The Local Employment Effects of Pollution Abatement Infrastructure Investments*

Taewon Hwang†

March 8, 2026

Abstract

Large-scale environmental regulations and infrastructure mandates are often criticized as costly, yet evidence on their potential to generate local economic co-benefits is scarce. This paper estimates the local labor-market effects of the Clean Water Act (CWA) Construction Grants Program using a county-level staggered difference-in-differences design exploiting variation in the timing of first grant-assisted construction starts. Initiating construction raises construction employment by approximately twenty jobs per county-year over the first two post-treatment years, with accompanying increases in establishments and wages. Effects are disproportionately concentrated in smaller counties, indicating that the local impact of federal infrastructure spending depends on the size of the receiving economy. I also find significant responses in total employment and per-capita personal income, consistent with spillovers beyond the construction sector. Overall, the results show that environmental infrastructure investment can also generate short-run local economic gains.

Keywords: Clean Water Act, pollution abatement, public infrastructure, local employment

JEL Classification: H45, H57, R11, E62, Q58

*First draft: March 8, 2026. I am grateful to Andrew Fieldhouse, Jorge Hirs Garzón, Yoon Joo Jo, Nawon Kang, Suhyeon Oh, Tatevik Sekhposyan, Soohyun Shin, Daeun Sung, Anna Ziff, and Sarah Zubairy for helpful comments and suggestions. I also thank Mark Mylin and Elisabeth Schlaudt at the EPA for assistance with data and documentation.

†Department of Economics, Texas A&M University, 2935 Research Parkway Suite 200, College Station, TX 77843, United States. E-mail: twhwang@tamu.edu.

1 Introduction

Water and wastewater systems are among the most capital-intensive components of local public infrastructure, and decades of underinvestment have left a large financing gap. Recent federal assessments put wastewater and stormwater needs at roughly $630 billion over a 20-year horizon (U.S. EPA, 2024). Yet large-scale water infrastructure programs are not without controversy; prior work suggests that measured environmental benefits may not always justify their substantial costs (Keiser & Shapiro, 2019; Keiser et al., 2019). A natural question is whether these programs generate additional economic co-benefits typically absent from cost-benefit accounting. Short-run local labor-market gains are one such channel: if construction-intensive infrastructure spending creates jobs and raises local incomes during implementation, this represents a meaningful offset to program costs that policymakers and researchers should take seriously.

Against this backdrop, water infrastructure has repeatedly been used as a policy lever in stimulus and recovery efforts. The ARRA directed $4 billion to water-quality infrastructure, and the IIJA provided an additional $12.7 billion in federal support. These policies are typically justified on two margins: (i) long-run public benefits (public health, environmental quality, and regulatory compliance), and (ii) short-run local stimulus, especially because much of the spending is construction-intensive (Dupor & Mehkari, 2016; H.R.3684 - 117th Congress (2021-2022), 2021). Yet evidence on short-run local labor-market impacts of major water infrastructure investments is thinner than for other categories of infrastructure, notably highways, and the extent to which water projects translate into broader local economic outcomes remains an open empirical question.

This paper provides new evidence on the short-run local economic effects of the Clean Water Act (CWA) Construction Grants Program, the large, grant-financed wastewater infrastructure program that operated primarily in the 1970s and 1980s. I assemble a county-level panel linking project construction timing to local labor-market outcomes and estimate dynamic treatment effects using a staggered difference-in-differences design keyed to the first year of grant-assisted construction in each county, comparing treated counties to not-yet-treated and never-treated counties. Grant allocation was determined by state priority lists based on documented pollution-control needs and affected population, and the CWA explicitly prohibited states from incorporating development objectives or projections of future population growth. Critically, the design exploits variation in the timing of first construction starts across counties rather than the fact of receipt itself; the determinants of pollution severity and population size do not have an obvious bearing on when a county first entered construction, and event-study pre-trends support this interpretation.

The first result is that CWA construction grants generate a clear and immediate expansion in the local construction sector, with construction employment rising by about 18 workers per county-year over the first two post-treatment years, corresponding to roughly 0.2 job-years per $100,000 of federal spending in 2017 dollars. For context, Garin (2019) estimates that ARRA highway construction generated approximately 0.5 job-years per $100,000; the smaller magnitude here is consistent with the CWA program operating during the tighter labor markets of the 1970s. These effects extend beyond construction: total employment rises by roughly 300 workers per county-year over the first two years, and per-capita personal income also increases, indicating broader local spillovers.

A key question is whether these effects are uniform across receiving economies. The evidence suggests they are not. Effects are disproportionately concentrated in smaller counties, where a federally financed project represents a larger shock relative to the existing construction base. In counties below the median pre-program population, construction employment increases by roughly 13 workers over the first two years, whereas effects in larger counties are statistically indistinguishable from zero. This heterogeneity suggests that the short-run economic returns to federal infrastructure spending may be systematically higher in smaller communities that often have the greatest infrastructure needs and the fewest alternative financing sources.

Taken together, these findings suggest that construction-intensive environmental infrastructure programs can generate meaningful short-run labor-market gains even when stimulus is not the explicit policy goal. Short-run local economic gains represent a quantifiable co-benefit of environmental infrastructure investment that is typically absent from cost-benefit accounting; incorporating them alongside long-run environmental and health benefits would provide a more complete picture of the overall value of large-scale water infrastructure programs.

Related Literature This paper sits at the intersection of three strands of research. The first is a growing empirical macroeconomics and public finance literature that uses geographic cross-sectional variation in fiscal spending to estimate regional employment and income effects. The second is a body of work in urban and spatial economics studying place-based policies and localized labor demand shocks in spatial equilibrium. These literatures provide the conceptual foundation for this paper: the macro literature motivates the focus on local fiscal multipliers, while the urban and spatial literature speaks to how local economic structure shapes the incidence of federal spending. The third strand assesses the Clean Water Act (CWA) and wastewater infrastructure investment, which has largely emphasized environmental quality and health benefits rather than short-run labor-market responses.

A large literature estimates the local incidence of fiscal spending using variation across regions in the timing and intensity of government spending. Seminal work by Nakamura and Steinsson (2014) exploits regional variation in military procurement to estimate an "open-economy relative multiplier" and finds sizable effects of government spending on local activity. Subsequent papers use a wide range of quasi-experimental sources of regional variation — including formula-driven transfers, public health expenditures, disaster relief, and other place-based programs — to estimate local employment and income effects (Chodorow-Reich, 2019; De Ridder et al., 2020; Dupor & Guerrero, 2021; Suárez Serrato & Wingender, 2016). A central lesson is that interpretation hinges on spatial adjustment and the mapping from spending to local outcomes: commuting and migration can reallocate employment across space, sectoral composition can shape where demand is absorbed, and local general equilibrium forces (e.g., wages and rents) can mediate measured effects. Empirical designs therefore emphasize plausibly exogenous variation in the distribution of spending and outcomes tightly linked to the proposed transmission mechanism.

Within this broader literature, evidence on *public investment* — and infrastructure investment in particular — is comparatively limited. Public investment differs from other government purchases because it is lumpy, project-based, and subject to implementation lags (Leeper et al., 2010; Ramey, 2020). Work on ARRA exploits cross-area variation in federal allocations, including "shovel-ready" infrastructure grants, to estimate local employment effects, but findings are mixed and can be sensitive to program category and measurement (Chodorow-Reich et al., 2012; Conley & Dupor, 2013;

Dupor & Mehkari, 2016; Dupor et al., 2023; Wilson, 2012). Related macro evidence also suggests that investment multipliers may differ from consumption multipliers, though results are not uniform across settings (Ilzetzki et al., 2013; Koh, 2025). Closer to the methodological spirit of this paper, Ghomi and Pappa (2026) use European Investment Bank loans to public infrastructure projects across EU countries and implement inverse-probability-score weighted regressions to address selection. They report cumulative output multipliers of about 1.1 after one year and 3.3 after five years, with larger effects under favorable financial conditions.

A recurring limitation in this literature is that researchers often do not directly observe the project-level implementation channel through which public investment is delivered, making it difficult to identify the mechanisms linking spending to local labor-market outcomes (Glaeser & Poterba, 2021). This motivates attention to settings in which the link between spending and local labor demand is tight and measurable. Leduc and Wilson (2013) and Garin (2019) emphasize the value of focusing on construction-intensive highway infrastructure programs and tracking outcomes in sectors that should respond mechanically if spending increases local activity. Related work studies the long-run local effects of large federal investment programs and discrete place-based shocks, including the Tennessee Valley Authority, wartime plant construction, and other infrastructure-driven investments (Garin & Rothbaum, 2025; Kang et al., 2025; Kline & Moretti, 2014; Leff Yaffe, 2020). This paper follows that logic by using project-level construction timing under the CWA Construction Grants Program and by focusing on construction employment, establishments, and wages as the most direct margins of adjustment, before examining whether these sectoral effects propagate to total employment and income.

A related body of work in urban and spatial economics studies place-based policies and localized labor-demand shocks in spatial equilibrium. Garin (2019) is especially close in spirit, as it examines whether federally financed highway infrastructure investment operates as a spatially targeted tool for increasing local payrolls in the construction sector. More broadly, work on local employment multipliers studies whether an exogenous increase in tradable employment generates additional non-tradable jobs locally (Kline & Moretti, 2014; van Dijk, 2017). This paper complements that literature by studying a large federal infrastructure program whose initial demand impulse is concentrated in construction and related sectors, allowing the data to speak to whether a localized construction shock is associated with broader county-level employment and income responses.

A separate literature evaluates the CWA and wastewater infrastructure investment primarily through environmental quality and welfare channels, including changes in water quality and valuation via housing markets (Keiser & Shapiro, 2019). More recent work also emphasizes health and other downstream benefits (Flynn & Marcus, 2023), and how local municipalities are affected by the environmental mandates set forth by the act (Flynn & Smith, 2022; Jerch, 2026). Relative to this literature, evidence on short-run local labor-market effects of CWA wastewater construction is scarce. To the best of my knowledge, the only closely related employment study is Ai et al. (2025), which uses a difference-in-differences design around the 1977 amendments and attainment versus nonattainment status to estimate impacts on sectoral employment. This paper is complementary: it focuses on the short-run macro-local dimension of the same policy domain; how federally financed wastewater construction affects local labor markets and incomes during implementation, and thereby adds an additional lens through which to interpret the benefits of large-scale water infrastructure investment.

The remainder of the paper proceeds as follows. Section 2 describes the CWA Construction Grants Program and the institutional features that link grants to local construction activity. Section 3 describes the data and the staggered difference-in-differences framework. Section 4 reports the main results for construction-sector outcomes, presents heterogeneity analyses, and evaluates threats to identification and robustness. Section 5 examines broader county-level responses in total employment and personal income, including sectoral and spatial patterns consistent with infrastructure-driven spillovers. Section 6 concludes.

2 The Clean Water Act Construction Grants Program

The Federal Water Pollution Control Act Amendments of 1972 (P.L. 92-500) — commonly known as the Clean Water Act (CWA) — constituted a major federal intervention to reduce water pollution and improve water quality across the United States (Congressional Research Service, 2019b). A central statutory objective was to restore and maintain the integrity of the nation's waters, requiring not only regulatory standards on pollutant discharges, but also large-scale public investment in wastewater treatment infrastructure.

This paper focuses on Title II of the 1972 amendments, which authorized the Construction Grants Program — a grant-financed program administered by the EPA that subsidized the construction and upgrading of publicly owned wastewater treatment works. From its inception in fiscal year 1973 through its phase-out in 1990, the program disbursed approximately $52 billion in nominal terms (about $145 billion in 2017 dollars), making it one of the largest public infrastructure investment programs in U.S. history (Congressional Research Service, 2019a). Federal grants covered a fixed share of eligible project costs — originally 75 percent, later reduced to 55 percent in 1981 — with the remainder financed by state and local matching contributions.

Grant allocation proceeded in two stages. The EPA first distributed appropriations to states using a formula based on wastewater treatment needs and population. States then awarded grants to municipalities using priority lists reflecting the severity of nearby water pollution, affected population size, and conservation needs. The CWA explicitly prohibited states from incorporating development objectives or projections of future population growth into these lists (U.S. EPA, 1980), and Flynn and Smith (2022) find no evidence that grants were allocated according to recipient governments' finances or spending preferences. The empirical design exploits differences in the timing of construction starts across counties rather than the fact of receipt, and evaluates identifying assumptions directly using event-time dynamics.[1]

Projects funded under the program proceeded through planning (Step 1), design (Step 2), and construction (Step 3) stages, with Step 3 awards accounting for the bulk of program expenditure. I focus on Step 3 and combined Step 2+3 grants, which include the on-site construction component and most directly translate federal funds into local labor demand.[2]

[1]Counties that ever receive grants differ from counties that never receive grants. I address these differences and assess robustness to alternative comparison groups in Section 4.2.

[2]A natural question is whether the local construction sector anticipated federal assistance when Step 1 and Step 2 grants were awarded. I address this in Section 4.2.

3 Empirical Framework and Data

This section describes the data and empirical design used to estimate local labor-market responses to federally assisted wastewater infrastructure construction under the CWA Construction Grants Program. The analysis links project-level grant records to county-level outcomes and exploits variation in the timing of construction starts across counties in a staggered difference-in-differences framework.

3.1 Data Sources

The analysis combines three primary data sources. First, I use administrative microdata on the universe of Clean Water Act (CWA) Construction Grants Program awards obtained from the EPA's Grant Information Control System through a Freedom of Information Act request. These records report project award dates, construction start and completion dates, project characteristics (e.g., grant step and project type), and funding amounts.[3] I merge the grant files to facility-level information from the Clean Watershed Needs Surveys (CWNS) using a unique facility identifier, and then aggregate to a county-year panel of grant receipt.[4]

The analysis focuses on construction-stage grants (Step 3 and combined Step 2+3 grants). Treatment timing is defined using the *construction start date*. For county i, the treatment year G_i is the first year in which any grant-assisted construction begins.[5] The main treatment indicator is binary: a county-year is treated if at least one grant-assisted construction project begins in that county in year t. Project cost information is used only for back-of-the-envelope calculations that translate estimated employment effects into job multipliers.

County-level employment outcomes come from the Quarterly Census of Employment and Wages (QCEW), which provides annual average employment, establishments, and wage measures by county

[3]As documented in Flynn and Smith (2022), the validity of the grant records is also assessed in Keiser and Shapiro (2019). I cross-check the extracted files against the datasets used in Flynn and Smith (2022), which rely on the same underlying system. Further details are provided in Appendix Section B.

[4]The Clean Watershed Needs Survey (CWNS) is conducted every two to four years and provides an assessment of the capital investments publicly owned wastewater treatment facilities require to meet Clean Water Act standards.

[5]When multiple projects begin in the same county, G_i is the earliest construction start year.

Table 1: Summary statistics of grants under the Clean Water Act Construction Grants Program

	Mean	SD	p5	Median	p95
Panel A. Counties across time					
Counties receiving grants ($N = 2437$)					
Number of grants	4.89	7.01	1.00	3.00	16.00
Project cost (million)	62.05	232.36	0.91	10.91	268.71
Federal grant award (million)	45.26	171.25	0.65	7.96	195.85
Per capita project cost	619.36	704.43	59.26	436.06	1684.97
Per capita federal grant award	454.56	524.23	44.81	322.30	1259.25
Panel B. County-year observations					
County-year observations with grants ($N = 8436$)					
Number of grants	1.41	1.07	1.00	1.00	3.00
Project cost (million)	17.93	57.58	0.37	4.65	72.13
Federal grant award (million)	13.07	42.34	0.26	3.32	52.70
Per capita project cost	178.92	307.80	5.25	85.93	656.24
Per capita federal grant award	131.32	229.88	3.72	61.34	485.22
Panel C. Individual projects/grants					
Total projects/grants ($N = 11921$)					
Project cost (million)	12.68	36.83	0.33	3.91	47.81
Federal grant award (million)	9.25	26.64	0.24	2.81	35.17
Per capita project cost	126.63	257.64	1.73	47.54	509.42
Per capita federal grant award	92.93	192.47	1.23	34.30	378.76

Notes: This table reports summary statistics of federally assisted projects under the CWA Construction Grants Program. Dollar values are in real 2017 USD. Panel A reports aggregations of projects at the county level. Panel B reports county-year observations. Panel C reports summary statistics at the project level. *Project cost* refers to total eligible cost for federal assistance, which includes federal assistance and an additional local cost share.

and industry. The primary outcomes are construction employment, total employment, average weekly wages, and the number of establishments. Because consistent county-by-industry QCEW series begin in 1975, the baseline sample excludes counties whose first grant-assisted construction start occurs prior to 1975. To measure broader local economic activity, I use the BEA Regional Economic Accounts, which report county-level personal income and population starting in 1968, allowing construction of personal income per capita. This outcome complements the QCEW labor-market measures and is also used to summarize pre-program differences between treated and untreated counties. Summary statistics of demographic characteristics and outcome variables of treated and untreated counties are provided in Appendix Table A.1.

Turning to the grant data, Figure A.1 plots the distribution of first treatment years across counties. Treatment cohorts are largest in the mid-to-late 1970s and decline steadily through the 1980s, consistent with the program's front-loaded spending profile; 20.88% of counties never receive a construction-stage grant. Figure A.2 shows the geographic distribution of grants. Panel (a) maps the first year of grant-assisted construction start by county. Timing is staggered across the full program period and geographically dispersed, with no single region dominating early adoption — the variation that underlies the empirical design described below. Panel (b) maps total real federal awards per capita, showing substantial heterogeneity in grant intensity across counties even conditional on treatment.

Table 1 reports summary statistics at the project, county-year, and county levels. Project size is highly right-skewed: the median project carries \$3.91 million in eligible cost and a \$2.81 million federal award, while the 95th percentile reaches \$47.81 million and \$35.17 million, respectively (2017 dollars). At the county-year level (Panel B; $N = 8{,}436$), grant activity is typically episodic, with the median grant-receiving county-year involving a single project and a median federal award of \$4.65 million. Aggregating to the county level over the full program period (Panel A; $N = 2{,}437$ ever-treated counties), the median treated county receives three grants totaling \$10.91 million, while the 95th percentile receives 16 grants totaling \$268.71 million. Together, these patterns motivate an empirical design that exploits variation in construction-start timing for identification and uses dollar magnitudes primarily to interpret effect sizes.

3.2 Empirical Framework

I estimate the effect of initiating grant-assisted construction using a staggered difference-in-differences design (Callaway & Sant'Anna, 2021). Let G_i denote the first year in which grant-assisted construction begins in county i (with $G_i = \infty$ for never-treated counties). The baseline estimand captures the response to a county *first* entering grant-assisted construction. To maintain a single-adoption setting, I drop all county-year observations that follow the start of a second grant-assisted construction episode in the same county.[6] I report dynamic event-study effects around G_i and average

[6]This restriction isolates the initial construction-start episode and avoids conflating subsequent construction waves with the first treatment.

post-treatment effects over two- and four-year windows.

Estimand and aggregation Let Y_{it} denote an outcome for county i in year t. For each cohort g and year $t \geq g$, the group-time average treatment effect on the treated is

$$ATT(g,t) \;=\; \mathbb{E}[\,Y_{it}(g) - Y_{it}(\infty) \mid G_i = g\,]\,, \qquad t \geq g, \tag{1}$$

where $Y_{it}(g)$ denotes the potential outcome at time t if county i first begins grant-assisted construction in year g, and $Y_{it}(\infty)$ denotes the potential outcome if it is never treated. Identification comes from comparing cohort-g counties to counties that are untreated at time t. In the baseline specification, the control set is

$$C_t = \{i : G_i > t\},$$

which includes both not-yet-treated counties (with finite $G_i > t$) and never-treated counties ($G_i = \infty$). The corresponding cohort-time DiD comparison is

$$\widehat{ATT}(g,t) = \Big(\overline{Y}_{g,t} - \overline{Y}_{g,g-1}\Big) - \Big(\overline{Y}_{C_t,t} - \overline{Y}_{C_t,g-1}\Big), \tag{2}$$

where $\overline{Y}_{g,t} = \mathbb{E}[Y_{it} \mid G_i = g]$ and $\overline{Y}_{C_t,t} = \mathbb{E}[Y_{it} \mid i \in C_t]$. I estimate and aggregate these effects using the Callaway and Sant'Anna (2021) DiD estimator (CSDiD). Event-study effects are obtained by re-indexing $ATT(g,t)$ in event time $e = t - g$ and aggregating across cohorts.

I also report average post-treatment effects over horizons $H \in \{2,4\}$, which summarize near-term responses and cover most of the construction cycle in the grants data.[7] Formally,

$$\widehat{ATT}^{H} \;=\; \sum_{g} \omega_g \left(\frac{1}{H} \sum_{h=0}^{H-1} \widehat{ATT}(g, g+h) \right), \tag{3}$$

where ω_g is proportional to the number of counties first treated in cohort g.

[7] Approximately 80 percent of projects are completed within four years of construction start; the average duration is 2.6 years.

Assumptions and threats to identification The staggered DiD design relies on three conditions: (i) no anticipation, (ii) parallel trends for treated cohorts relative to the chosen control group, and (iii) no spillovers across counties. In this setting, *no anticipation* means that outcomes do not change systematically *prior* to the first year of grant-assisted construction start in ways attributable to the forthcoming project (e.g., pre-construction mobilization that occurs before the recorded construction start date). I assess this condition using pre-treatment event-time estimates.

The *parallel trends* condition requires that, absent treatment, outcomes for counties first treated in year g would have evolved similarly to outcomes for the comparison group over the same periods. A potential threat is endogenous targeting: because the allocation process incorporated population and documented treatment needs, treated counties may differ systematically from untreated counties in ways correlated with differential trends in outcomes. However, as discussed in Section 2, the CWA explicitly prohibited states from incorporating development objectives or projections of future population growth into priority lists, which severs the most direct link between economic trajectory and grant receipt. Critically, the empirical design exploits variation in the *timing* of first construction starts rather than the fact of receipt itself, and the determinants of pollution severity and population size do not have an obvious bearing on when a county first entered construction within the program period. Under the baseline control definition $C_t = \{i : G_i > t\}$, the identifying restriction can be stated as

$$\mathbb{E}[\,Y_{it}(\infty) - Y_{i,g-1}(\infty) \mid G_i = g\,] = \mathbb{E}[\,Y_{it}(\infty) - Y_{i,g-1}(\infty) \mid i \in C_t\,], \quad \forall\, t \geq g. \tag{4}$$

A weaker version excludes never-treated counties from the identifying comparison and requires parallel trends only among *eventually treated* counties:

$$\mathbb{E}[\,Y_{it}(\infty) - Y_{i,g-1}(\infty) \mid G_i = g\,] = \mathbb{E}[\,Y_{it}(\infty) - Y_{i,g-1}(\infty) \mid G_i > t,\ G_i < \infty\,], \qquad \forall\, t \geq g. \tag{5}$$

This weaker restriction corresponds to specifications that restrict the control group to not-yet-treated *eventually treated* counties and thus avoid relying on the comparability of never-treated counties. I assess the plausibility of these assumptions using event-study pre-trends and by reporting robustness to alternative comparison-group definitions, including specifications that condition on pre-program county characteristics using a doubly robust estimator (Sant'Anna & Zhao, 2020).

Finally, *no spillovers* requires that treatment in one county does not affect outcomes in other counties. Spillovers are plausibly limited for construction employment because projects are tied to specific wastewater treatment facilities and on-site work, although cross-county contractor activity could attenuate estimated effects if labor is reallocated across county lines. I investigate these concerns in Section 5.1, where I estimate the effect of grant-assisted construction on neighboring never-treated counties within the same commuting zone and find no meaningful spillover effects on either the construction sector or aggregate outcomes.

A separate concern is crowd-out: grant-assisted projects could displace other local construction activity rather than increasing total construction demand. Flynn and Smith (2022) provide evidence that CWA grants increased sewerage capital spending approximately dollar-for-dollar when compliance requirements were binding, consistent with limited crowd-out on the relevant margin.

4 Local Employment Effects

I begin by examining employment effects in counties where federally assisted wastewater construction projects take place, starting with the direct response of the construction sector. Following Garin (2019), the construction sector is a natural first margin of adjustment: grant-assisted projects mechanically raise demand for on-site labor and contracting services, so a construction response is a necessary condition for broader local effects. Table 2 reports baseline CSDiD estimates of the construction-sector labor-market response to the initiation of grant-assisted wastewater construction. The reported coefficients summarize average post-treatment effects over the first two and first four years after the county's initial construction-start event (as defined by the first grant-assisted Step 3 or Step 2+3 construction start in the county). Standard errors are clustered at the county level. Figure 1 presents dynamic event-time estimates for each outcome from three years prior to construction start through four years after. The primary construction-sector outcomes are employment, the number of establishments, and average weekly wages.

Table 2: Local construction sector effects

	Construction emp.		# of establishments		Avg. weekly wage	
	2 yrs	4 yrs	2 yrs	4 yrs	2 yrs	4 yrs
Post avg.	17.851*	7.428	1.613**	0.347	5.685*	2.886
	(8.322)	(11.326)	(0.549)	(0.775)	(2.669)	(3.141)
N	38,992	38,992	38,994	38,994	38,992	38,992
Pre-treatment avg.	778.56	778.56	107.82	107.82	624.64	624.64

Notes: This table reports average post-treatment effects over the first 2 years and first 4 years after the start of grant-assisted construction. Outcomes are from the QCEW and include construction employment, the number of construction establishments, and average weekly construction wages. Standard errors are clustered at the county level and reported in parentheses. *Pre-treatment averages* are computed using 1975 data, the first year with consistent QCEW coverage. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The estimates indicate a meaningful short-run increase in construction activity. Averaged over the first two post-treatment years, construction employment increases by 17.851 workers per county-year. In contrast, the four-year post-treatment averages are smaller and less precisely estimated, suggesting that the construction response is concentrated early in the post-treatment period rather than persisting uniformly over longer horizons. To gauge magnitudes, the pre-treatment mean construction employment is 778.56 workers per county, implying that the two-year average effect corresponds to an increase of 2.29 percent.[8] A back-of-the-envelope translation implies 35.7 job-years per treated county over the first two post-treatment years and 29.7 job-years over the first four post-treatment years. Scaling these job-years by the mean federal award per treated county-year ($17.93 million in 2017 dollars; Panel B of Table 1) yields 0.199 job-years per $100,000 over two years and 0.166 job-years per $100,000 over four years.[9] As a benchmark, Garin (2019) estimates that federal investment in highway construction under the ARRA generated 0.5 job-years per $100,000 over four years. While the institutional setting and spending composition differ, the magnitudes are of the same order; one likely reason Garin's estimates are larger is that ARRA highway spending occurred in the immediate post — Great Recession period, when construction labor markets exhibited substantially greater slack, raising the likelihood that federal spending translated into net job creation.

[8]Because treatment is defined by the first construction-start episode and the estimates are reported in levels, percentage changes are calculated relative to the pre-treatment county mean.

[9]Because the median treated county-year includes one project, scaling instead by the mean federal award per project ($9.25 million; Panel C) yields 0.386 job-years per $100,000 over two years and 0.321 job-years per $100,000 over four years.

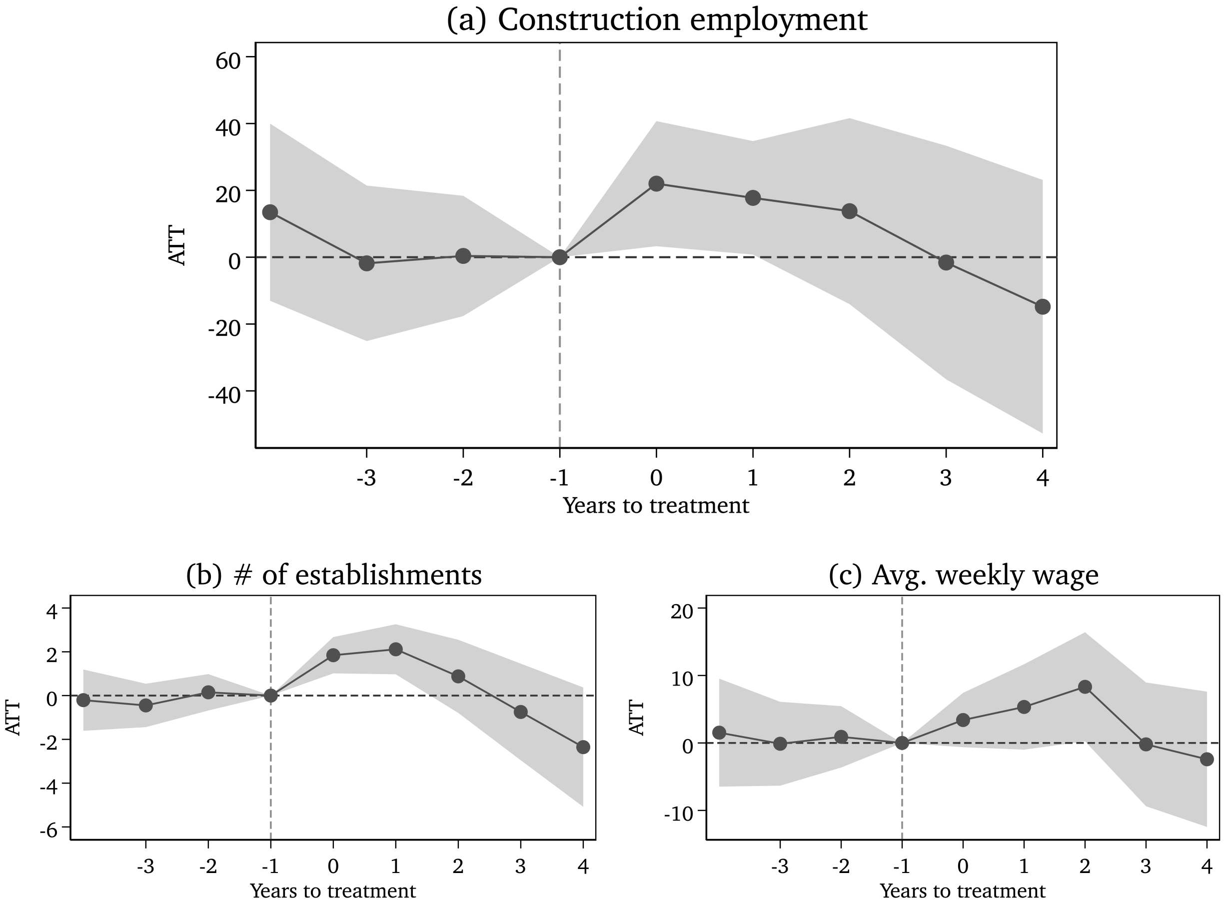

Figure 1: Event-study CSDiD estimates of county construction-sector outcomes (employment, establishments, and average weekly wages) using QCEW data. Event time is measured in years relative to the construction start year. Shaded areas show 95% confidence intervals. Control group includes never-treated and not-yet-treated counties.

Consistent with increased sectoral demand for construction labor, the latter columns of Table 2 and panels (b) and (c) of Figure 1 show increases in both establishments and wages. Two years after construction begins, the number of construction establishments rises by 1.613 and average weekly wages increase by $5.685. Point estimates remain positive through year four, although they are less precisely estimated. Relative to pre-treatment means, these effects correspond to approximately a 1.49 percent increase in establishments and a 0.9 percent increase in weekly wages. Together, these results indicate that CWA construction grants generated a direct, construction-sector-intensive local labor demand response, increasing construction employment and, to a lesser extent, the number of establishments and wages.

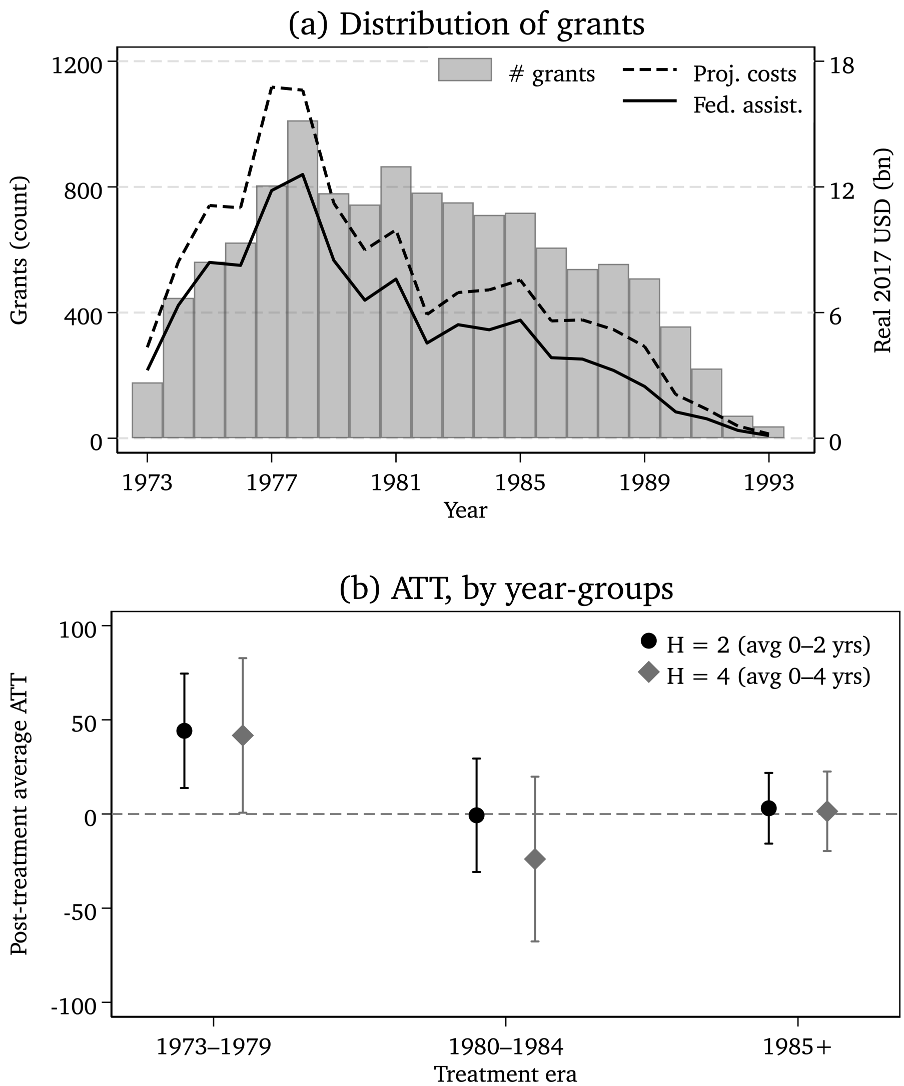

Figure 2: Event-study CSDiD estimates of county construction employment by era. Panel (a) plots the annual distribution of construction starts (left axis) and total project costs and federal assistance (right axis; real 2017 USD billions), aggregated by construction-start year. Panel (b) reports 2-year and 4-year post-average ATT estimates by era (1973–1979, 1980–1984, 1985+); spikes denote 95% confidence intervals.

4.1 Treatment Effect Heterogeneity

4.1.1 Early versus late program eras

Federal assistance under the Construction Grants Program was heavily front-loaded: both the number of construction starts and the volume of federal awards peak in the late 1970s (Figure 2, Panel (a)). This time variation motivates examining whether local employment responses differed across program eras, potentially reflecting changes in program rules as well as learning-by-doing in the specialized construction segment that builds and upgrades wastewater infrastructure.

Guided by data availability and major legislative revisions to the program, I partition counties by the era of their *first* grant-assisted construction start. The three eras are: (i) 1973–1979, (ii) 1980–1984, and (iii) 1985 and later. These cutoffs are chosen to align, with a modest lag, to the timing of program amendments (notably those enacted in 1977 and 1981), allowing for the delay between grant approval and construction start.[10]

The amendments potentially matter for both the scale and composition of projects. The 1977 amendments expanded authorizations, formalized priority rules, and broadened grant eligibility.[11] The 1981 amendments tightened the program along several margins—restricting eligible project categories after 1984, limiting grants for excess reserve capacity, emphasizing cost effectiveness, and reducing the federal share to 55 percent beginning in fiscal year 1985 (Congressional Research Service, 2019b). Together, these changes could affect the size and timing of projects initiated, the stringency of local matching requirements, and the extent to which construction activity is concentrated in compliance-driven upgrades rather than expansions.

Panel (b) of Figure 2 reports estimated effects on construction employment by era. For each era-specific estimate, I restrict the treated sample to counties whose first construction start occurs in that era, and I use counties that are untreated at time t as controls (never-treated counties, with not-yet-treated counties included when available). The results indicate that the baseline construction employment effects are primarily driven by the early program period: for counties first treated between 1973 and 1979, estimated impacts at both the two-year and four-year horizons are positive and statistically significant. In contrast, corresponding estimates for counties first treated in 1980–1984 and after 1985 are smaller and generally not statistically distinguishable from zero. This pattern is consistent with the program's front-loaded intensity and with the possibility that later-era projects were smaller, more incremental, or implemented in a markedly different fiscal and

[10]Because treatment is defined by construction start, the era classification reflects realized construction timing rather than statutory enactment dates.

[11]Changes included provisions for collection systems and single grants combining design and construction for smaller projects.

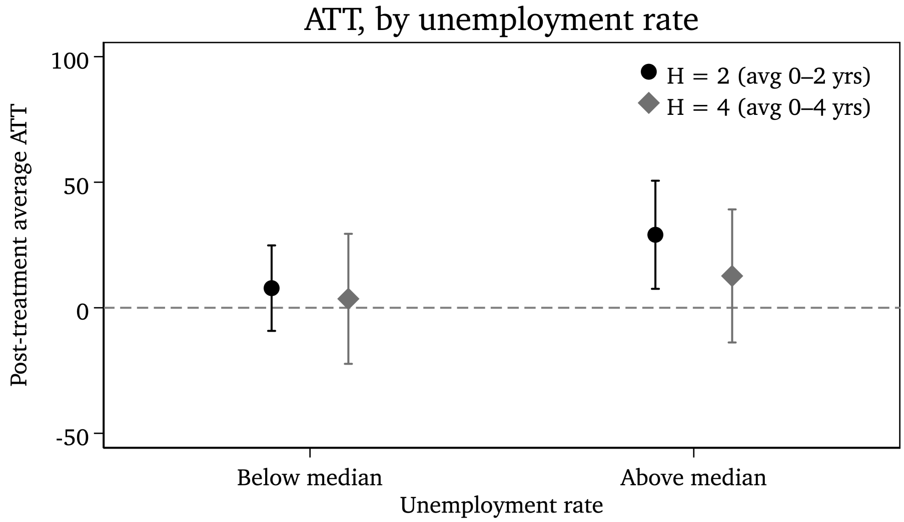

Figure 3: CSDiD estimates of county construction employment by unemployment rate. County-level unemployment rates are from the BLS LAUS and are measured in 1976, the earliest year available. Points report 2-year and 4-year post-treatment average ATT estimates. The median county unemployment rate in 1976 is 6.25 percent. Spikes denote 95% confidence intervals.

regulatory environment.

4.1.2 Slack in the economy

A recurring hypothesis in the fiscal multipliers literature is that government spending generates larger real effects when there is greater economic slack (Auerbach & Gorodnichenko, 2012; Nakamura & Steinsson, 2014; Owyang et al., 2013; Ramey & Zubairy, 2018). To examine this mechanism in the context of CWA construction starts, I use cross-county variation in local labor market slack measured by county unemployment rates from the BLS Local Area Unemployment Statistics (LAUS). Because LAUS county unemployment data begin in 1976, I use the 1976 unemployment rate—the earliest available year—as a predetermined proxy for baseline slack.[12] I split counties into two groups based on whether their 1976 unemployment rate is above or below the median and estimate the baseline CSDiD specification separately for each group.[13]

[12] County-level LAUS unemployment data prior to 1990 are not publicly posted and were obtained from the BLS by request. The accompanying documentation notes a break in the LAUS series in 1990. This is unlikely to affect my analysis because I use only cross-sectional variation within a single year (1976), rather than changes over time.

[13] The median county-level unemployment rate in 1976 is 6.25 percent.

Figure 3 shows clear state dependence. The construction-employment response is positive and statistically significant in counties with above-median unemployment, while estimates for below-median counties are close to zero and statistically indistinguishable from no effect. This pattern is consistent with the slack-based mechanism emphasized in the fiscal multiplier literature: when local labor markets have greater slack, public construction spending can expand employment with limited crowding out of private activity. In contrast, when unemployment is low and labor markets are tighter, additional public demand is more likely to bid workers away from existing jobs or face capacity constraints, yielding smaller net employment gains.

4.1.3 County size

A natural additional dimension of heterogeneity is baseline county size. Because CWA construction grants finance large, discrete capital projects, the same dollar amount may represent a relatively larger shock in smaller counties than in larger, more diversified local economies. In addition, smaller counties may have had more limited pre-existing wastewater infrastructure and fewer local funding sources prior to the CWA, implying that federal assistance could translate into a more visible construction-sector response.

To examine this margin, I divide counties into two groups based on the median of average pre-treatment population over 1969–1971.[14] I then estimate the baseline CSDiD specification separately within each group. Figure 4 presents the corresponding subgroup ATT estimates, and the corresponding 2-year and 4-year subgroup ATT estimates are reported in Appendix Table A.2. The results indicate substantial heterogeneity. For counties below the median population, construction employment increases by 12.99 workers on average over the first two post-treatment years (s.e. 5.91). Relative to a pre-treatment mean of 127.25 construction workers, this corresponds to an increase of approximately 10.2 percent. Four-year average effects are similar in magnitude, though less precisely estimated.

In contrast, counties above the median population show estimates that are economically small and statistically indistinguishable from zero. While the point estimates for larger counties are impre-

[14] The median county has an average 1969-1971 population of approximately 18,900.

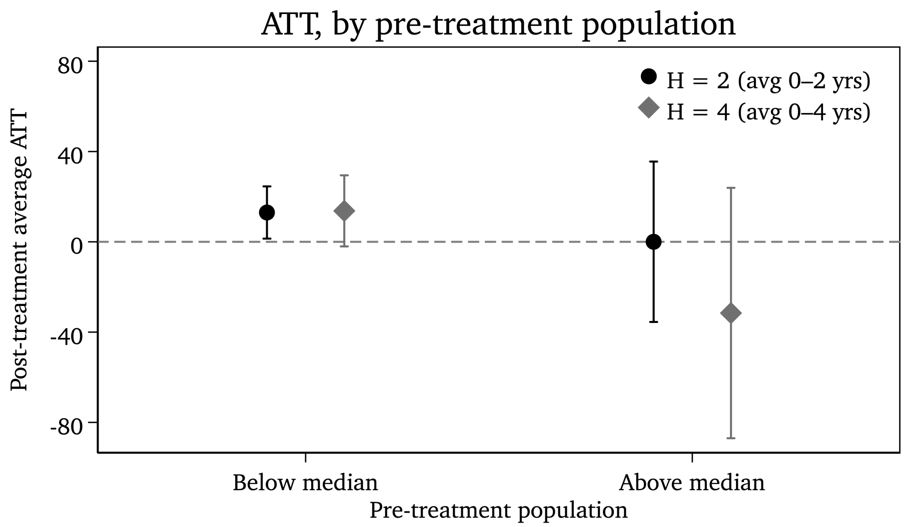

Figure 4: CSDiD estimates of county construction employment by pre-program (1969–1971) average county population. Population data are from the BEA Regional Economic Accounts. Points report 2-year and 4-year post-treatment average ATT estimates. Spikes denote 95% confidence intervals.

cise—partly reflecting the smaller effective sample after restricting to initial construction episodes—the divergence in magnitudes suggests that the local employment response is concentrated in smaller counties.

This pattern is consistent with larger proportional local impacts in smaller counties, where a grant-financed wastewater project represents a larger shock relative to the size of the existing construction sector and local labor market. The findings therefore indicate that the short-run construction effects of the CWA Construction Grants Program were disproportionately borne by smaller, less populous counties. This heterogeneity reflects differential local impact rather than differential treatment intensity per se: similar projects can generate larger proportional effects in smaller local economies. The estimates do not by themselves distinguish scaling effects from differences in project intensity across county types.

4.2 Threats to Internal Validity and Robustness

Anticipation A key requirement for the staggered DiD design is that outcomes do not respond *before* grant-assisted construction begins (no anticipation). In this setting, economically meaningful

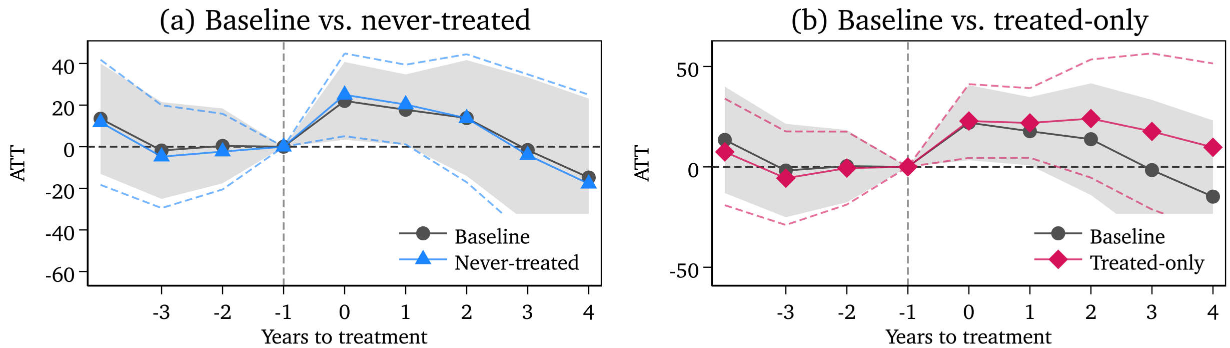

Figure 5: Event-study CSDiD estimates of county construction employment under alternative control group definitions. Panel (a) compares the baseline specification (dark line; never-treated and not-yet-treated controls) to estimates using only never-treated counties as controls (blue line). Panel (b) compares the baseline specification (dark line) to estimates using only not-yet-treated counties as controls (red line). Shaded areas and dotted lines denote 95% confidence intervals.

anticipation in construction employment would most plausibly arise if (i) the recorded construction start date systematically lags actual on-site activity (e.g., early mobilization or site preparation), or if (ii) pre-construction phases generated sizable construction labor demand. Both mechanisms appear limited in this context: the analysis defines treatment using Step 3 (and Step 2+3) construction starts, and pre-construction grant phases are small relative to construction-stage awards. Consistent with this interpretation, event-study estimates for construction-sector outcomes in Figure 1 show that pre-treatment coefficients are small and statistically indistinguishable from zero, providing no evidence of differential pre-trends or anticipatory effects prior to the construction start year.

As an additional check, I re-estimate the event study using only never-treated counties as the comparison group. This exercise removes not-yet-treated counties from the control set, addressing the concern that counties treated later might exhibit anticipatory adjustments as construction starts approach. Panel (a) of Figure 5 shows that the estimated post-treatment effects and the absence of pre-trends are similar to the baseline specification, suggesting that the main results are not driven by contamination of the not-yet-treated control group.

Targeting and selection A second concern is endogenous targeting: counties receiving grant-assisted projects may differ systematically from counties that never receive assistance. As described in Section 2, municipalities applied for funding through state priority lists that were required to focus on documented pollution-control needs and were prohibited from incorporating factors such as geographic location, projections of future population growth, or broader development objectives unrelated to pollution abatement; these lists were reviewed and approved by the EPA (U.S. EPA, 1980). Nonetheless, selection remains a plausible threat because allocation rules were explicitly related to population and measured treatment needs. Consistent with this, Appendix Table A.1 shows that pre-program characteristics differ between ever-treated and never-treated counties. Such cross-sectional differences would bias estimates if they were associated with differential counterfactual trends.

Two pieces of evidence suggest that this concern is unlikely to drive the results. First, the baseline event-study patterns show no evidence of differential pre-trends, which directly addresses the key threat posed by systematic selection on levels. Second, I implement robustness exercises that alter the comparison group and flexibly account for observed baseline differences.

The first exercise restricts the control group to *not-yet-treated* counties. Under this design, identification comes from comparisons among counties that eventually receive treatment at different times, rather than comparisons between ever-treated and never-treated counties. Panel (b) of Figure 5 shows that the estimated effects are very similar to the baseline. This implies that the results are not driven by unobserved, time-varying differences between ever-treated and never-treated counties and that the estimated effects are robust when inference relies primarily on timing variation within the ever-treated sample.

The second exercise introduces county-level covariates using the doubly robust IPW-OLS (DRIPW) implementation of the Callaway–Sant'Anna estimator (Sant'Anna & Zhao, 2020). *Demographic and economic controls* include the mean population and personal income per capita over 1969–1971, together with their growth rates over 1969–1971. *Construction controls* include the mean construction-sector income per capita over 1969–1971 and its growth rate over 1969–1971. Table 3 reports average post-treatment effects over the first two years (columns 1–2) and the first four years (columns 3–4). Columns (1) and (3) include demographic and economic controls only, and columns (2) and (4) additionally include construction controls. Across specifications, the estimated treatment effects remain close to the baseline magnitudes, although precision declines in some columns. This loss of

Table 3: Construction employment effects, including controls

	2 yrs		4 yrs	
	(1)	(2)	(3)	(4)
Post avg.	16.436	16.946*	7.545	8.875
	(8.610)	(8.605)	(10.678)	(10.718)
N	38,946	38,235	38,946	38,235
Pre-treatment avg.	778.56	778.56	778.56	778.56
Demographic, economic controls	Yes	Yes	Yes	Yes
Construction controls		Yes		Yes

Notes: This table reports average post-treatment effects over the first 2 years and first 4 years after the start of grant-assisted construction. The outcome variable is construction employment. Estimates are obtained using the CSDiD estimator with covariates via the doubly robust IPW-OLS (DRIPW) approach, which combines stabilized inverse-probability weighting with outcome regression (Sant'Anna & Zhao, 2020). *Demographic and economic controls* include the mean population and personal income per capita over 1969–1971, and the corresponding growth rates over 1969–1971. *Construction controls* include the mean construction-sector income per capita over 1969–1971 and its growth rate over 1969–1971. Standard errors are clustered at the county level and reported in parentheses. *Pre-treatment averages* are computed using 1975, the first year with consistent QCEW coverage. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

precision is expected: the DRIPW procedure relies on estimated propensity scores and outcome regressions, and incorporating several baseline covariates can increase sampling variability (especially when treatment probabilities are near zero or one for some counties), which widens standard errors even when point estimates are stable. Overall, these exercises indicate that the main results are not sensitive to selection concerns captured by baseline observables or to the inclusion of never-treated counties in the comparison group.

5 Aggregate Effects of the Construction Grants Program

This section examines whether grant-assisted wastewater infrastructure projects generate broader responses along two margins: beyond the construction sector within treated counties, and across county boundaries within commuting zones. While Section 4 focuses on construction labor markets, total employment and personal income capture the extent to which these projects translate into wider local economic activity. I also summarize sectoral responses to assess whether employment gains concentrate in industries plausibly linked to infrastructure construction and its local supply chain. Finally, I investigate whether grant-assisted construction in one county affects labor-market

Table 4: Aggregate effects

	Total employment			Per-capita personal income		
	2 yrs	4 yrs	6 yrs	2 yrs	4 yrs	6 yrs
Post avg.	316.315***	284.093***	223.525*	0.056	0.105*	0.093
	(74.328)	(84.881)	(108.579)	(0.034)	(0.041)	(0.050)
N	38,746	38,746	38,746	63,898	63,898	63,898
Pre-treatment avg.	12611.59	12611.59	12611.59	19.94	19.94	19.94

Notes: This table reports average post-treatment effects over the first 2, 4, and 6 years after the adjusted treatment year, defined as two years prior to construction start to account for general pre-construction economic activity. *Total employment* is from the QCEW and *per-capita personal income* is from the BEA Regional Economic Accounts. Pre-treatment averages are computed using 1975 data for total employment, and 1973 data for per-capita personal income. Standard errors are clustered at the county level and reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

outcomes in neighboring counties within the same commuting zone.

For aggregate and sectoral outcomes, I adjust treatment timing to account for the lag between grant award and the start of on-site construction. Sectors outside construction may respond during the pre-construction phase through planning, engineering, procurement, and subcontracting relationships that precede physical construction (Leeper et al., 2010; Ramey, 2020). To capture this buildup, I shift the treatment date earlier by two years, which corresponds to the median lag between grant award and construction start in the grants data. Formally, I define $\widetilde{G}_i = G_i - 2$, where G_i is the construction-start treatment year used in the baseline analysis. This timing adjustment is used only in this section; the baseline construction-sector estimates remain anchored on G_i. Importantly, this alternative timing is also informative for identification: event-study dynamics around $\widetilde{G}_i$ provide a diagnostic check on whether grant awards were systematically aligned with contemporaneous county-wide economic movements. Consistent with the identifying assumptions, the pre-treatment event-time coefficients around $\widetilde{G}_i$ are small and statistically indistinguishable from zero (see Figure 6).[15]

Table 4 reports baseline CSDiD estimates using the adjusted treatment timing. Total employment rises by 316.3 workers per county-year on average over the first two post-treatment years (s.e. 74.3), a 2.51 percent increase relative to the pre-treatment mean of 12,611.6. Interpreted in job-

[15] In Appendix Section C.1, I compare event study estimates across alternative timing adjustments and find that the dynamic responses are qualitatively robust across specifications.

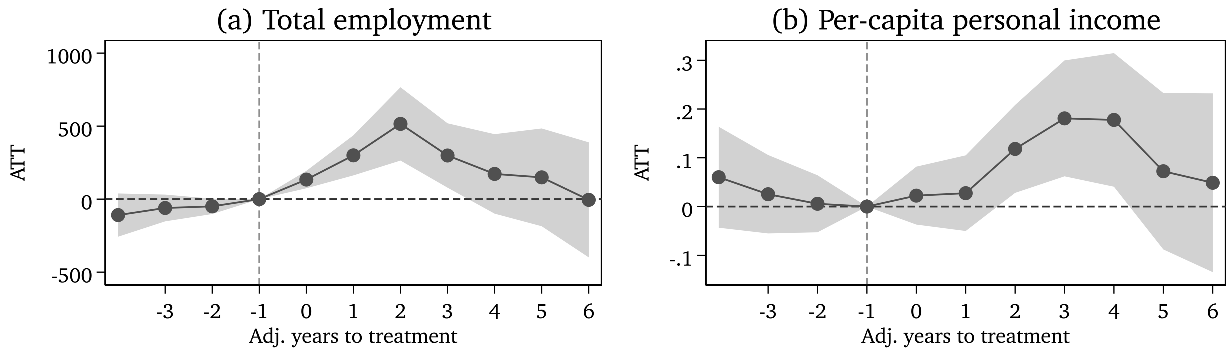

Figure 6: Event-study CSDiD estimates of aggregate effects using adjusted treatment timing. Panel (a) reports total employment from the QCEW and Panel (b) reports per-capita personal income from the BEA. Treatment timing is shifted two years earlier than construction start to capture general pre-construction economic activity. Shaded areas denote 95% confidence intervals.

years, the implied cumulative effects are 632.6 job-years per treated county over two years, 1,136.4 job-years over four years, and 1,341.2 job-years over six years. Scaling by the mean federal award per treated county-year ($17.93 million in 2017 dollars; Panel B of Table 1), these correspond to 3.53, 6.34, and 7.48 job-years per $100,000 over the 2-, 4-, and 6-year windows (equivalently, about $28,300, $15,800, and $13,400 per job-year, respectively).[16] Per-capita personal income also increases, with the strongest evidence appearing in the four-year window (0.105, s.e. 0.041), consistent with broader gains beyond the construction sector.

The personal income per capita estimates can be used to construct a reduced-form benchmark for the local income return to federal assistance. Interpreting the coefficients in Table 4 as changes in thousands of 2017 dollars per person per year, the implied cumulative per-capita income gain over horizon H is $\widehat{\beta}_H \times H \times 1000$. Scaling by the mean federal award per treated county-year episode expressed on a per-capita basis ($178.92; Panel B of Table 1) yields back-of-the-envelope income-based local multiplier of approximately 2.35 using the four-year estimate.[17] Although the two-year

[16]The cost-per-job-year estimates generally exceed those reported in Chodorow-Reich (2019) for the 2009 American Recovery and Reinvestment Act (ARRA), which range from approximately 0.70 to 3.51 job-years per $100,000 in 2017 dollars, as well as those from studies of defense spending, which are around 0.3 job-years per $100,000 in 2017 dollars (**briganti_2025**; Park et al., 2026). Two factors may help explain the larger estimates here. First, the CWA Construction Grants Program directly subsidized on-site wastewater construction, creating a tight and immediate link between federal spending and local labor demand that may generate larger employment responses per dollar than broader fiscal transfers or defense procurement, where the translation from spending to local jobs is more diffuse. Second, the program operated primarily during the 1970s, when labor markets in many recipient counties likely had greater slack than in the post–Great Recession period covered by much of the ARRA literature, suggesting that federal spending may have translated more directly into net job creation with more limited crowding out of private activity.

[17]Calculated as $(0.105 \times 4 \times 1000)/178.92 \approx 2.35$.

and six-year income estimates are less precise, they provide suggestive bounds of 0.63 over two years and 3.12 over six years.[18] Taken together, these estimates suggest that wastewater infrastructure investment generated economically meaningful cross-sector spillovers in the short run.

To investigate sectoral heterogeneity in transmission, I estimate dynamic responses for county-by-sector employment using the adjusted treatment date. The resulting event-study estimates are reported in Appendix Figure A.3. The sectoral estimates suggest that employment gains are concentrated in industries plausibly connected to infrastructure investment and local demand spillovers (e.g., public utilities and transportation, services, manufacturing, and wholesale/retail trade) while sectors less directly related to wastewater construction, such as mining and agriculture, show little systematic response. This pattern provides a qualitative check that the aggregate employment effects are not driven by unrelated sectoral trends.

5.1 Spatial Spillovers

A potential concern with county-level estimates is that CWA construction grants may generate employment effects that spill across county boundaries, either inflating or attenuating the measured local impact. This section studies such spillovers both as a robustness check on the county-level design and as evidence on the geographic scope of infrastructure-driven employment gains.

I use commuting zones to define neighboring labor markets. Introduced by Tolbert and Sizer (1996) and based on cross-county commuting patterns, commuting zones are designed to capture economically integrated local labor markets: counties within the same commuting zone exhibit strong within-market labor-market dependence, whereas counties in different commuting zones are comparatively less connected. Relative to Core-Based Statistical Areas (CBSAs), commuting zones provide more complete geographic coverage of the United States, including rural and non-metropolitan areas that are well represented in the CWA Construction Grants Program (Dupor & McCrory, 2018; Park et al., 2026). Commuting zone delineations are from the USDA Economic Research Service, which harmonizes the Tolbert and Sizer (1996) boundaries to county FIPS codes.

[18] This back-of-the-envelope calculation benchmarks cumulative income gains within treated counties per dollar of federal awards and does not account for financing, inter-county spillovers, general equilibrium adjustments, or the full incidence of benefits and costs.

Construction grant receipt in one county may affect labor-market outcomes in neighboring counties within the same commuting zone through two competing channels, with opposite implications for the interpretation of the county-level estimates. First, if commuting-zone-based contractors and subcontractors expand hiring across county lines, neighboring counties may absorb part of the regional employment gain, causing the county-level estimate to understate the program's true effect on the local labor market. Second, if workers relocate from neighboring counties to fill construction demand at project sites, the county-level estimate may partly capture spatial reallocation rather than net job creation, in which case the true aggregate employment effect would be smaller than the county-level estimate suggests. Distinguishing between these channels has direct implications for whether the county-level estimates should be interpreted as lower or upper bounds on the program's true employment impact.

To assess the importance of these channels, I apply the same CSDiD estimator to a sample restricted to counties that never directly receive a CWA construction grant. For each never-treated county, I define treatment as the first year any other county in the same commuting zone initiates grant-assisted construction. The control group consists of never-treated counties whose commuting zone also contains no treated county, augmented by not-yet-CZ-exposed counties. Because the sample excludes all own-grant recipients by construction, any estimated effect on construction employment, establishments, or wages in these counties cannot reflect direct grant activity and must instead capture spillovers from neighboring treated counties.

Table 5 reports the spillover estimates. Across all specifications, the results indicate no meaningful employment effects in neighboring counties within the same commuting zone. Construction employment, establishments, and average weekly wages in never-treated counties have point estimates close to zero and are statistically indistinguishable from zero at conventional significance levels over both the two- and four-year horizons. Aggregate outcomes tell a similar story: total employment and per-capita personal income in neighboring counties show no systematic response to grant-assisted construction in adjacent counties. These null results are consistent with the nature of the CWA Construction Grants Program as a highly place-specific infrastructure policy. Because projects are tied to particular wastewater treatment facilities at fixed locations, the labor-demand shock appears to be absorbed largely within the directly treated county rather than diffusing across county lines. Taken

Table 5: Spatial spillovers

	Construction			Aggregate	
	Emp.	Estab.	Wkly wage	Total emp.	Income pc.
Post avg. (2 yrs)	-4.796	0.099	3.705	25.970	0.071
	(9.904)	(0.385)	(7.727)	(24.110)	(0.125)
Post avg. (4 yrs)	-3.981	0.022	3.284	9.854	0.090
	(11.652)	(0.506)	(8.150)	(28.308)	(0.131)
N	8,300	8,300	8,300	7,240	19,680
Pre-treatment avg.	176.68	27.15	555.21	2,567.15	19.05

Notes: This table reports average post-treatment effects over the first 2 and first 4 years after the first grant-assisted construction start in a neighboring county within the same commuting zone. Outcomes include construction employment, the number of construction establishments, average weekly construction wages, total employment, and per-capita personal income. Treatment timing is adjusted by two years earlier than construction start for the latter two aggregate outcomes. Standard errors are clustered at the county level and reported in parentheses. Pre-treatment averages are computed using 1975 for QCEW outcomes and 1973 for per-capita personal income. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

together, the spillover estimates support the interpretation of the county-level results in Section 4 as reflecting genuine local employment gains rather than spatial reallocation from neighboring areas, and they suggest that the county-level estimates are not meaningfully biased by cross-county contractor or worker mobility.

Although the results suggest minimal spillovers across county boundaries, one may be interested in whether the program generated effects at broader geographic aggregations. A commuting-zone-level analysis is not pursued because the Construction Grants Program achieves near-universal geographic penetration even at the commuting-zone level, leaving insufficient statistical power for a credible parallel-trends design. Of the 741 commuting zones in my sample, 687 (92.7 percent) contain at least one treated county, leaving only 54 never-treated commuting zones as potential controls. These never-treated commuting zones are structurally atypical: nearly half (48 percent) consist of a single county, making them systematically smaller and less economically integrated than ever-treated commuting zones. Moreover, by 1975, 46.7 percent of all commuting zones already contain a treated county, so the pool of not-yet-treated commuting zones is nearly exhausted precisely in the cohorts where most identification would occur. The county-level neighbor design avoids these problems: the unit of observation remains the county, and the control pool of never-directly-treated counties residing in untreated commuting zones is substantially larger and more comparable to the treated group.

6 Conclusion

This paper studies the short-run local labor-market effects of the Clean Water Act (CWA) Construction Grants Program, one of the largest federal infrastructure investment programs in U.S. history. Using project-level construction start dates linked to county-level labor-market outcomes, I estimate dynamic treatment effects in a staggered difference-in-differences framework. The identifying logic is straightforward: if federally financed infrastructure spending stimulates local economic activity, the most immediate margin of adjustment should be the construction sector, with broader spillovers to follow.

The results support this mechanism. Counties initiating grant-assisted wastewater construction experience statistically meaningful increases in construction employment, along with increases in construction establishments and average weekly wages, consistent with a localized labor-demand shock in the sector that directly implements the projects. Beyond the construction sector, total employment rises following project initiation and per-capita personal income increases over medium horizons, with employment gains concentrating in plausibly related industries such as transportation, public utilities, manufacturing, and wholesale and retail trade. These effects are disproportionately borne by smaller counties: in counties below the median pre-program population, construction employment rises by roughly ten percent relative to baseline, whereas effects in larger counties are economically small and statistically indistinguishable from zero. This pattern is consistent with a scaling mechanism in which a federally financed project represents a larger shock relative to the existing construction base in smaller local economies.

These findings speak to a growing literature on the local economic effects of public infrastructure investment and place-based policy. The Construction Grants Program was not designed as countercyclical stimulus; its primary objectives were regulatory compliance and pollution abatement. Nonetheless, the evidence indicates that large-scale, construction-intensive public investment can generate meaningful short-run labor-market and income responses through a direct sectoral channel. This provides historical empirical grounding for the intuition, visible in ARRA- and IIJA-era infrastructure expansions, that infrastructure investment can support local economic conditions even when stimulus is not the explicit program goal.

The results also speak to the ongoing debate over the costs and benefits of large environmental infrastructure programs. Prior work documents longer-run benefits of the Construction Grants Program through improved water quality, capitalization into housing values, and health improvements, but also suggests that program costs may exceed measured benefits (Flynn & Marcus, 2023; Keiser & Shapiro, 2019; Keiser et al., 2019). The short-run local labor-market gains documented here represent a quantifiable economic co-benefit that is typically absent from cost-benefit accounting. Incorporating such effects alongside long-run environmental and public health benefits would provide a more complete picture of the overall value of large-scale water infrastructure investment and may help reconcile the tension between substantial program costs and measured benefits.

Several limitations are important for interpretation and point to directions for future work. The estimates are local and partial-equilibrium in nature, capturing within-county changes relative to comparison counties. As is standard in the local labor markets literature, these relative effects should not be directly interpreted as aggregate government spending multipliers: they do not account for general equilibrium adjustments such as commuting, migration, and local prices that mediate how measured employment gains translate into economy-wide welfare, nor do they capture the macroeconomic effects of financing the program. More broadly, integrating the short-run local economic effects documented here with longer-run environmental and public health benefits, as well as financing and incidence considerations, would help clarify when large-scale water infrastructure investment is simply costly versus costly but socially productive.

References

Ai, J., Wan, X., Yang, A., & Zhang, W. (2025). The employment effects of environmental regulations: Evidence from the 1977 clean water act amendments. *Working Paper*.

Auerbach, A. J., & Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, *4*(2), 1–27.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Chodorow-Reich, G. (2019). Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, *11*(2), 1–34.

Chodorow-Reich, G., Feiveson, L., Liscow, Z., & Woolston, W. G. (2012). Does state fiscal relief during recessions increase employment? evidence from the american recovery and reinvestment act. *American Economic Journal: Economic Policy*, *4*(3), 118–145.

Congressional Research Service. (2019a, March 6). *Funding for EPA water infrastructure: A fact sheet* (legislation No. R43871).

Congressional Research Service. (2019b, April 10). *Water infrastructure financing: History of EPA appropriations* (legislation No. 96-647).

Conley, T. G., & Dupor, B. (2013). The american recovery and reinvestment act: Solely a government jobs program? *Journal of Monetary Economics*, *60*(5), 535–549.

De Ridder, M., Hannon, S., & Pfajfar, D. (2020, August). The multiplier effect of education expenditure.

Dupor, B., & Guerrero, R. (2021). The aggregate and local economic effects of government financed health care [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ecin.12951]. *Economic Inquiry*, *59*(2), 662–670.

Dupor, B., Karabarbounis, M., Kudlyak, M., & Saif Mehkari, M. (2023). Regional consumption responses and the aggregate fiscal multiplier. *The Review of Economic Studies*, *90*(6), 2982–3021.

Dupor, B., & McCrory, P. B. (2018). A cup runneth over: Fiscal policy spillovers from the 2009 recovery act [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ecoj.12475]. *The Economic Journal*, *128*(611), 1476–1508.

Dupor, B., & Mehkari, M. S. (2016). The 2009 recovery act: Stimulus at the extensive and intensive labor margins. *European Economic Review*, *85*, 208–228.

Flynn, P., & Marcus, M. (2023). A watershed moment: The clean water act and birth weight. *Journal of Human Resources*.

Flynn, P., & Smith, T. (2022). Rivers, lakes and revenue streams: The heterogeneous effects of clean water act grants on local spending. *Journal of Public Economics*, *212*, 104711.

Garin, A. (2019). Putting america to work, where? evidence on the effectiveness of infrastructure construction as a locally targeted employment policy. *Journal of Urban Economics*, *111*, 108–131.

Garin, A., & Rothbaum, J. (2025). The long-run impacts of public industrial investment on local development and economic mobility: Evidence from world war II*. *The Quarterly Journal of Economics*, *140*(1), 459–520.

Ghomi, M., & Pappa, E. (2026). Stimulating avenues: EIB loans and returns to public investment. *Working Paper*.

Glaeser, E. L., & Poterba, J. M. (Eds.). (2021, November 11). *Economic analysis and infrastructure investment*. University of Chicago Press.

H.R.3684 - 117th Congress (2021-2022): Infrastructure Investment and Jobs Act (2021, November 15).

Ilzetzki, E., Mendoza, E. G., & Végh, C. A. (2013). How big (small?) are fiscal multipliers? *Journal of Monetary Economics*, *60*(2), 239–254.

Jerch, R. L. (2026). Municipal governments under the clean water act. *Journal of the Association of Environmental and Resource Economists*.

Kang, N., Rad, M. R., Nayga Jr., R. M., Hrozencik, A., & Perez-Quesada, G. (2025). The reclamation act and regional growth: How canals amplified the impact of USBR dams in the arid west [Num Pages: 54]. *Working Paper*.

Keiser, D. A., Kling, C. L., & Shapiro, J. S. (2019). The low but uncertain measured benefits of US water quality policy. *Proceedings of the National Academy of Sciences*, *116*(12), 5262–5269.

Keiser, D. A., & Shapiro, J. S. (2019). Consequences of the clean water act and the demand for water quality. *The Quarterly Journal of Economics*, *134*(1), 349–396.

Kline, P., & Moretti, E. (2014). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority *. *The Quarterly Journal of Economics*, *129*(1), 275–331.

Koh, K. W. (2025, October 30). Regional government consumption and investment multipliers.

Leduc, S., & Wilson, D. (2013). Roads to prosperity or bridges to nowhere? theory and evidence on the impact of public infrastructure investment. *NBER Macroeconomics Annual*, *27*, 89–142.

Leeper, E. M., Walker, T. B., & Yang, S.-C. S. (2010). Government investment and fiscal stimulus. *Journal of Monetary Economics*, *57*(8), 1000–1012.

Leff Yaffe, D. (2020). *Essays on the effects of highway spending* [Doctoral dissertation, UC San Diego].

Nakamura, E., & Steinsson, J. (2014). Fiscal stimulus in a monetary union: Evidence from US regions. *American Economic Review*, *104*(3), 753–792.

Owyang, M. T., Ramey, V. A., & Zubairy, S. (2013). Are government spending multipliers greater during periods of slack? evidence from twentieth-century historical data. *American Economic Review*, *103*(3), 129–134.

Park, G., Zhou, X., & Zubairy, S. (2026, January). Subcontracting in federal spending: Micro and macro implications.

Ramey, V. A. (2020, July). The macroeconomic consequences of infrastructure investment.

Ramey, V. A., & Zubairy, S. (2018). Government spending multipliers in good times and in bad: Evidence from US historical data. *Journal of Political Economy*, *126*(2), 850–901.

Sant'Anna, P. H. C., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, *219*(1), 101–122.

Suárez Serrato, J. C., & Wingender, P. (2016, July). Estimating local fiscal multipliers.

Tolbert, C. M., & Sizer, M. (1996). *U.s. commuting zones and labor market areas: A 1990 update* (Num Pages: 158).

U.S. EPA. (1980). Handbook of procedures.

U.S. EPA. (2024). *2022 clean watersheds needs survey* (Report to Congress No. EPA 832-R-24-002).

van Dijk, J. J. (2017). Local employment multipliers in u.s. cities. *Journal of Economic Geography*, *17*(2), 465–487.

Wilson, D. J. (2012). Fiscal spending jobs multipliers: Evidence from the 2009 american recovery and reinvestment act. *American Economic Journal: Economic Policy*, *4*(3), 251–282.

Appendices

A Supplemental Tables and Figures

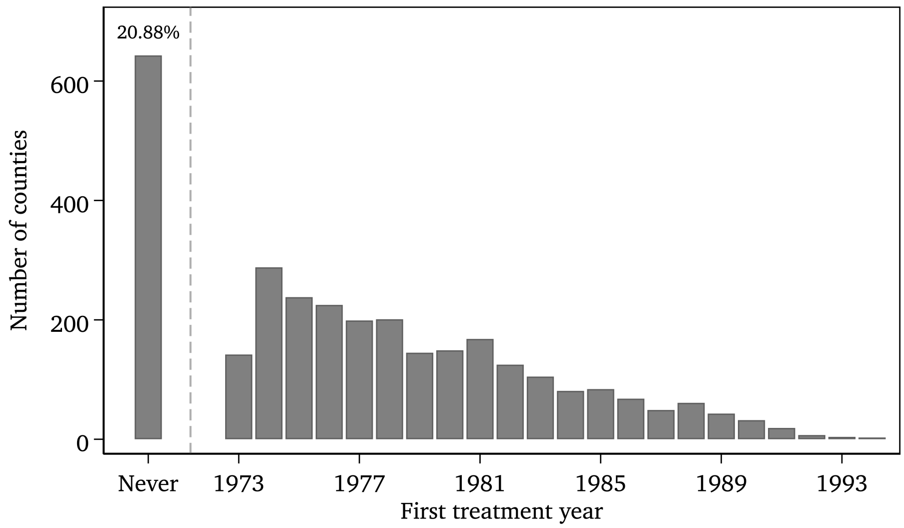

Figure A.1: County-level distribution of first treatment year. *Treatment* is defined as the first year in which grant-assisted construction begins under the CWA Construction Grants Program. *Never* denotes counties that never receive federal assistance under the program.

(a) First year of grant-assisted construction start, by county

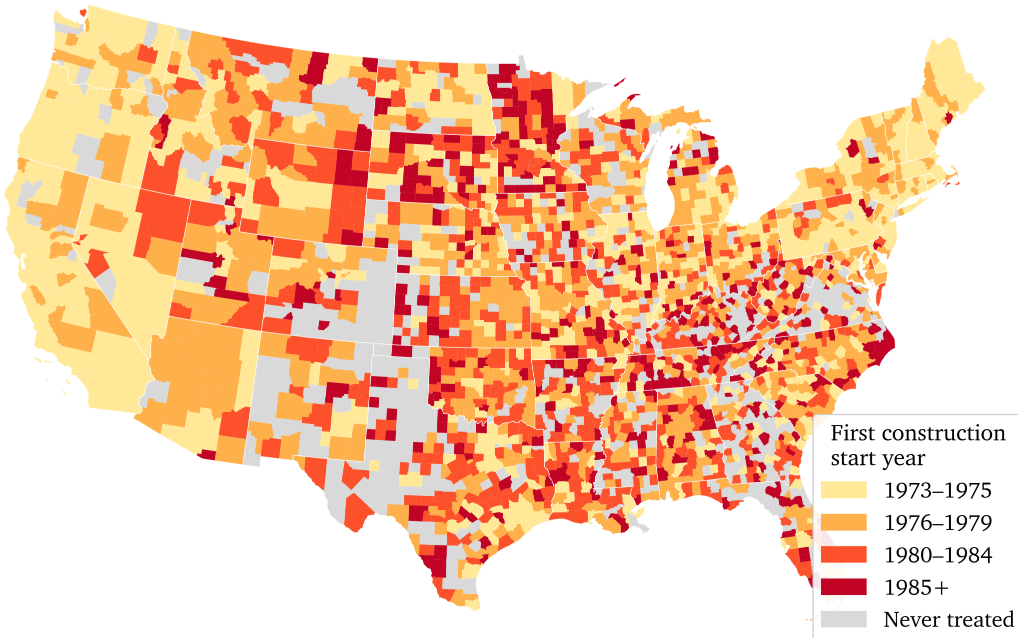

(b) Total real federal award per capita, by county

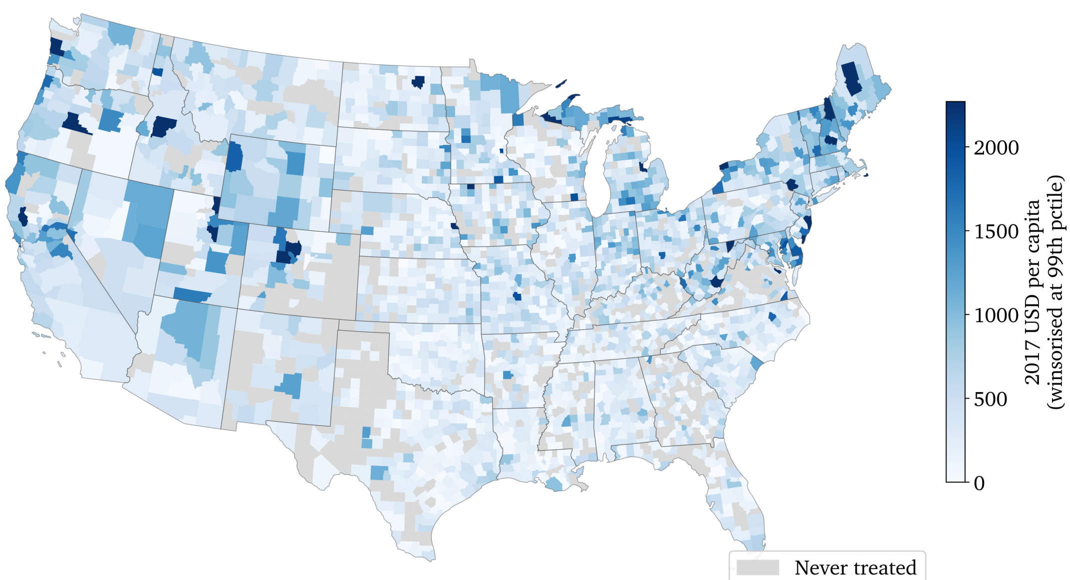

Figure A.2: Geographical distribution of grants under the CWA Construction Grants Program. Panel (a) plots the first year of grant-assisted construction start, by county, based on Step 3 and Step 2+3 construction-start dates from EPA administrative records. Counties never receiving a construction-stage grant are shown in gray. Panel (b) plots total real federal award per capita (2017 USD), by county. Color scale winsorized at the 99th percentile ($2,269 per capita).

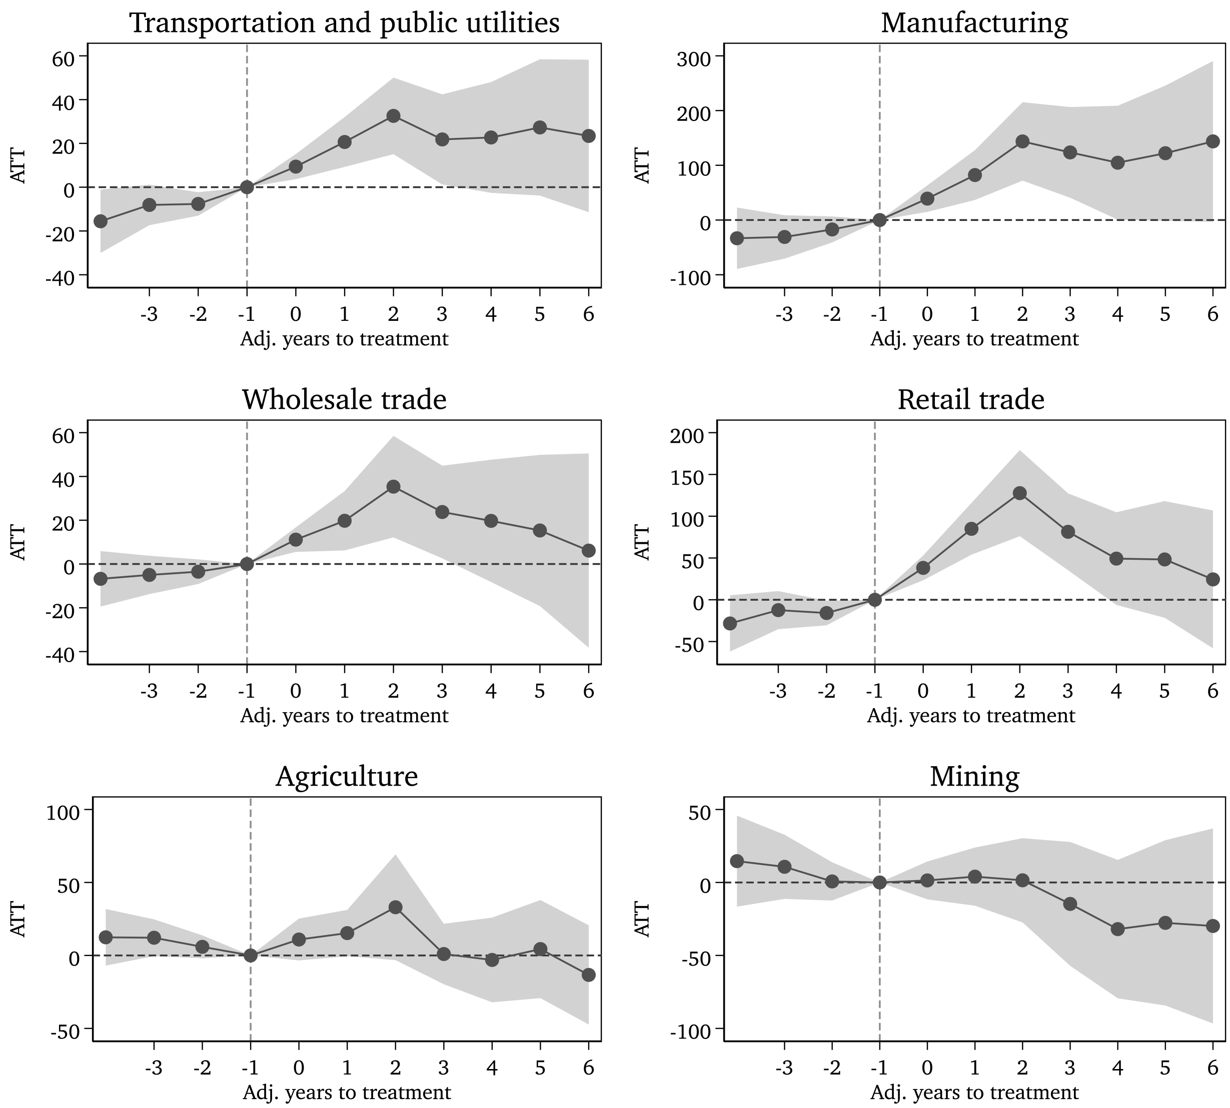

Figure A.3: Event-study CSDiD estimates of sectoral employment effects using adjusted treatment timing. Treatment timing is shifted two years earlier than construction start to capture preconstruction economic activity. Shaded areas denote 95% confidence intervals.

Table A.1: Summary statistics by Construction Grants Program treatment status

	Never treated	Ever treated
Panel A. BEA Personal Income (1969–1971; Real 2017 USD)		
Avg. population	27,329.23	76,108.59
	(145,793.71)	(247,149.94)
Population growth (%)	1.72	2.58
	(5.95)	(4.56)
Avg. personal income (thousands)	516,449.89	1,593,235.68
	(3,273,038.41)	(6,112,136.49)
Personal income growth (%)	8.31	7.78
	(14.08)	(7.49)
Avg. personal income per capita	15,518.94	16,784.04
	(4,086.11)	(3,785.10)
Personal income per capita growth (%)	6.85	5.17
	(17.74)	(7.31)
Avg. construction-sector personal income	20,212.69	85,978.31
	(95,065.32)	(325,843.73)
Construction-sector personal income growth (%)	21.17	16.93
	(107.02)	(89.13)
Avg. construction personal income / population	0.68	0.82
	(0.61)	(0.65)
Construction personal income per capita growth (%)	19.14	13.93
	(106.55)	(84.31)
Panel B. QCEW Employment Statistics (1975)		
Total employment	5,619.90	15,664.68
	(26,559.01)	(42,898.89)
Total establishments	444.40	1,066.04
	(2,072.88)	(2,525.62)
Avg. weekly wage, all sectors	523.13	535.02
	(182.51)	(112.61)
Construction employment	351.62	950.69
	(1,312.91)	(2,371.92)
Construction establishments	49.23	132.12
	(158.45)	(285.75)
Construction avg. weekly wage	595.07	640.50
	(223.24)	(182.59)
Observations	639	2,436

Notes: This table reports demographic and economic summary statistics for counties that were never treated (i.e., counties with no recorded federally assisted project under the Construction Grants Program) and counties that were ever treated. Panel A reports averages over 1969–1971, and the corresponding growth measures are percentage changes from 1969 to 1971. Panel B reports QCEW outcomes in 1975, the earliest year available in the QCEW. Because 1975 occurs after the start of the Construction Grants Program, these employment statistics may reflect early program effects and should be interpreted cautiously as baseline differences.

Table A.2: Local construction sector effects, by pre-treatment population

	Below median		Above median	
	2 yrs	4 yrs	2 yrs	4 yrs
Post avg.	12.985*	13.706	0.021	-31.550
	(5.907)	(8.034)	(18.126)	(28.296)
N	25,601	25,601	13,320	13,320
Pre-treatment avg.	127.25	127.25	1466.62	1,466.62

Notes: This table reports average post-treatment effects over the first 2 years and first 4 years after the start of grant-assisted construction, by pre-treatment (1969-1971) average county population. The median county has an average pre-treatment population of approximately 18,900. Standard errors are clustered at the county level and reported in parentheses. *Pre-treatment averages* are construction sector employment in 1975, the first year with consistent QCEW coverage. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B Sample Construction

The primary source for Construction Grants Program awards is the EPA's Grant Information Control System (GICS), obtained through a Freedom of Information Act request. The dataset records each grant's recipient municipality, the associated wastewater treatment facility, the grant step (planning, design, or construction), award and construction start dates, and the total federal award and eligible project cost. Keiser and Shapiro (2019) use the same data, and Appendix Section B.4 of that paper documents its accuracy.

I restrict the analysis to Step 3 (construction) and Step 2+3 (combined design and construction) grants, as these steps correspond directly to physical construction activity at the treatment facility. Treatment timing is defined by the construction start date (`wwt_start_date_dt`), since construction start more precisely captures when local labor demand would have increased. When the construction start date is missing, I use the award date instead. I further restrict the sample to grants initiated between 1973 and 1990, corresponding to the active years of the program under study.

County identifiers are recovered by merging grant records to the EPA's Clean Watershed Needs Survey (CWNS) facility registry using a unique facility code and by standardizing the county-name strings in the GICS against a county FIPS crosswalk.[19] A small number of grant records cannot be matched to a county FIPS code and are excluded. For Virginia, independent cities are consolidated with their surrounding counties following the BEA's combined-area convention, which is also used for the employment and income outcomes.

C Robustness checks

C.1 Robustness to treatment timing

In Section 5, I examine whether grant-assisted wastewater infrastructure projects generate broader county-level responses beyond the construction sector. Because sectors outside construction may

[19]The CWNS facility registry is obtained from the replication files of Flynn and Marcus (2023).

Table C.3: Aggregate effects under alternative timing adjustments

	Total employment			Per-capita personal income		
	$G-1$	$G-2$	$G-3$	$G-1$	$G-2$	$G-3$
Post avg. (2 yrs)	303.145***	316.315***	134.430*	0.088*	0.056	0.008
	(61.008)	(74.328)	(54.803)	(0.034)	(0.034)	(0.036)
Post avg. (4 yrs)	242.596**	284.093***	203.071*	0.101*	0.105*	0.064
	(75.457)	(84.881)	(84.741)	(0.044)	(0.041)	(0.042)
Post avg. (6 yrs)	169.069	223.525*	197.152	0.081	0.093	0.097
	(96.636)	(108.579)	(115.288)	(0.055)	(0.050)	(0.050)
N	40,332	38,746	36,770	63,898	63,898	63,898
Pre-treatment avg.	12,611.59	12,611.59	12,611.59	19.94	19.94	19.94

Notes: This table reports average post-treatment effects over the first 2, 4, and 6 years after the start of grant-assisted construction under alternative anticipation windows: $\widetilde{G}_i = G_i - 1$, $\widetilde{G}_i = G_i - 2$, $\widetilde{G}_i = G_i - 3$. Standard errors are clustered at the county level and reported in parentheses. *Pre-treatment averages* are outcome levels in 1975 for total employment and 1973 for per-capita personal income. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

respond during the pre-construction phase — through planning, engineering, procurement, and subcontracting relationships that precede physical construction (Leeper et al., 2010; Ramey, 2020) — I shift the treatment date two years earlier than construction start, corresponding to the median lag between grant award and construction start in the grants data. Here, I assess robustness to this timing choice by comparing estimates under $\widetilde{G}_i = G_i - t$ for $t \in \{1, 2, 3\}$, where G_i is the construction-start year used in the baseline specification.

Table C.3 reports CSDiD estimates for the first 2, 4, and 6 years after each adjusted treatment date. Relative to the baseline case of $t = 2$, the coefficients for both total employment and per-capita personal income are qualitatively similar across all three specifications. Differences in magnitude at shorter horizons are expected because shifting the reference date mechanically displaces the response along the event-time axis.

Figure C.4 makes this shift explicit. The event-study profiles for $t = 1$, $t = 2$, and $t = 3$ trace the same hump-shaped dynamic in the post-treatment period, offset from one another by approximately one year, while the pre-trends are broadly flat and similar across all three specifications. The similarity of the pre-trends is consistent with the anticipation window being short enough that a one-year shift does not materially alter the pre-period fit. The post-treatment profiles are, however, visibly displaced: under $t = 1$, the response peaks earlier relative to the normalized zero, whereas under

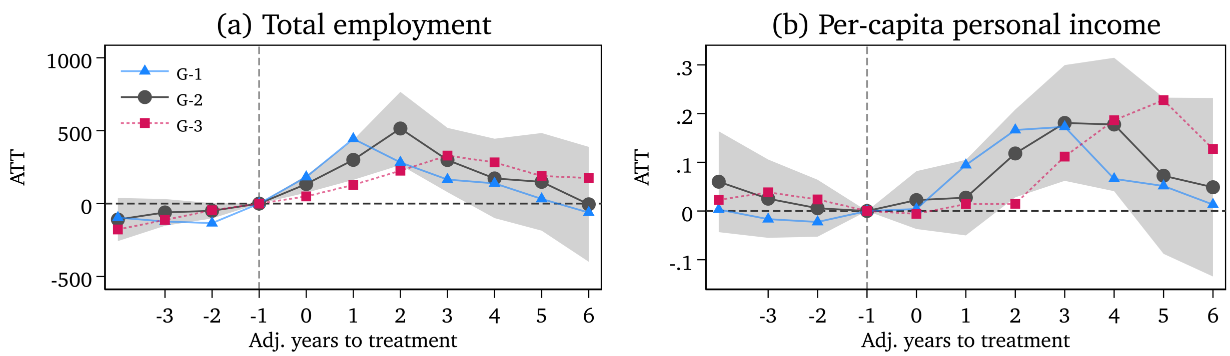

Figure C.4: Event-study CSDiD estimates of aggregate outcomes under alternative anticipation windows. Panel (a) reports total employment from the QCEW; Panel (b) reports per-capita personal income from the BEA Regional Economic Accounts. Treatment timing is shifted one year (blue), two years (gray; baseline), and three years (red) prior to construction start. Shaded areas denote 95% confidence intervals for the baseline specification.

$t = 3$, the peak is delayed, with $t = 2$ falling in between. The qualitative pattern of a gradual rise and hump-shaped response is present under all three specifications, supporting the conclusion that the baseline results are not sensitive to the precise choice of anticipation window.