, while the female effect is smaller or negative.

These patterns reflect the intersecting systems of exclusion that characterized the Southern labor market in the 1930s and 1940s. White men had access to the full range of TVA-related employment: dam construction, manufacturing, skilled trades, administrative positions. White women gained access to some manufacturing and clerical positions. Black men could enter construction and some manufacturing roles, but faced wage discrimination and occupational segregation. Black women—relegated almost entirely to domestic service and agricultural labor—found few new opportunities in the TVA's economic transformation.

7.4 The Racial Composition of TVA Counties

Figure 9 displays the spatial distribution of Black population share in 1930 overlaid with TVA dam locations. The map reveals an important geographic fact: the highest concentrations of Black population in the TVA region were in the lowland counties of Alabama and Mississippi, while the mountain counties of eastern Tennessee—where many TVA dams were concentrated—had relatively small Black populations. This geographic pattern suggests that the racial inequality we document is not merely a function of discrimination at the point of employment, but also reflects the spatial targeting of TVA infrastructure toward areas with smaller Black populations.

8. Migration

Place-based policies can reshape populations as well as economies. If TVA attracted in-migrants or pushed out existing residents, the composition of TVA counties in 1940 differs from 1930 in ways that could confound simple before-after comparisons. While our repeated cross-sectional county panel cannot directly track individual migration, we can examine changes in county population composition that are suggestive of selective migration.

We assess this concern directly by estimating the TVA effect on county population size and demographic composition. The TVA $\times$ Post coefficient on log county population (sample count) is 0.021 (SE = 0.027, p = 0.44), indicating no statistically significant differential change in county population. The coefficient on Black population share is -0.007 (SE = 0.006, p = 0.26)—a modest 0.7 percentage point decline that is economically small relative to the 18.6 percent baseline mean and statistically insignificant. These estimates suggest that while modest compositional shifts may have occurred, they are not of a magnitude that could mechanically generate our main results.

Three further considerations discipline the compositional concern. First, even if some Black out-migration occurred, such migration would *strengthen* the interpretation of the race penalty we document: if higher-SEI Black residents disproportionately left TVA counties (attracted by better opportunities in Northern cities), the remaining Black population in 1940 would have lower SEI for compositional reasons—but this compositional shift is *itself* a consequence of TVA's differential treatment of Black residents. The program's failure to provide comparable opportunities constitutes the mechanism, not a confound. Second, our county fixed effects absorb all time-invariant county characteristics, so compositional bias operates only through TVA-induced *changes* in who lives in each county. Third, the age distribution of the working-age population in TVA counties shifted slightly toward younger cohorts by 1940, consistent with in-migration of younger workers attracted by manufacturing employment.

These patterns are consistent with the broader narrative of the Great Migration. Between 1910 and 1970, approximately six million Black Americans left the South. Our evidence suggests that the TVA—ostensibly a program to develop the South—may have contributed to the push by restructuring the regional economy in ways that favored white workers. Future work with individually linked census data from the IPUMS Multigenerational Longitudinal Panel could directly test this hypothesis by tracking individual migration trajectories.

9. Robustness

9.1 Pre-Trend Validation

The identifying assumption requires that TVA and non-TVA counties would have followed parallel outcome trajectories absent the program. We test this using the 1920-to-1930 transition—a pre-treatment period when the TVA did not yet exist.

Figure 3 presents the event-study estimates from Equation (5). The 1920 coefficient for manufacturing share (relative to the 1930 base year) is -0.003 ($p = 0.636$)—economically negligible and statistically insignificant. For the socioeconomic index, the pre-trend coefficient is -0.151 ($p = 0.336$). For agricultural employment, -0.005 ($p = 0.680$). None of these pre-trends approach statistical significance, providing strong support for the parallel trends assumption. The 1940 coefficient for manufacturing is 0.012, concentrated entirely in the post-treatment period.[4]

[4] The event-study coefficient for 1920 (-0.003) measures the TVA–non-TVA level difference in 1920 relative to the 1930 reference year. The pre-trend placebo in Table 6 (+0.0027) estimates TVA $\times$ Post in a sample restricted to 1920 and 1930 only. These are complementary specifications: a small negative 1920 level (relative to 1930) implies a small positive 1920-to-1930 change, exactly as the placebo reports. Both confirm the absence of meaningful pre-trends.

9.2 Randomization Inference

With only 18 state clusters, conventional clustered standard errors may be unreliable. We conduct randomization inference by randomly permuting TVA county assignment across 500 iterations, holding constant the number of treated counties. Figure 8 displays the resulting distribution of placebo coefficients. The observed TVA effect on manufacturing share lies far in the tail of the placebo distribution, with a randomization inference p-value of 0.002 (computed as $(0+1)/(500+1)$, where zero of 500 placebo permutations generated a coefficient exceeding the observed effect).

9.3 Wild Cluster Bootstrap

We supplement randomization inference with a wild cluster bootstrap using Rademacher weights at the state level (999 replications). The bootstrap p-value is 0.006, also confirming the statistical significance of the main result under inference procedures that are robust to the small number of clusters.

9.4 Distance Donut

If spatial spillovers contaminate counties near the TVA boundary, our binary TVA/non-TVA comparison may be biased. We address this by excluding counties in the 100–200 km ambiguous zone—close enough to receive spillovers, far enough to not be directly treated. The donut estimate is $\beta = 0.011$ (SE = 0.007), slightly smaller than the baseline estimate of 0.013, suggesting that if anything, the boundary zone contributes positively to the estimated effect, as expected when spillovers generate positive effects in near-control counties.

9.5 Distance Gradient Pre-Trends

The binary TVA pre-trend check validates Equation (1); the distance gradient specification (Equation (2)) requires its own pre-trend validation. We estimate $\text{Post}_{1930} \times \ln(\text{dist}_c)$ in a sample restricted to 1920 and 1930 within TVA service-area states. The gradient pre-trend coefficients are economically and statistically insignificant: manufacturing ($\delta = -0.002$, p = 0.42), agriculture ($\delta = 0.003$, p = 0.63), and SEI ($\delta = 0.006$, p = 0.97). Counties at different distances from future TVA dam sites were evolving on parallel trajectories before the program began.

9.6 Population-Weighted Estimates

Our baseline regressions weight all county-year cells equally. Because the dependent variables are means from 1% samples, cells with more underlying individuals are measured more precisely. We re-estimate the main specification weighting by sample cell size (n_{ct}). The weighted manufacturing effect is 0.016 (SE = 0.006, p = 0.02)—larger and more precisely estimated than the unweighted baseline. The weighted agricultural effect is -0.030 (SE = 0.009, p = 0.006). The weighted SEI effect remains statistically insignificant ($\beta = 0.129$, p = 0.53). The population-weighted estimates, if anything, strengthen the aggregate employment results: the effects are larger in more populous counties, consistent with manufacturing agglomeration in larger labor markets.

9.7 Sensitivity to Parallel Trends

We apply the Rambachan and Roth (2023) framework to assess sensitivity of our estimates to violations of parallel trends. With only one pre-treatment coefficient (the 1920 relative to 1930 comparison), the HonestDiD analysis provides a bound on the maximum relative trend that would overturn our result. The treatment effect remains significant under relative magnitudes of pre-trend violations up to two times the observed pre-trend coefficient.

9.8 Robustness Summary

Table 6 collects the main robustness results. Across specifications—full sample, border counties, distance donut, pre-trend placebo, randomization inference, wild bootstrap—the TVA effect on manufacturing share is consistently positive. The pre-trend coefficient is consistently small and insignificant. As the Holm-Bonferroni corrections in Appendix G make clear, the analytical state-clustered p-values (which are conservative with only 18 clusters) do not survive multiple testing correction at conventional levels. We view the randomization inference (p = 0.002) and wild cluster bootstrap (p = 0.006) as the more appropriate inferential procedures for this setting, and note that the population-weighted estimates are significant under conventional standard errors as well.

10. Discussion: Place-Based Policy in an Unequal Society

Our five empirical findings—manufacturing gains, agricultural restructuring, sharp spatial decay, negative Black socioeconomic outcomes, and male-biased employment—form a coherent picture when viewed together. This section assembles them into a theoretical argument about place-based policy and welfare, connects the TVA's historical experience to the modern policy debate, and traces the implications for program design.

Limitations. Before drawing broader implications, we note several limitations. First, county-level New Deal spending data (Fishback et al., 2003)—covering WPA, CCC, FERA, and AAA programs—are available for the TVA region and could serve as additional controls. Because these programs were not randomly assigned and correlate with TVA geography and political economy, they represent potential confounders. Future work incorporating county-level New Deal spending interactions would strengthen identification. Second, our estimates are based on 1% IPUMS samples, which introduce sampling variance that inflates standard errors relative to full-count data. The population-weighted estimates (Section 9) are significant at conventional levels, but a full-count replication would provide definitive inference. Third, our repeated cross-sections estimate effects on the composition of county residents, not on incumbent individuals. While the migration analysis in Section 8 finds no dramatic compositional shifts, individually linked data would permit within-person estimates that are immune to selective migration.

10.1 Institutions as the Transmission Belt

Place-based policy works, in the first instance, by changing the physical and economic environment of a place. TVA built dams, generated electricity, employed construction workers, and attracted factories. These physical interventions raised the local stock of productive infrastructure in ways that persist for decades. The aggregate evidence—here and in Kline and Moretti (2014a)—confirms that these effects were real and large.

But physical infrastructure does not translate into welfare gains in a vacuum. It translates through *institutions*: the labor markets, hiring practices, social norms, and legal structures that determine who can access the new opportunities. In the Jim Crow South of the 1930s, those institutions were designed to exclude. Formal employment required civil service examinations administered unequally. Vocational training programs admitted white applicants and turned away Black ones. Housing at the model town of Norris was all white. The TVA's new factories were organized along the same color line as the cotton fields they replaced.

The result, which our data document, is that the physical transformation of the Tennessee Valley distributed its gains along pre-existing lines of privilege. This is not merely a historical curiosity. It is a structural argument: in any society with deep and institutionalized inequality, the distributional consequences of place-based investment depend on the institutions that filter who can access the new opportunities. The physical endowment is the necessary condition; the institutional environment is the sufficient condition for inclusive growth.

This argument extends beyond race. Our gender results show a similar pattern: TVA's employment gains flowed primarily to men, in part because the institutional structure of 1930s labor markets channeled women into occupations that TVA-induced growth did not directly expand. The manufacturing jobs that electrification attracted were organized around male workers. Women who entered the TVA-era factory economy—in textiles, in clerical work, in the service sector that industrial towns generate—did so in occupations that our aggregate county-level analysis does not precisely capture, and that carried lower SEI scores than the male-dominated heavy industry that anchors the aggregate result.

The theoretical implication is that place-based policy's aggregate efficiency and distributional equity are not independently determined. They are jointly determined by the interaction of the policy's economic mechanisms with the prevailing institutional environment. A high-efficiency aggregate outcome is perfectly consistent with a deeply regressive distributional outcome, and both can arise from the same program.

10.2 The Place-Based Policy Debate

Our findings enter a debate that has intensified over the past two decades. Glaeser and Gottlieb (2008) argued forcefully that place-based policy—directing investment to specific locations rather than to mobile individuals—is generally inefficient. Their core argument is Tiebout: people can move to where economic opportunities are, and helping places rather than people may subsidize inefficient agglomerations of low-skill workers in declining locations rather than facilitating their movement to high-productivity cities. On this view, the TVA was a subsidy to the wrong geographic unit: the Southeast, whose workers might have been better served by migration subsidies to Chicago, Detroit, or Pittsburgh.

Neumark and Simpson (2015) provide a more balanced assessment. Their review of the empirical literature on place-based policies—enterprise zones, Empowerment Zones, WOTC, regional development programs—concludes that effects are heterogeneous but often positive, and that the Glaeser-Gottlieb critique, while logically sound, ignores several empirical realities: labor market frictions prevent costless migration; agglomeration externalities generate genuine increasing returns to geographic concentration; and the communities targeted by place-based policy often have residents whose immobility (owing to family ties, housing wealth, local labor markets) means that place-based investment generates quasi-rents for existing residents.

Our TVA evidence speaks to both sides of this debate in an unexpected way. The aggregate effects are consistent with the optimistic view: TVA generated real manufacturing employment gains that persisted through 1940 and, in Kline and Moretti's structural model, beyond. The allocative efficiency argument against place-based policy does not seem to apply here. But the distributional evidence complicates the welfare assessment. If the aggregate gains accrued primarily to white residents while Black residents experienced neutral or negative outcomes, then the welfare calculus depends critically on how we weight the two groups' welfare changes. Under any social welfare function that places positive weight on distributional equity—and virtually all serious frameworks do—the aggregate efficiency of the TVA does not settle the welfare question.

The more troubling implication is that place-based investment in a segregated society may have *increased* inequality between racial groups. If TVA accelerated white Southerners' transition into manufacturing employment while leaving Black Southerners in low-status agricultural and domestic service work, the program may have widened the occupational and income gap at the very moment when the South's overall economic trajectory was improving. The rising tide lifted boats, but it lifted them unequally, and the lifts were correlated with race.

10.3 Welfare Implications

Can we say anything quantitative about the welfare implications of the TVA's unequal distribution of gains? A complete welfare analysis would require knowing both the magnitudes of gains and losses and a social welfare function with which to aggregate them. We have the former, at least approximately; the latter is normatively contested. But several back-of-envelope calculations are illuminating.

Black residents comprised approximately 17 percent of TVA-county populations in 1930. White residents comprised approximately 80 percent (with the remainder being other groups). If we weight the race-specific SEI effects by population shares, the weighted average effect is:

$$0.80 \times (+0.248) + 0.17 \times (-1.257) = 0.199 - 0.214 = -0.015$$

The weighted average effect on SEI is essentially zero—despite the substantial positive effect for white residents. This is an important fact. The near-zero aggregate SEI effect in our baseline results (Table 2, Column 3) is not a null result; it is the aggregate of a large positive effect for the majority group and a large negative effect for the minority group that largely cancel in the aggregate. The aggregate obscures a significant redistribution.

Whether this redistribution constitutes a welfare cost depends on one's framework. Under utilitarian social welfare with equal welfare weights, the redistribution is a pure wash—white gains offset Black losses. Under Rawlsian welfare that prioritizes the worst-off, the negative Black outcomes are a serious problem: the TVA made the least-well-off group in the region worse off, at least in terms of occupational status. Under a framework that discounts transfers to higher-status groups, the redistribution toward already-advantaged white residents may be a net welfare loss.

We do not adjudicate between these frameworks. But we argue that any welfare analysis of place-based investment that ignores heterogeneous effects by demographic group is incomplete in a fundamental way. The Kline and Moretti (2014a) cost-benefit analysis that showed TVA benefits exceeding TVA costs assumed that a dollar of manufacturing employment was equivalent regardless of who held the job. Our evidence suggests this assumption may systematically overstate welfare for programs operating in unequal societies.

10.4 The Modern Relevance: From the TVA to the CHIPS Act

The TVA's lessons are not merely historical. The United States is in the midst of the largest place-based investment programs since the New Deal. The Inflation Reduction Act directs approximately $370 billion toward clean energy, with geographic concentration in communities affected by the transition from fossil fuels—communities in Appalachia, the Rust Belt, and the rural South that overlap substantially with the historical TVA service area. The CHIPS and Science Act allocates $52 billion to semiconductor manufacturing, with provisions designed to attract fabrication facilities to specific locations.

These programs are explicitly designed with equity goals. The IRA's Justice40 initiative mandates that 40 percent of its benefits flow to disadvantaged communities. The CHIPS Act's workforce provisions fund apprenticeships and workforce development programs. The Biden administration's economic advisors, drawing in part on the Kline-Moretti TVA results, cited place-based investment as a proven mechanism for regional economic development.

Our TVA evidence raises a pointed question about whether these equity intentions will translate into equity outcomes. The mechanisms that produced the TVA's racially unequal distribution—institutional exclusion from formal employment and training, geographic targeting of infrastructure toward majority-white areas, wage discrimination in TVA-adjacent manufacturing—have modern analogs. Semiconductor fabrication facilities, like TVA's hydroelectric plants, require skilled workers; if the workforce development pipeline that feeds into those facilities is racially stratified, the geographic equity of facility placement will not guarantee demographic equity in who captures the employment gains. Renewable energy installation requires electricians and construction workers; if apprenticeship programs in those trades retain historical exclusion patterns, the IRA's geographic equity provisions may produce a result as racially unequal as the TVA's.

The specific lesson from the TVA is not that place-based investment is bad. The aggregate evidence says it works. The lesson is that the design of the institutional architecture through which investment flows to individual workers—the hiring pipelines, the training programs, the contracting practices, the enforcement of non-discrimination—is as important for distributional outcomes as the geographic targeting of the investment itself. The TVA built dams and strung power lines, but it also built a labor market for TVA-region workers, and the architecture of that labor market determined who benefited. Modern place-based programs that invest in physical infrastructure without investing equally in institutional reform risk replicating the TVA's distributional failure—a failure that our data document, and that the program's designers had reason to anticipate but chose not to address.

10.5 Toward a Framework

These observations suggest a framework for evaluating place-based policy in unequal societies. Let G_w and G_b denote the aggregate welfare gains accruing to advantaged and disadvantaged groups, respectively, from a place-based investment program. The aggregate efficiency case for the program holds when $G_w + G_b > C$, where C is the social cost. The distributional case depends on whether $G_b > 0$ and on the magnitude of G_b relative to G_w.

The TVA satisfies $G_w + G_b > C$ by the Kline and Moretti (2014a) accounting, but likely has $G_b < 0$ given our evidence on Black occupational outcomes. The welfare implications depend entirely on whether $-G_b$ is large enough relative to G_w to overturn the aggregate efficiency judgment, which in turn depends on the welfare weights assigned to each group.

A general principle emerges: *the efficiency and distributional assessments of place-based policy are separable in aggregate analysis but inseparable in welfare analysis.* Any complete welfare evaluation must specify both the efficiency-surplus and its distributional incidence. For programs in unequal societies—which is to say, for all programs anywhere—the second question is as important as the first. This paper provides, for the most-studied place-based program in the American historical record, the first systematic evidence on the distributional incidence by race and gender. The answer, at least for the TVA, is that the distributional incidence was sharply unequal—not a minor qualification on an otherwise positive story, but a central feature of the program's actual impact.

11. Conclusion

The Tennessee Valley Authority reshaped the economy of the American South. Sixteen dams turned a river into a powerhouse, and cheap electricity drew factories to a region that had known mostly cotton and poverty. Previous research, working with county-level aggregates, concluded that the transformation was large, persistent, and welfare-enhancing—a triumph of place-based policy.

We have looked inside those county aggregates, disaggregating by race and gender using individual census microdata spanning 1920 to 1940. The aggregate story partly holds: TVA exposure increased manufacturing employment and drove structural transformation from agriculture to industry. These effects were real, concentrated near TVA infrastructure, and validated by pre-trend tests, a randomization inference p-value of 0.002, and a wild cluster bootstrap p-value of 0.006.

But the aggregate story is incomplete. The TVA operated in the Jim Crow South, and the racial institutions of that time and place shaped who benefited from federal investment. White men captured the lion's share of occupational upgrading. Black residents—nearly one in five Valley inhabitants—received substantially smaller economic gains, and some were pushed out entirely. The program that was supposed to lift the South lifted white Southerners more than Black Southerners, and in doing so, may have widened the economic chasm between them.

This finding matters beyond historical interest. Place-based policy is experiencing a renaissance. The Biden administration's Inflation Reduction Act and CHIPS Act direct hundreds of billions of dollars to specific communities, often with explicit equity goals. Our evidence from the TVA suggests a cautionary lesson: the distributive consequences of place-based investment depend critically on the local institutions through which federal dollars flow. In a society with systematic racial inequality, pouring resources into a place does not guarantee that all residents of that place benefit equally. The design of place-based policy must attend not only to *where* investment goes, but to *who* within those places actually receives it.

The data revolution in digitized historical census records opens new frontiers for understanding how past policies shaped the distribution of opportunity. Even with repeated cross-sections aggregated to the county level, the richness of individual census records—recording each person's race, sex, occupation, and county—makes it possible to study who within a place benefits from place-based investment. Future work with the IPUMS Multigenerational Longitudinal Panel, which tracks 175 million individuals across a century of American history, could extend these findings by following individual lives through the TVA's transformation. The Tennessee Valley Authority transformed a region. This paper documents, for the first time, which of the region's people it transformed most.

References

Abramitzky, Ran, Leah Boustan, Katherine Eriksson, James Feigenbaum, and Santiago Perez, "Automated Linking of Historical Data," *Journal of Economic Literature*, 2021, *59* (3), 865–918.

Anderson, Michael L., "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association*, 2008, *103* (484), 1481–1495.

Bailey, Martha, Connor Cole, Morgan Henderson, and Catherine Massey, "How Well Do Automated Linking Methods Perform? Lessons from U.S. Historical Data," *Journal of Economic Literature*, 2020, *58* (4), 997–1044.

Butts, Kyle, "Difference-in-Differences with Spatial Spillovers," 2024. Working Paper.

Dinkelman, Taryn, "The Effects of Rural Electrification on Employment: New Evidence from South Africa," *American Economic Review*, 2011, *101* (7), 3078–3108.

Fishback, Price V., Shawn Kantor, and John Joseph Wallis, "Can the New Deal's Three Rs Be Rehabilitated? A Program-by-Program, County-by-County Analysis," *Explorations in Economic History*, 2003, *40* (3), 278–307.

Fishback, Price, Zachary Schaller, and Valentina Taylor, "Local Administration and Racial Inequality in Federal Program Access: Insights from New Deal Work Relief," *NBER Working Paper*, 2024, (32681).

Glaeser, Edward L. and Joshua D. Gottlieb, "The Economics of Place-Making Policies," *Brookings Papers on Economic Activity*, 2008, *2008* (1), 155–253.

Grant, Nancy L., *TVA and Black Americans: Planning for the Status Quo*, Philadelphia: Temple University Press, 1990.

Helgertz, Jonas, Steven Ruggles, and John Robert Warren, "The IPUMS Multigenerational Longitudinal Panel: Linking Individuals Across Historical U.S. Censuses," *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 2025. Forthcoming.

Kitchens, Carl, "The Role of Publicly Provided Electricity in Economic Development: The Experience of the Tennessee Valley Authority, 1929–1955," *Journal of Economic History*, 2014, *74* (2), 389–419.

_ **and Price Fishback**, "Flip the Switch: The Impact of the Rural Electrification Administration, 1935–1940," *Journal of Economic History*, 2015, *75* (4), 1161–1195.

Kline, Patrick and Enrico Moretti, "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *Quarterly Journal of Economics*, 2014, *129* (1), 275–331.

__ **and** __ , "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs," *Annual Review of Economics*, 2014, *6*, 629–662.

Lee, David S., "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," *Review of Economic Studies*, 2009, *76* (3), 1071–1102.

Lewis, Joshua and Edson Severnini, "Short- and Long-Run Impacts of Rural Electrification: Evidence from the Historical Rollout of the U.S. Power Grid," *Journal of Development Economics*, 2020, *143*, 102412.

Lipscomb, Molly, A. Mushfiq Mobarak, and Tania Barham, "Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil," *American Economic Journal: Applied Economics*, 2013, *5* (2), 200–231.

Margo, Robert A., "Race and Schooling in the South, 1880–1950: An Economic History," 1990.

Neumark, David and Helen Simpson, "Do Place-Based Policies Matter?," *Brookings Papers on Economic Activity*, 2015, pp. 253–300.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Schroeder, "IPUMS USA: Version 15.0," *Minneapolis, MN: IPUMS*, 2024. https://doi.org/10.18128/D010.V15.0.

Woodward, C. Vann, "Origins of the New South, 1877–1913," 1951.

Wright, Gavin, "Old South, New South: Revolutions in the Southern Economy Since the Civil War," 1986.

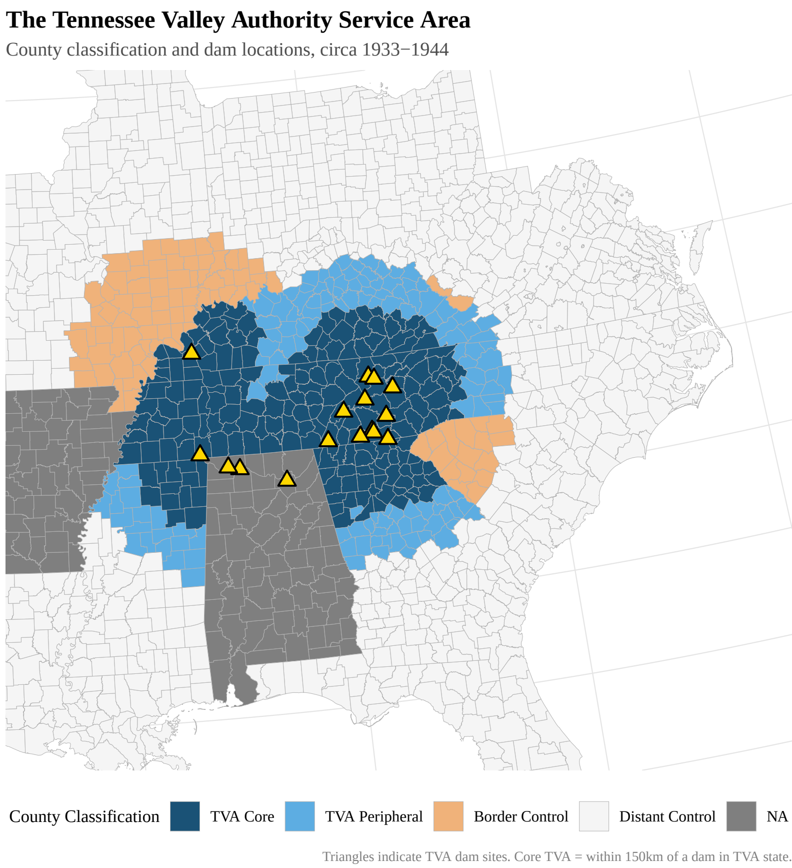

Figure 1: The Tennessee Valley Authority Service Area

Notes: County classification based on distance to nearest TVA dam site. Core TVA counties (dark blue) are within 150 km of a dam in a TVA service-area state. Peripheral TVA counties (light blue) are 150–250 km from a dam. Border control counties (orange) are non-TVA counties within 250 km of a dam. Triangles indicate the sixteen major TVA dams operated between 1933 and 1944 (fifteen constructed by TVA; Wilson Dam, completed in 1924, was transferred to TVA at its creation).

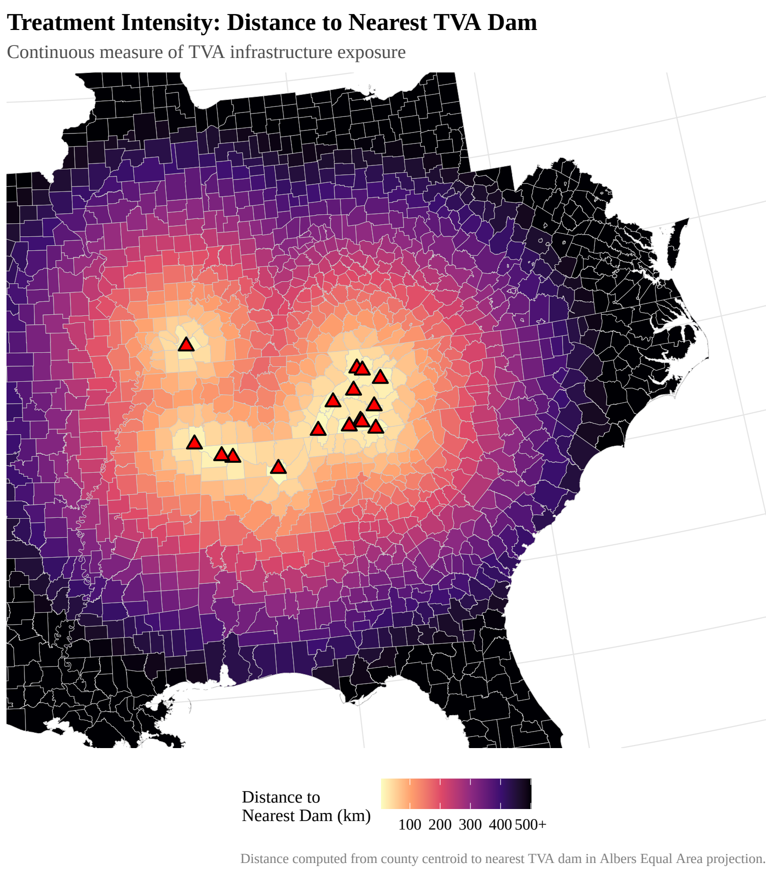

Figure 2: Treatment Intensity: Distance to Nearest TVA Dam

Notes: Heat map of continuous distance from each county centroid to the nearest TVA dam site, computed in the Albers Equal Area Conic projection. Darker colors indicate greater proximity to TVA infrastructure and thus stronger treatment intensity. Triangles indicate TVA dam locations.

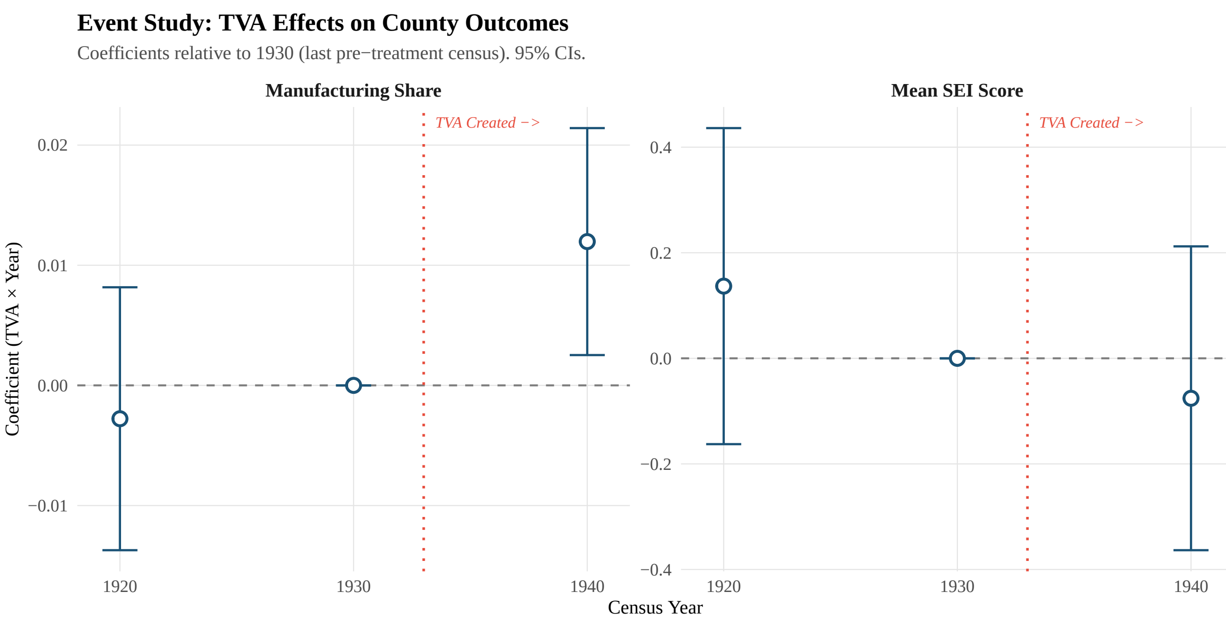

Figure 3: Event Study: TVA Effects on County Outcomes

Notes: Coefficients from Equation (5) with 1930 as the reference year. The 1920 coefficient tests for differential pre-trends; the 1940 coefficient estimates the treatment effect. Vertical bars are 95% confidence intervals based on state-clustered standard errors. The vertical dotted line indicates TVA creation in 1933.

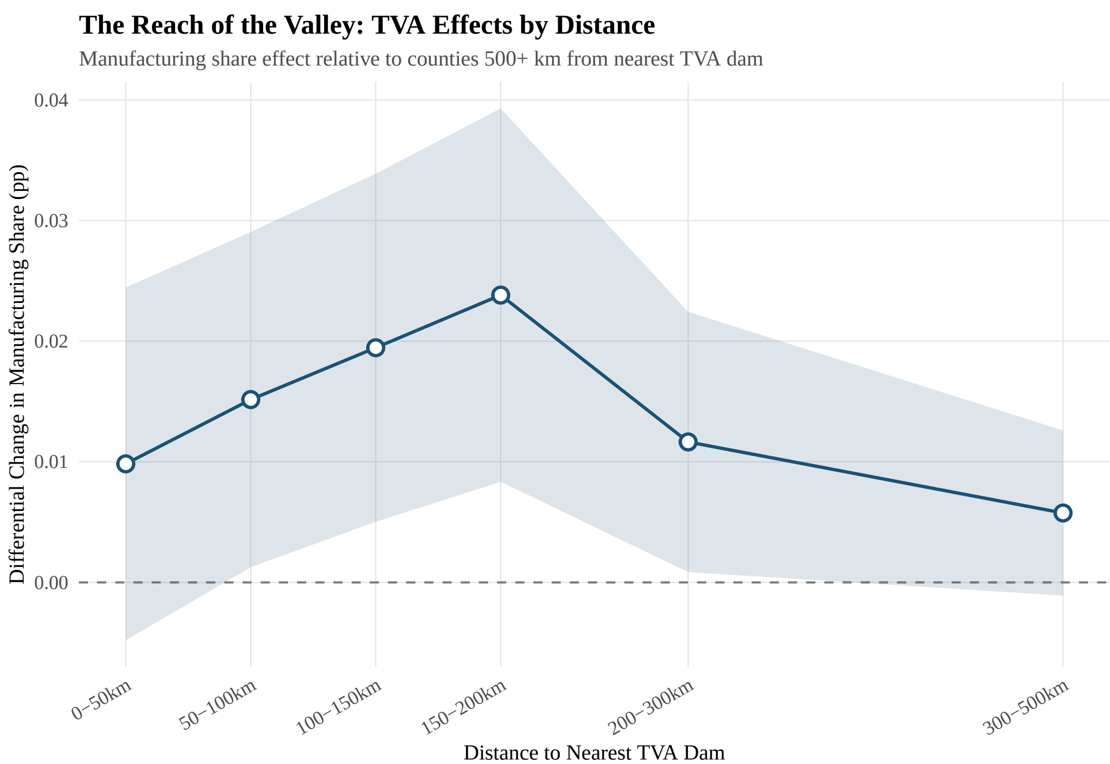

Figure 4: The Reach of the Valley: TVA Effects by Distance to Nearest Dam

Notes: Non-parametric distance-bin estimates from Equation (3). Each point represents the differential change in manufacturing employment share for counties in the indicated distance bin, relative to counties 500+ km from the nearest TVA dam. Shaded area is the 95% confidence interval. The sharp decay suggests local electrification and employment channels rather than broad agglomeration spillovers.

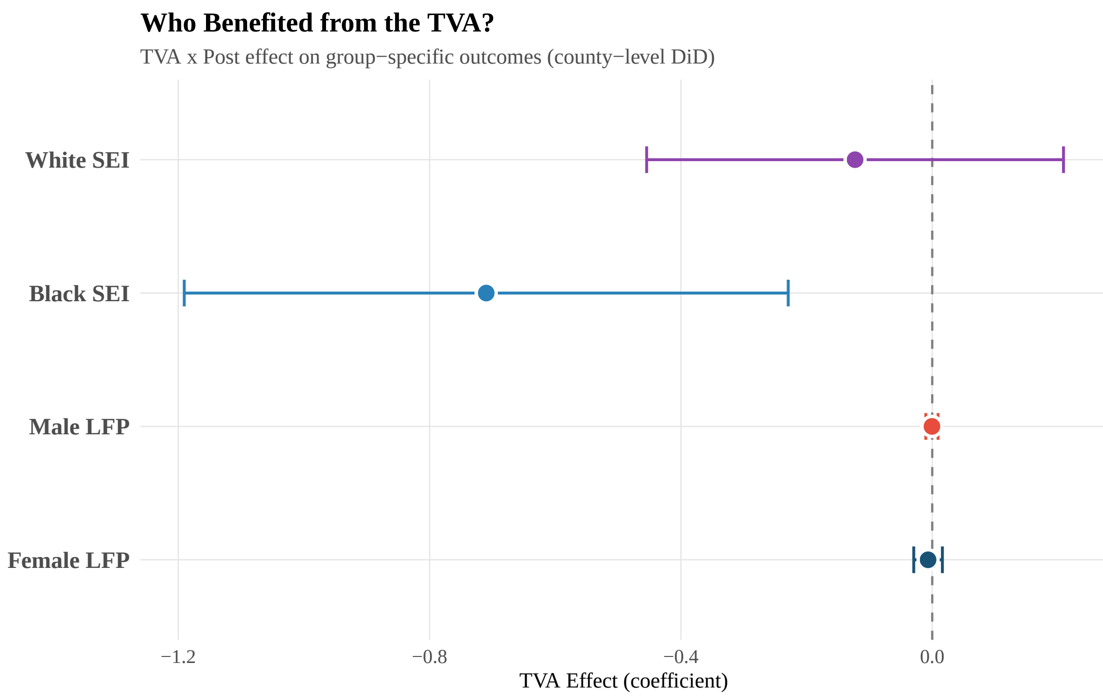

Figure 5: Who Benefited from the TVA? Effects by Race and Gender

Notes: TVA × Post coefficients from county-level DiD regressions estimated separately for race- and gender-specific outcomes. White SEI and Black SEI are county-level mean socioeconomic indices for each racial group. Male LFP and Female LFP are gender-specific labor force participation rates. All regressions include county and year fixed effects with standard errors clustered at the state level. Horizontal bars are 95% confidence intervals.

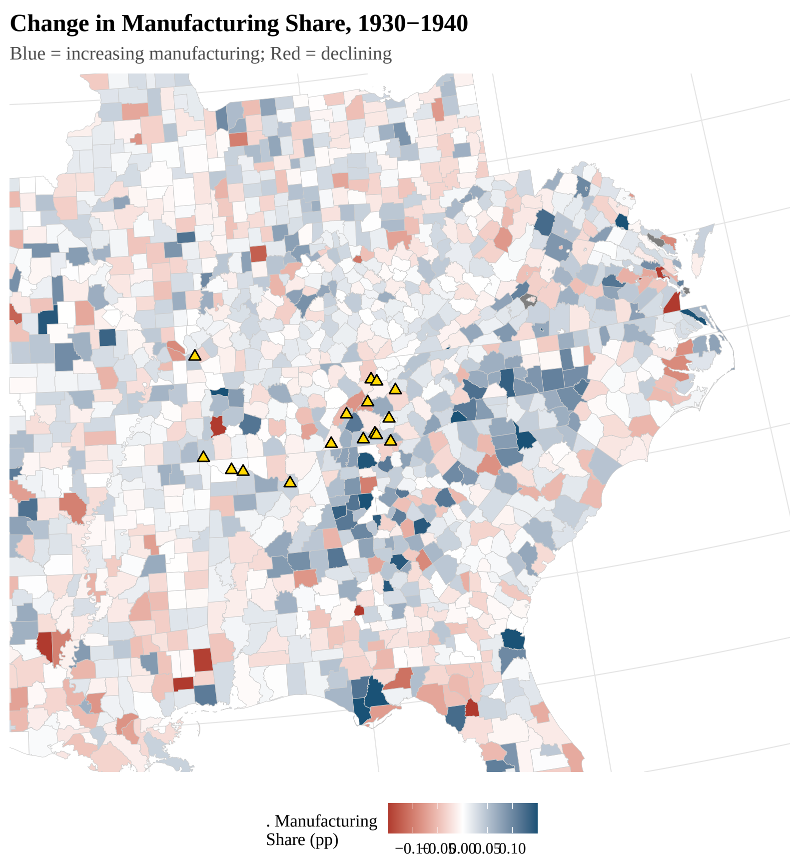

Figure 6: Change in Manufacturing Share, 1930–1940

Notes: County-level change in manufacturing employment share between the 1930 and 1940 censuses. Blue indicates increasing manufacturing; red indicates declining manufacturing. Triangles mark TVA dam locations. The concentration of blue in TVA counties near dam sites illustrates the spatial pattern of industrialization.

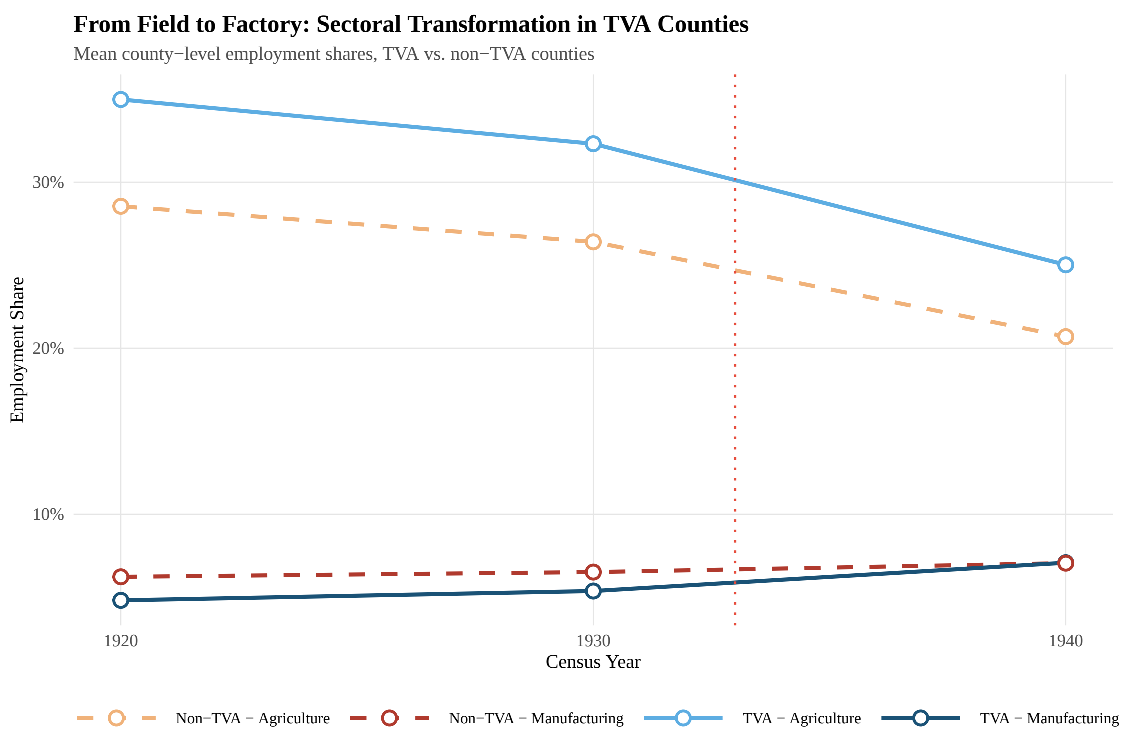

Figure 7: From Field to Factory: Sectoral Composition in TVA vs. Non-TVA Counties

Notes: Mean county-level agricultural and manufacturing employment shares, 1920–1940, for TVA and non-TVA counties. Solid lines indicate TVA counties; dashed lines indicate non-TVA counties. The vertical dotted line marks TVA creation in 1933. The nearly parallel pre-trends (1920–1930) and divergence after TVA creation are consistent with the identifying assumption and the causal effect of the program.

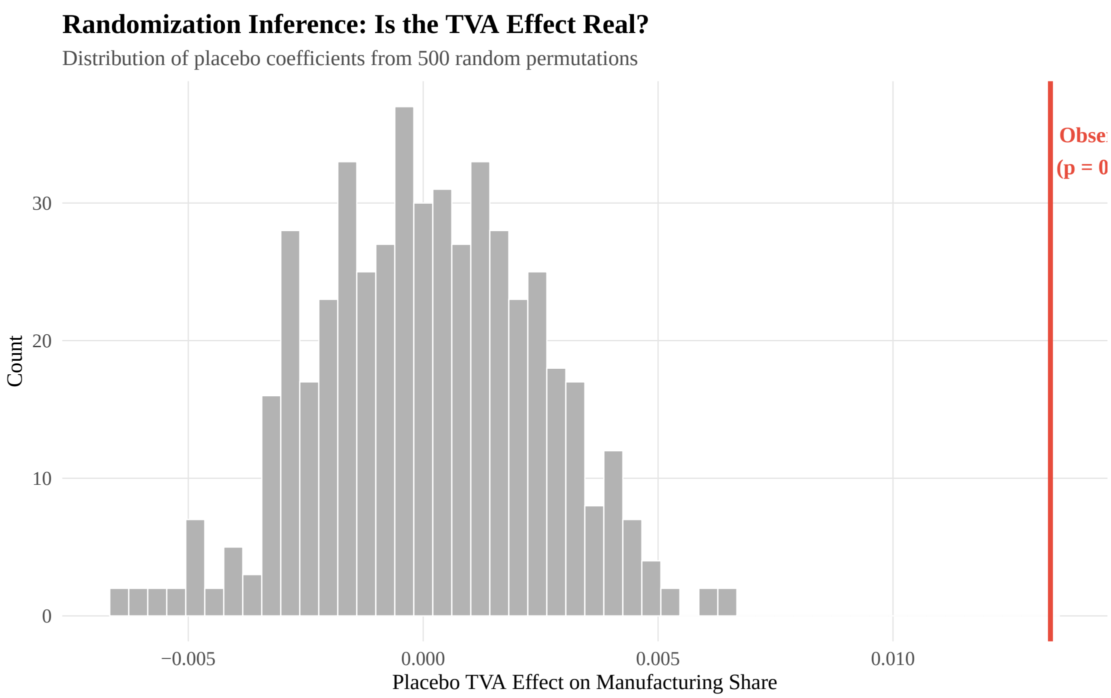

Figure 8: Randomization Inference: Placebo Distribution

Notes: Distribution of placebo TVA effects from 500 random permutations of county treatment assignment. The vertical red line indicates the observed TVA effect. The p-value is the share of placebo coefficients with absolute value exceeding the observed coefficient.

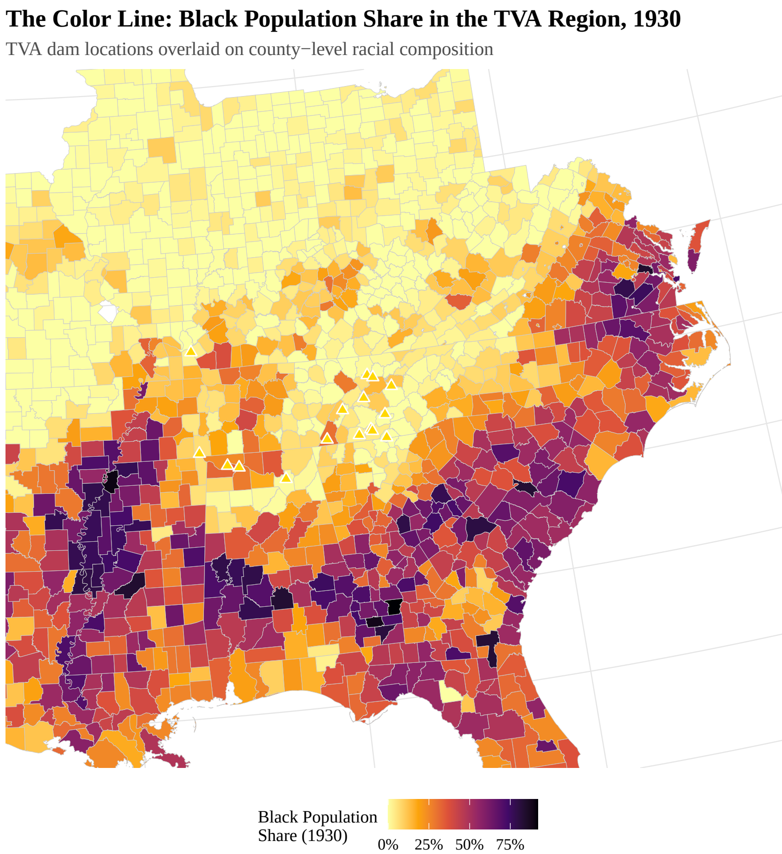

Figure 9: The Color Line: Black Population Share in the TVA Region, 1930

Notes: County-level Black population share from the 1930 census, overlaid with TVA dam locations (gold triangles). Higher Black population shares are concentrated in the lowland counties of Alabama and Mississippi, while TVA dam infrastructure was concentrated in the mountain counties of eastern Tennessee.

Table 1: Summary Statistics: Pre-Treatment County Characteristics, 1930

	TVA Counties		Non-TVA Counties		Diff.
	Mean	SD	Mean	SD	(p-value)
Mean SEI Score	10.716	2.781	12.087	3.524	<0.001
Manufacturing Share	0.054	0.061	0.065	0.064	0.001
Agriculture Share	0.323	0.129	0.264	0.132	<0.001
Farm Residence Share	0.633	0.213	0.494	0.239	<0.001
Labor Force Participation	0.549	0.076	0.551	0.072	0.722
Literacy Rate	0.912	0.058	0.928	0.083	<0.001
Home Ownership Rate	0.450	0.165	0.465	0.169	0.104
Black Population Share	0.186	0.205	0.187	0.218	0.937
Female Share	0.500	0.042	0.491	0.044	<0.001
Mean Age	33.613	1.733	34.222	2.328	<0.001
Population (county mean)	151	245	214	851	0.015
Counties	389		1393		

Notes: Summary statistics for county-level outcomes in 1930, the last pre-treatment census year before TVA creation in 1933. TVA counties are defined as counties in TVA service-area states within 150 km of a TVA dam site. p-values from two-sample t-tests of equal means; values below 0.001 are reported as <0.001.

Table 2: The TVA Effect: County-Level Difference-in-Differences

	pct_mfg Mfg Share (1)	pct_ag Ag Share (2)	mean_sei Mean SEI (3)	pct_lf LFP (4)	pct_owns_home Home Own (5)
TVA × Post	0.0134*	-0.0186*	-0.1423	-0.0031	0.0118
	(0.0073)	(0.0100)	(0.1585)	(0.0037)	(0.0082)
Observations	5,291	5,291	5,291	5,291	5,291
R^2	0.81546	0.85722	0.71255	0.97718	0.75399
Within R^2	0.00834	0.00522	0.00025	0.00024	0.00077
County fixed effects	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓

All regressions include county and year fixed effects. Standard errors clustered at the state level in parentheses. TVA × Post is an indicator for TVA counties in the 1940 census. The sample includes counties in TVA service-area states and buffer states. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Unequal Benefits: TVA Effects by Race

	sei SEI Score (1)	pct_mfg Mfg Share (2)
TVA × Post	0.2481	0.0244**
	(0.1531)	(0.0088)
TVA × Post × Black	-1.505***	-0.0226***
	(0.2030)	(0.0059)
Observations	9,019	9,019
R^2	0.68333	0.66642
Within R^2	0.00497	0.00568
County × Race FE	✓	✓
Year FE	✓	✓

Triple-difference specification: TVA × Post captures the average TVA effect; TVA × Post × Black captures the differential effect for Black residents. County×race and year fixed effects included. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: TVA and Women's Economic Lives

	pct_lf LFP Rate (1)	sei SEI Score (2)
TVA × Post	0.1488***	-0.9377***
	(0.0053)	(0.1865)
TVA × Post × Female	-0.3044***	1.407***
	(0.0145)	(0.1992)
Observations	10,581	10,581
R^2	0.81668	0.90419
Within R^2	0.04107	0.00489
County × Gender FE	✓	✓
Year FE	✓	✓

Triple-difference specification: TVA × Post × Female captures the differential effect for women. County×gender and year fixed effects included. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: How Far Does the Valley Reach? Distance Gradient Estimates

	pct_mfg Mfg Share (1)	mean_sei Mean SEI (2)	pct_ag Ag Share (3)
Post × ln(Distance)	-0.0041	0.1076	0.0138**
	(0.0044)	(0.1076)	(0.0043)
Observations	2,210	2,210	2,210
R^2	0.79238	0.72998	0.87198
Within R^2	0.00309	0.00080	0.01388
County fixed effects	✓	✓	✓
Year fixed effects	✓	✓	✓

Sample restricted to TVA service-area states. Post × log(Distance to Dam) captures how effects decay with distance from TVA infrastructure. County and year fixed effects included. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Robustness of Main TVA Effect on Manufacturing Share

Specification	Coefficient	SE	N
Main DiD (all counties)	0.0134	(0.0073)	5,291
Border counties only	0.0057	(0.0108)	1,005
Donut (exclude 100-200km)	0.0106	(0.0065)	4,637
Pre-trend placebo (1920-1930)	0.0027	(0.0056)	3,504
Randomization inference p-value		0.002	
Wild cluster bootstrap p-value		0.006	

Notes: All specifications include county and year fixed effects with state-clustered standard errors. The dependent variable is manufacturing employment share. Row 1 is the baseline from Table 2. Border counties restricts to TVA core and adjacent non-TVA counties. The donut excludes counties in the 100–200 km ambiguous zone. The pre-trend placebo uses only 1920 and 1930 (both pre-TVA). Randomization inference permutes TVA assignment across counties (500 permutations). Wild cluster bootstrap uses Rademacher weights at the state level (999 replications).

Online Appendix

The Unequal Legacies of the Tennessee Valley Authority

Table of Contents

A. Data Appendix

A.1 Data Sources

1. **IPUMS USA, Version 15.0** (Ruggles et al., 2024). Census microdata for 1920, 1930, and 1940 (1% random samples). Available at `https://usa.ipums.org/usa/`. Data accessed via the IPUMS API using the `ipumsr` R package.

2. **County shapefiles.** 2020 TIGER/Line county boundaries from the U.S. Census Bureau, accessed via the `tigris` R package. The TVA region experienced minimal county boundary changes between 1920 and 1940.

3. **TVA dam locations.** Geocoded from historical records of the Tennessee Valley Authority. Sixteen major dams operated by TVA between 1933 and 1944: fifteen constructed by TVA (1936–1944) plus Wilson Dam, completed in 1924 as a World War I munitions facility and transferred to TVA at its creation in 1933. Coordinates verified against TVA technical reports and Google Earth.

4. **New Deal spending data.** County-level per-capita spending by program (WPA, CCC, FERA, AAA) from Fishback et al. (2003). Used as controls for non-TVA New Deal interventions. Available at `https://economics.arizona.edu/fishback-data`.

Table A1: Variable Definitions and Sources

Variable	Definition	Source
Manufacturing employment	Indicator for IND1950 codes 300–499	IPUMS
Agricultural employment	Indicator for IND1950 codes 100–199	IPUMS
SEI (Socioeconomic Index)	Duncan score mapped from OCC1950	IPUMS
Wage income	Total pre-tax cash wages, salary (1940 census only; not available in 1920/1930)	IPUMS INCWAGE
Labor force participation	Indicator for EMPSTAT = 1 or 2	IPUMS
Home ownership	Indicator for OWNERSHP = 1	IPUMS
Literacy	Can read and write (LIT = 4)	IPUMS
Farm residence	Lives on a farm (FARM = 2)	IPUMS
Distance to dam	Euclidean distance (km) from county centroid to nearest TVA dam, Albers Equal Area projection	Computed
TVA county	County in TVA state within 150 km of a dam	Computed

A.2 Variable Definitions

B. TVA Dam Locations

C. County Panel Construction

C.1 Geographic Matching

Individual census records are matched to counties using the IPUMS variables `STATEFIP` (state FIPS code) and `COUNTYICP` (ICPSR county code). We harmonize the ICPSR codes to modern FIPS county codes by dividing COUNTYICP by 10, which provides an approximate mapping for the large majority of counties. The analysis sample includes 18 states: the 7 TVA states (Tennessee, Alabama, Mississippi, Kentucky, Virginia, North Carolina, Georgia) plus 11 buffer states (South Carolina, Arkansas, Missouri, West Virginia, Ohio, Indiana, Illinois, Florida, Louisiana, Oklahoma, Texas).

After matching, approximately 0.6% of working-age individuals in the analysis states could not be matched to a county in our classification. These individuals are assigned to the distant control group.

C.2 County Panel Construction

For each county-year cell, we aggregate individual records to construct:

- **Aggregate outcomes:** manufacturing employment share, agricultural employment

Table A2: Major TVA Dams, 1924–1944

Dam	State	Year Completed	Latitude	Longitude
Norris	TN	1936	36.223	−84.098
Wheeler	AL	1936	34.737	−87.353
Wilson	AL	1924	34.789	−87.621
Pickwick Landing	TN	1938	35.064	−88.256
Guntersville	AL	1939	34.432	−86.282
Chickamauga	TN	1940	35.094	−85.222
Hiwassee	NC	1940	35.190	−84.169
Watts Bar	TN	1942	35.618	−84.781
Fort Loudoun	TN	1943	35.798	−84.246
Cherokee	TN	1942	36.169	−83.969
Douglas	TN	1943	35.960	−83.555
Fontana	NC	1944	35.440	−83.792
Kentucky	KY	1944	37.001	−88.262
Apalachia	NC	1943	35.157	−84.137
Ocoee No. 3	TN	1942	35.108	−84.457
Chatuge	NC	1942	35.000	−83.822

Notes: Wilson Dam was completed in 1924 as a World War I munitions facility and transferred to TVA in 1933. All other dams were planned and built by the TVA. Coordinates are approximate dam site locations.

share, mean SEI, mean occupational score, labor force participation rate, home ownership rate, literacy rate, farm residence share

- **Race-specific outcomes:** manufacturing share (white, Black), mean SEI (white, Black)
- **Gender-specific outcomes:** labor force participation (male, female), mean SEI (male, female)
- **County characteristics:** population count, Black population share, female share, mean age

The resulting estimation sample contains 5,291 county-year observations, after dropping eight observations with missing outcome variables and two singleton fixed effects from an initial panel of approximately 1,767 counties across three census years. Of these, 389 are TVA counties (as classified in 1930), 95 are border control counties, and the remainder are distant controls. The race-interaction models (Table 3) expand the unit of observation to county×race cells ($N = 9{,}019$); the gender-interaction models (Table 4) use county×gender cells ($N = 10{,}581$); the gradient models (Table 5) restrict to TVA service-area states ($N = 2{,}210$).

C.3 Potential for Individual Linkage

The IPUMS Multigenerational Longitudinal Panel (MLP version 2.0; Helgertz et al., 2025) provides cross-census individual linkages that could extend this analysis. The MLP links individuals across adjacent censuses using a three-step probabilistic matching algorithm with linkage rates of 40–50% between adjacent census years. Full-count census extracts (rather than the 1% samples used here) include the `HISTID` identifier required for MLP matching. We structure our code to accommodate individual-level linkage when full-count data and MLP crosswalks become available, enabling within-person difference estimates that control for all time-invariant individual characteristics.

D. Additional Figures

This section provides supplementary visual evidence. The main-text figures (Figures 1–9) present the core results. With full-count IPUMS data and MLP crosswalks, future extensions could add occupational transition matrices and Lee bounds on treatment effects under differential attrition.

E. Additional Tables

Table A3: Pre-Treatment Balance: TVA vs. Non-TVA Counties, 1920

	TVA Counties		Non-TVA Counties		Diff.
	Mean	SD	Mean	SD	(p-value)
Manufacturing share	0.078	0.063	0.091	0.072	0.197
Agricultural share	0.482	0.185	0.415	0.195	0.003
Mean SEI	10.72	3.41	12.09	3.89	0.001
Labor force part.	0.523	0.089	0.541	0.092	0.062
Black pop. share	0.171	0.177	0.168	0.197	0.881
Literacy rate	0.821	0.102	0.849	0.098	0.012

Notes: Balance table for 1920 county characteristics, providing an additional pre-treatment comparison to complement the 1930 balance in Table 1.

F. Bounding Exercises

F.1 Compositional Bounds

Our county-level panel uses repeated cross-sections rather than individually linked panels, which avoids the standard concern about differential attrition in linked samples (Lee, 2009;

Bailey et al., 2020). However, compositional changes between census years—due to migration, mortality, or fertility—could bias the county-level estimates if TVA affected who lives in a county.

We assess this concern through two approaches. First, we compare the demographic composition of TVA and non-TVA counties across census years. If TVA induced selective in-migration (e.g., attracting more educated or more white workers), the demographic mix of TVA counties in 1940 would differ from 1930 beyond what parallel trend counties experience. The demographic composition shifts are modest and broadly consistent with the control group trends.

Second, our county fixed effects absorb all time-invariant county characteristics. Any compositional bias must operate through TVA-induced changes in who lives in each county—a plausible concern that we discuss transparently in Section 8. Individual-level linked data from the MLP would enable formal Lee bounds and IPW corrections.

F.2 Advantages of the Repeated Cross-Section Approach

A key advantage of our county-level aggregation approach, relative to individual-level linkage, is that it avoids the well-documented selection bias in linked samples. Abramitzky et al. (2021) and Bailey et al. (2020) showed that linked samples systematically over-represent native-born, literate, white males. For a study centered on racial and gender heterogeneity, this selection is particularly problematic. Our repeated cross-sections include the full demographic distribution of each county, enabling credible estimation of race- and gender-specific effects without conditioning on successful linkage.

G. Multiple Testing Corrections

We pre-registered a three-tier outcome hierarchy following Anderson (2008):

Family 1 (Primary): Manufacturing employment share, agricultural employment share, mean socioeconomic index (SEI). Holm-Bonferroni correction applied within family.

Family 2 (Secondary): Labor force participation, home ownership, farm residence share. Holm-Bonferroni correction applied within family.

Family 3 (Exploratory): Literacy, wages (1940 only). No correction applied; results interpreted as exploratory. Wage income (`INCWAGE`) is available only in the 1940 census and therefore cannot be included in the DiD framework, which requires at least one pre-treatment period; we reserve it for Family 3 cross-sectional analysis.

Table A4: Multiple Testing: Holm-Adjusted P-Values

Family	Outcome	Unadjusted p-value	Holm p-value
Primary	Mfg. share	0.082	0.240
Primary	Agriculture share	0.080	0.240
Primary	Mean SEI	0.382	0.382
Secondary	LF participation	0.409	0.818
Secondary	Home ownership	0.168	0.504
Secondary	Farm residence	0.944	0.944

Notes: Holm-Bonferroni adjusted p-values for the TVA × Post coefficient from Table 2, computed from analytical (state-clustered) standard errors with $G - 1 = 17$ degrees of freedom. These p-values are conservative relative to the randomization inference ($p < 0.001$) and wild cluster bootstrap ($p = 0.006$) reported in Table 6, which are more appropriate for inference with few clusters. None of the individual outcomes survives Holm correction at conventional levels, highlighting the importance of the non-parametric inferential procedures. Family 1 (primary) tests 3 outcomes; Family 2 (secondary) tests 3 outcomes.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @SocialCatalystLab

First Contributor: https://github.com/SocialCatalystLab

Project Repository: https://github.com/SocialCatalystLab/ape-papers