

Putting America to Work, Where? Evidence on the Effectiveness of Infrastructure Construction as a Locally Targeted Employment Policy

Andrew Garin*

This Version: May 29, 2018

Abstract

Is infrastructure construction an effective way to boost employment in distressed local labor markets? I use new, geographically detailed data on highway construction funded by the American Recovery and Reinvestment Act to study the relationship between construction work and local employment growth. The method for allocating funds across space facilitates a plausible selection-on-observables strategy. I find that highway funding impacted construction employment at the county level: A dollar of additional Recovery Act spending on local construction increased local construction payrolls by thirty cents during the five years after the Act's passage. The magnitude of this effect matches the national labor share of construction revenues, suggesting that targeted spending did not crowd out other local construction. These effects are most pronounced among counties with smaller populations and smaller shares of residents that commute to outside counties for work. However, when testing for general equilibrium effects on local employment and payroll aggregates, I find effects close to zero with very wide confidence intervals across all specifications. Although the Recovery Act was a significant enough intervention to have a sizable impact on the construction sector in counties with low mobility, these findings suggest that the local variation in highway spending was too small relative to baseline regional volatility to detect a local employment "multiplier."

*Harvard University. Email: agarin@fas.harvard.edu. Support from Department of Transportation Universities Transportation Centers Grant DTRT13-G-UTC31 and a grant from Taubman Center for State and Local Government at the Harvard Kennedy School is gratefully acknowledge. I would like to thank Edward Glaeser, Nathaniel Hendren, Larry Katz, Gabe Chodorow-Reich, Daniel Shoag, Raj Chetty, Oren Ziv, Xavier Jaravel, Matthew Turner, Alex Bartik, James Lee, John Coglianese, Luca Maini, Kirill Borusyak, Ethan Kaplan, and participants at the Harvard University Public/Labor Economics and Macroeconomics Workshops for helpful comments and advice. I would also like to thank Brian Deery at AGC and Alison Black at ARTBA for helpful conversations about the road construction industry.

1 Introduction

Within countries, differences in labor market performance across places create a rationale to target labor demand enhancing policies to underperforming locales. For example, during the Great Recession in the United States, there were major differences in how different communities felt this downturn: Construction and total private employment grew by 2.3 and 20.4 jobs per thousand residents in Harris County, TX (home of Houston), while they respectively fell by 18.1 and 44.4 jobs per thousand residents in Clark County, NV (home of Las Vegas). In response to these disparities in regional performance, many policymakers believed government stimulus should be directed to hard-hit areas to facilitate recovery in those areas most in need. Even if the desire to provide greater support to underperforming communities is well-founded, the use of targeted stimulus spending to meet that end raises a more fundamental question: in an open economy, how much scope is there for policy to affect local outcomes using localized spending?

This paper studies spending on public infrastructure construction projects—in particular, “shovel-ready” projects that can commence immediately upon funding—a frequently proposed policy intended to boost local construction employment and overall economic health. Such projects can target particular geographies in a highly visible way, making them intuitively appealing as way to boost employment in distressed areas. These projects inherently create a demand for construction laborers to go to work in precise locations. As a result, one might expect approving a project in a place will increase local employment of in the construction sector, and perhaps increase employment more broadly in turn as construction workers spend their earnings locally. Despite the intuitive appeal of these arguments, however, it is important to evaluate whether the data supports them. To test these hypotheses, I study the the 2009 American Recovery and Reinvestment Act (the “Recovery Act” or ARRA for short), which authorized \$27 billion for supplemental “shovel-ready” road-construction projects that could commence promptly, with priority given to economically distressed areas. In contrast to standard federal road expenditures, the Recovery Act required detailed reporting about all stimulus road construction projects nationwide. This provides a unique opportunity to study the local employment effects of infrastructure spending.

Using the Recovery Act’s spatially-detailed data on infrastructure spending, I test whether places that received relatively more funding experienced more favorable employment outcomes than those that received relatively less using a variable treatment intensity difference-in-differences design. To the extent that construction worker might have been engaged in other construction work in the absence of the stimulus measure, counting bodies at project sites is insufficient; one must determine causal effects of spending relative to the no-spending counterfactual. To this end, I use a rich local-level dataset to consider the plausibility of a selection-on-observables methodology. Contrary to concerns that funds were systematically targeted to places with unobservably worse downturns, I find surprisingly little evidence of any targeting based on observable employment trends.

There is a clear transmission mechanism by which expenditures should affect local employment: all local employment effects should stem from the first-order “direct effect” on construction sector employ-

ment, which may in turn prompt a “local multiplier” effect on broader employment.¹ While it would be ex ante implausible to find a large employment effect with zero effect on the construction sector, whether or not there is any direct effect on local construction employment—or, given such an effect, whether there are any multiplier effects—are ultimately empirical questions. Even if many workers are engaged on stimulus projects, there may be no causal direct effect on local construction employment. This could occur due to crowd-out: Recovery Act spending on highway projects in a locality may result in the redirection of state and local funds to other localities—or even other uses—resulting in a small or even zero net increase in total local construction spending. Alternatively, if firms have mobile employees, then the direct effect could be dispersed across workers based in diffuse localities, attenuating any *localized* direct effects on construction employment. When direct effect does exist, there may in turn be a “local multiplier” effect if the added construction employment and income supports further jobs in the *same* locality (Moretti, 2010). If, however, that income is spent on goods and services produced in dispersed locations, there may “diffuse multiplier” effects—yet even if such an effect exists, it cannot be identified from cross-sectional comparisons of places.

I find that highway construction did have a direct effect on construction employment at the county level. In particular, a dollar of additional Recovery Act spending on local construction increased local construction payrolls by thirty cents during the five years after the Act’s passage, nearly exactly labor’s share of construction revenues nationwide. These labor-market effects were largest in 2010 and dissipated gradually over the following years. I find no evidence of differential pre-period trends across differently treated counties, supporting the identification assumption that all effects are casual results of additional stimulus spending. The finding that the magnitude of the direct effect is roughly what one would expect with zero crowd-out suggests that targeted Recovery Act spending during the Recession did not crowd out other local construction. However, I find that commuting matters in that local spending only impacts local employment in more isolated locales with smaller populations and smaller fractions of residents who travel to outside counties for work. When I test for effects in nearby commuting origins or destinations, I find some evidence of commuting spillovers, but the estimates are highly imprecise.

When I test for general equilibrium effects on the total employment and payroll levels within locales, I find effects close to zero, with very wide confidence intervals across all specifications. I find evidence suggesting that places with less commuting penetration have larger total employment effects, but the results are too imprecise to conclude that there is a large local multiplier anywhere. In sum, although I do not find evidence of large local multipliers, neither can I rule out large local multipliers. A local multiplier of zero could simply indicate that the multiplier effects of localized spending are highly spatially diffuse; yet, an imprecise zero estimate may also merely reflect from lack of statistical power. While the Recovery Act was a significant enough intervention to have a sizable impact on the construction sector in low-mobility counties, the local variation in highway spending may have been too small relative to baseline regional volatility to detect a local multiplier.

¹There may also be direct effects on construction firm profits, but these are unlikely to have localized multiplier effects, and, because of data limitations, cannot be directly measured at the local level.

While previous research has attempted to assess aggregate employment effects of government purchases, this paper stands apart in two respects: its direct focus on local employment effects, and its tests for distinct direct and multiplier effects of government spending. A growing strand of research, summarized by Chodorow-Reich, (forthcoming) aims to estimate the macroeconomic “multiplier” effects of fiscal expenditures using cross-regional variation in government expenditures and in economic outcomes across regional macroeconomies. Most of these studies exploit some exogenous source of variation in state-level government spending to directly estimate the effect of spending on regional aggregate employment and income—though typically without providing evidence on the transmission mechanism, as I do. Like this paper, several of the papers surveyed by Chodorow-Reich, (forthcoming) study cross-state variation in Recovery Act spending, instrumented by pre-Recession policy obligations Chodorow-Reich et al. (2012), highway spending and tax obligations (Conley and Dupor, 2013; Wilson, 2012), and seniority of legislators Feyrer and Sacerdote (2011) among other approaches, and they find the cost per job-year may land anywhere between \$25,000 and \$150,000.² Most relevant to this paper is work by Leduc and Wilson (2017) that directly evaluate the effects Federal highway spending both during and prior to the Recovery Act. Their work tests transmission mechanisms of highway spending by exploiting institutional rules governing the allocation of funds, but their primary focus is on the potential crowding-out of state government expenditures by federal expenditures—in fact, they find evidence of crowd-*in* of state spending. In their Appendix, however, they do test for direct effects on highway construction-sector employment at the state-level and find that \$1 million creates only two jobs (in 2010 only).

Recently, several papers have studied the employment effects of government spending, including the Recovery Act, at finer levels of geographic detail. Suarez Serrato and Wingender (2014) studies county-level spending multipliers using variation induced by Census revisions that result in unexpected adjustments in Federal spending levels; they estimate a cost per job of \$30,000 per year in the affected county, but no effect in adjacent counties. More closely related to this paper is recent work by Dube et al. (2014) and Dupor and McRory (2017) that draw from the same underlying data on Recovery Act awards used in my analysis to study county-level macroeconomic multipliers. These papers examine all types of awards by location of recipient—not just highway construction grants—and study effects on total employment both in recipient counties and nearby counties.³ These studies find an own-county effect of 5-10 job-years per \$1 million dollars, plus an additional 20-30 job-years per \$1 million in nearby counties (measured either by commuting flows or geographic proximity) corresponding to a total cost per job-year of \$30,000⁴. Studying the local

² Nakamura and Steinsson (2014) demonstrate that the relationship between regional multipliers and aggregate multipliers is ambiguous in a model with immobile agents, price rigidities, and non-tradable goods. In their model, local spending can drive up short-run prices for local non-tradables; depending on the monetary policy regime, aggregate multipliers may either be larger *or* smaller than local multipliers, in contrast to simple intuition suggesting that aggregate multipliers must be larger than regional multipliers.

³For many types of Federal awards that support State programs, the recipient location is the headquarters office of the local agency.

⁴This paper differs from their work in important ways. In general, the site of an “award” is not informative about the location

employment effects of a national school infrastructure investment program in Germany, Buchheim and Watzinger (2017) find a comparable effect on local employment. While these papers find large employment effects, none distinguish between a direct effect of spending on recipients and an implied multiplier effect. However, other studies in the urban economics literature have attempted to directly estimate the local-employment multiplier effects of additional jobs, usually motivated by models in which individuals are partially geographically mobile. Both Moretti (2010) and van Dijk (2016) find that an exogenously-added job in a county or metropolitan area results in one additional service-sector job in the same locale, with larger effects when the exogenous job is in a high-skill sector. To be consistent with these latter studies, then, the studies of expenditure impacts would need to have very substantial “direct” effects to be plausible.

What warrants such explicit focus on local labor-market effects? First, there are substantive reasons: Place-based policies are commonly employed as strategies that benefit local residents and workers; hence, is important to know whether or not the impacts of location-based policies effectively target local individuals at all. To that end, this paper joins a growing literature evaluating the local incidence of place-based policies in open economies, notably the study of the Empowerment Zone program by Busso, Gregory, and Kline (2013); the classic study of local taxation by Feldstein and Wrobel (1998); and other research surveyed by Kline and Moretti (2014b). During the Great Recession, there was particular desire to target policy to specific regions due to the spatially heterogeneous nature of the downturn; for instance, Yagan (2016) found that where one lived and worked during the onset Great Recession had major implications for one’s longer-run income and employment prospects. Yagan’s finding reinforces earlier work by Blanchard and Katz (1992) demonstrating that labor markets are spatially segmented, giving rise to cross-regional variation in underemployment. In theory, this could create an opportunity for a spatially targeted countercyclical employment policy to improve on unemployment insurance extensions or Keynesian policies. In models with price rigidities and labor market frictions, spending and UI create larger aggregate demand externalities when labor markets are slack (Farhi and Werning, 2012; Michaillat and Saez, 2015; Kekre, 2016). This creates a rationale to target stimulus at hard-hit regions—so long as the policy actually boosts local employment. Recent work by Monte et al. (2016) notes that the local employment elasticity in response to local expenditures may vary significantly in different settings, depending crucially on commuting behavior.

This paper assesses the plausibility of using targeted infrastructure construction to boost local employment given the mobility of the agents involved. Such approaches come at a cost, however; while more disaggregated datasets offer richer variation, the relationship between regional and aggregate outcomes

of the economic activity—for example, purchases of new buses will be listed as an award taking place at the state’s Department of Transportation. A large share of stimulus spending is nominally designated as being “spent” in state government offices, even when these funds are not actual paid to workers or firms nearby. I focus on highway construction because award affected-population ZIP codes have clear economic meaning—they are the physical construction site. Likewise, the time variation in the data based on “award date” does not necessarily correspond to the timing of any induced demand increase; rather, it is distinct from the date of any purchases or performed services. Therefore, I do not rely on this information. Finally, while spending could plausibly have larger impacts in slack markets, they test this hypothesis by stratifying the sample based on *post-2009* slack, which may be an outcome of the Recovery Act itself. Because stratification on outcomes may confound causal interpretations, I have only examined interactions of highway spending with slackness in the first year of the recession (that is, the change in the private-sector employment rate from 2006–2008), yet I find no evidence of a differential effect in slack areas to report in this paper.

becomes less clear at finer levels of aggregation. In county-level analyses, one must also consider the implications of labor mobility: If local job gains are due to in-migration, then one region's gain is another's loss.⁵ Thus, while the local detail of the data employed here facilitates transparent tests of local employment effects, I do not directly estimate an aggregate multiplier. Nonetheless, I can and do speak directly to a crucial step in the transmission mechanism of government spending onto aggregate outcomes.

The rest of the paper proceeds as follows: Section 2 provides background information about the institutional features of public highway construction funded by the Recovery Act and about the overall performance of the highway-construction sector in the context of the broader construction boom, bust, and recovery. Section 3 describes the empirical methodology and the data sources. Section 4 presents baseline estimates effects of direct effects and local multiplier effects at the county level. Section 5 examines how effects might propagate through space, and how the answer may differ across settings. Section 6 concludes.

2 The Local Distribution of Recovery Act Highway Spending: Background and Data

2.1 Background

After the 2008 onset of the Great Recession in the United States, the United States Congress passed the American Recovery and Reinvestment Act (henceforth the Recovery Act or the "stimulus bill") on January 6, 2009. The bill authorized \$821 billion of emergency supplemental expenditure by the United States federal government as an attempt to stimulate macroeconomic growth and reduce unemployment; by 2011, nearly \$500 billion had been spent directly and \$184 billion had been expended through tax reductions (CBO, 2012). Proponents of the law argued that funds could boost employment expansion in the construction sector by funding "shovel-ready" infrastructure projects, plans that could be implemented immediately if public funds were made available.⁶ The Great Recession and the associated housing bust had particularly impacted the construction sector; between 2006 and 2009, the annual rate of expenditure in residential construction fell by over \$300 billion, representing a fifty percent decline. Over the same period construction annual employment fell by over a million jobs on, contributing significantly to five million annual jobs lost in total over the same period. Thus, policymakers hoped that a rapid expansion of government infrastructure spending might stem these losses.

In practice, only a small fraction of the Recovery Act authorizations (roughly \$100 billion in total) were designated for public infrastructure construction. Of this portion, \$43 billion was apportioned for some type of transportation spending and \$27 billion was apportioned for highway construction. While this \$27 billion was small relative both to the overall stimulus both and to the aggregate decline in annual construction revenues, it representing a substantial increment over typical pre-Recession annual highway

⁵This point is made formally in Kline and Moretti (2014a)

⁶The term "shovel-ready" was first used to describe the proposed stimulus measure by President-Elect Barack Obama during a televised interview on December 8, 2008.

appropriations, which averaged around \$40 billion. It was also large relative to the overall size of the highway construction sector, the total revenues of which were \$102 billion in 2007.⁷

As with standard Federal highway expenditures, the federal government did not directly execute or even select the projects funded by Recovery Act. The distribution of Recovery Act highway funds across states was determined by preexisting apportionment formulas⁸. While it was up to States to decide which projects to fund with their allotment, the stimulus bill nonetheless provided several directions as to how projects should be selected. Specifically, the bill stated that “priority shall be given to projects that are projected for completion within a 3-year time frame, and are located in economically distressed areas”(US GPO, 2009). The first directive established an enforceable “shovel-readiness” requirement—all funds had to be obligated (that is, to have had construction contracts signed) within a year or else they would be retracted. This generally ruled out the use of funds for projects that required extensive new planning or design. In practice, stimulus funds were spent fairly quickly: By May of 2011, about 60 percent of Recovery Act transportation infrastructure funds had been spent, and 95 percent had been obligated for specific projects (GAO, 2011). As a result, the vast majority of Recovery Act highway spending went to pavement improvement, road-widening, and resurfacing projects that could be completed in a short timeframe. Appendix Figure 1 shows that over two-thirds of spending fell into this category. The second directive was to prioritize spending in “economically distressed areas.” However, the law provided no criteria by which one might determine whether economically distressed areas had been sufficiently prioritized—no quotas were set, and there was no prohibition of projects in non-distressed areas. More importantly, states had relatively free reign in designating cities, towns, and counties as “economically distressed,” and thus they applied the term liberally.⁹ Although states had to report all selected projects to the Federal Highway Administration (FHWA) for approval, no projects were rejected in practice.

While publicly funded roadbuilding may evoke memories of Works Progress Administration (WPA) workers building roads in the 1930s, procurement for Recovery Act projects followed a strikingly different model. The WPA was an explicit workforce program that directly hired unskilled workers during the Great Depression. While states and local governments were responsible for selecting and financing construction projects, the WPA provided them with free labor from workers on its payroll. WPA projects typically centered around labor-intensive work like clearing simple dirt roads and covering them with gravel using shovels. As a result compensation of workers accounted for a full 69 percent of spending on WPA projects (US GPO, 1946).

By contrast, modern public road construction is very rarely conducted by public employees. Rather, after states authorities choose project sites, they are required to select vendors through standard compet-

⁷See: <https://www.fhwa.dot.gov/policy/olsp/financingfederalaid/approp.cfm>.

⁸See Leduc and Wilson (2017) for detailed discussion of these formulas

⁹State agencies were generous in designating regions as “distressed,” and therefore particularly worthy of funds. For example, on a website providing states with implementation guidance, the FHWA posted a map of West Virginia illustrating the designation of “economically distressed areas” (EDAs) (see Appendix Figure 2) in the state. Likewise, the Commonwealth of Massachusetts lists towns designated as EDAs on its website—a list which comprises a majority of towns in the state, including Boston and Cambridge. In this light, the language urging prioritization of projects in EDAs has little bite.

itive bidding processes. Likewise, during implementation of the Recovery Act, firms for these projects were chosen strictly on competitive cost-bid bases, not by discretion.¹⁰ The Recovery Act did not give federal, state, or local authorities any special ability to interfere with the procurement process. Moreover, modern private road construction firms are relatively high-tech and rely heavily on expensive, specialized machinery. In 2012, compensation of employees represented only 28 percent of the cost of construction¹¹. Highway construction is now a relatively skill-intensive, high-wage sector; the average annual earnings for all employees in 2012 was \$56,276. At that salary, the mechanical effect (in the absence of crowd-out) of an additional million dollars of road work on construction employment should be 5.5 job-years.¹²

Appendix Figures A.1 and A.2 display how total highway construction spending and employment evolved in the United States in the context of the broader construction boom and bust. Appendix Figure A.1 shows that highway spending rose dramatically during the housing boom and that it actually continued to grow through 2008—*before* the Recovery Act was enacted—despite the dramatic decline in residential construction. However, the years following Recovery Act did not see any net increase in the total level of highway construction expenditure. In the 2007 Census of Construction, 64 percent of all highway construction was on infrastructure owned by governments (federal, state, or local); by 2012, that number had risen to 78 percent. Appendix Figure A.2 plots the evolution of employment in the same three categories during the boom, bust, and recovery. The boom-and-bust cycle is particularly apparent in the construction sector, where employment grew by 15 percent (about one million jobs) from 2000 through 2006 and then declined by double that amount between 2006 and 2011, at which point the sector began to recover. Highway construction employment also grew during the boom period, albeit less rapidly. Interestingly, although highway *spending* did not decline in 2009 or 2010, highway *employment* declined significantly on net—despite the Recovery Act—shedding 80,000 jobs between 2007 and 2013. It should be observed that the decline in highway employment was both less rapid and more persistent than the decline in the broader construction sector, potentially due to cushioning effects of the Recovery Act.

2.2 Data Sources

To study the local distribution of stimulus funds, I use project-level data obtained from the website of

¹⁰Title 23 U.S.C. 112 requires that all Federal-aid highway funds be awarded to firms based on competitive bidding, and the Recovery Act maintained this requirement. Politicians and bureaucrats may not legally interfere with procurement processes to favor certain firms, and there is no particular reason to suspect systematic corruption at this point in history.

¹¹This figure is calculated as the total cost of payroll for all employees (including white collar workers) plus the value of fringe benefits, divided by the total value of receipts *net* of subcontracts let out. Importantly, netting for sub-contracts amounts to considering the labor share inclusive of all subcontracts, under the assumption that subcontractors are also in the same sub-sector. Figures are taken from the 2012 Census of Construction.

¹²An effect of 5.5 job-years per \$1 million spent implies the \$27 billion spent during ARRA would have supported 150,000 construction job-years. If 5.5 is the correct figure, one should note that even in this no-crowd-out case, each additional construction would need to in turn to support at least five additional jobs to yield an aggregate regional multiplier as large as the ones estimated in Suarez-Serrato and Wingender(2014).

the Recovery Act Accountability and Transparency Board.¹³ Although the federal government does not generally collect or report data about the projects that state and local governments finance with federal funds, the Recovery Act explicitly required public disclosure of all awards made under the law to ensure ongoing transparency. Awardees were required to disclose the ZIP codes where and the purposes for which the funds were used, so local representatives could account for the amounts and uses of funds spent in each district and jurisdiction. From this repository, I have assembled a comprehensive dataset of each highway construction project funded by the Recovery Act, including information on the amount spent, the precise geographic location of the project, the data of completion and the nature of the project.¹⁴ Although I only observe the date on which the funds were obligated and not the dates of construction, I do observe whether work was complete by 2013 when the final version of the data was released. To ensure that my spending measures do not include large amounts of unused funds, I restrict the sample to projects that were completed by 2013.

In addition, local governments were requested, but not required, to report information about the vendors contracted to work on each project. I am therefore able to determine which firms won general-contractor bids in most instances.¹⁵ Even when projects involved multiple prime vendors, the data reports how much was paid to each vendor, as well as the name, industry, and address of the responsible office at the vendor firm.¹⁶ This unique dataset allows not only for analysis of project impacts at fine geographic levels, it also permits one to observe where contracting establishments are located relative to the projects.¹⁷ To my knowledge, this is the first such dataset with national coverage.

Table I describes the data on stimulus-funded road construction projects and county-level aggregates. Comprehensive data on projects are available for 46 of the continental states, accounting for \$23 billion in expenditures.¹⁸ Because vendor reporting was voluntary, only \$15.4 billion of these funds can be attributed to contracts with vendors. I drop all projects with no information on vendors from my primary analysis, because the quality of the location data in these observations is questionable—such cases are frequently placed in state capitol complexes. This limits my primary sample to \$16.8 billion of spending on projects with at least some vendor information listed. However, for reference, I examine my results with the inclusion of these projects in the Appendix. In order to study the impacts of these projects on local employment,

¹³While the website, www.recovery.gov, is no longer active, the full award-level dataset is available from the author upon request.

¹⁴A “project” is a grant sub-award. In the case of highway construction, a sub-award is functionally equivalent to a “project”—that is, improvements to a specified piece of infrastructure at a specified site. The data lists two ZIP codes for each sub-award, one corresponding to the location of the “recipient” authority overseeing the project and one corresponding to the “population” affected by the project. While for many types of purchases these ZIPs are the same and correspond to the local agency overseeing the use of the grant, in the case of highway construction, the “population” ZIP is intended to reflect the location of the construction work. I therefore use the “population” ZIP to designate project locations.

¹⁵The locations correspond to the regional office of the firm in charge of supervising the project. According to conversations with industry professionals, these offices are the same locations where the construction workers involved in individual projects would have been reported as employed.

¹⁶The data list recipient ZIPs for each vendor that are distinct from the award recipient and correspond to the location of the office of the vendor supervising construction.

¹⁷An important limitation is the lack of information on subcontractors, who may perform a significant part of the work.

¹⁸Michigan is excluded from these data, as most projects were erroneously reported as located in the State Capitol complex, and Illinois lacks any data on vendors.

I aggregate spending measures at the county level, which is the finest geographic unit available with detailed industry-level data. However, because individuals frequently commute across county lines, I also examine effects using alternative definitions of local labor markets, including both Metropolitan Statistical Areas (MSAs) and Commuting Zones (CZs), which are groups of counties that represent distinct and self-contained labor markets based on commuting patterns (Tolbert and Sizer, 1996).¹⁹

The data report addresses the primary vendors for each project, which allows for direct inspection of where these firms' offices were reported to have been located. In the analysis sample fully 78 percent of vendors report offices based in different counties than where their project sites are.²⁰ The county-level correlation coefficient of total per-capita spending on Recovery Act projects and total per-capita vendor receipts for Recovery Act projects is only 0.039. Even larger CZs may not be sufficiently self-contained; in fact, in my data, 55 percent of vendors have offices in different CZ than the project sites. Even at this broader level, the coefficient of correlation between per-capita spending on projects in a county and per capita payments to firms in the same county is only 0.079. The mean distance between project sites and vendor offices in the analysis sample is 79 miles, roughly the same distance as Philadelphia is from New York City. Though such distances seem large, they are actually consistent with standard bidding behavior for large construction projects, for which vendors offices routinely bid at considerable distance.²¹ The state may be the lowest level of geography at which vendor offices and projects are co-located; indeed, 90 percent of vendors are in the same state as the project sites in my data. However, it is important to note that the vendor's administrative offices may not be the establishments at which the on-site construction workers are reported as employed; this consideration is even more important to the extent that prime vendors subcontract out tasks to third-party employers.

The primary outcomes I study are private-sector employment levels within a locality, broken out by industry. I obtain annual county-by-sector employment totals from the Quarterly Census of Employment and Wages (QCEW), released by the Bureau of Labor Statistics. The QCEW data is compiled from administrative establishment-level records collected by state unemployment insurance systems. The resulting dataset includes the annual average employment and salary levels, broken out by county and industry. While data is available for detailed industries (distinguished by NAICS codes) within counties, data at fine levels of

¹⁹These are similar to Metropolitan Statistical Areas, but unlike MSAs CZs cover *all* counties.

²⁰This figure does not change significantly if it is calculated on a percent-of-spending instead of a percent-of-projects basis. The reported vendor addresses in these data reflect the location of the office supervising the project—the same office that should be reporting the employees in the administrative datasets that underly this data. While one might suspect that firms are giving addresses for a national headquarters office, this does not appear to be the case in the data—vendors who appear more than once in the data often report from different ZIP codes. Conversations with industry professionals indicate that it is standard for large multi-establishment firms to have a headquarters and “regional” offices that supervise bids, work, and employment on projects in their area. These regional offices appear to be what appear in my data.

²¹It is not uncommon for construction workers to drive 90 minutes to a project site on a daily basis. An illustrative example can be found at a road construction site near Harvard Square. A worker at the site in Cambridge informed the author that his employer (the vendor on the project) was about fifty miles north in New Hampshire—and therefore in a separate Commuting Zone. Interestingly, this worker lived fifty miles to the south in the Providence Commuting Zone. In this case, the worker would appear as employed in New Hampshire in employer reported data, though the worker actually lived in a CZ different than *both* the employer and the project. This case illustrates the difficulty in attempting to define a “self-contained” local labor market where workers, employers, and worksites would be co-located.

disaggregation with small numbers of establishments are frequently suppressed for confidentiality reasons. Thus, my primary analysis focuses on total construction sector figures, which are nearly always available at the county level. I also study effects on finer sub-industries, though results are likely attenuated due to censoring.

Importantly, my employment concept is based on place of work, not place of residence. To the extent that individuals commute beyond counties, the level of employment at local establishments may differ substantially from the level of employment among local residents. Both concepts are meaningful measures of local economic health; however, only scant data is available on employment by place-of-residence, which are primarily collected through household surveys with too few respondents to conduct local-level analyses.²²

Additionally, I supplement these data with additional control variables and with outcome variables from different sources. County-level unemployment and labor-force participation levels and rates are available in the BLS Local Area Unemployment Series; however, these data are largely imputed from sparse Current Population Survey micro-data and should be interpreted with caution. Additional demographic variables are taken from the 2000 Decennial Census Long Form survey and the combined 2005–2009 American Community Surveys (Minnesota Population Center, 2011). Data on new permits issued for the construction of housing units are obtained from the Department of Housing and Urban Development. County-level income data are obtained from IRS Statistics of Income Data. Data on road mileage and quality are assembled from road-segment data in the Highway Performance Management System compiled by the U.S. Department of Transportation. Throughout, I scale variables by 2010 population from the full-county Decennial Census, which I refer to as “per capita” units.

3 Methodology

Research Design

To evaluate the impact of government construction spending on local labor markets, I implement a difference-in-differences design allowing for heterogeneity in treatment intensity. The basic principle of this methodology is to test how employment outcomes evolved differently in places with higher versus lower levels of per-capita stimulus construction spending. This approach is valid provided unobserved determinants of employment evolutions are uncorrelated with spending levels—that is, so long places with different spending levels would have evolved identically in the absence of the Recovery Act.

I begin with a partial equilibrium approach, testing first for effects in the construction sector. Any plausible effect on the broader labor market should stem from the project’s direct effect on construction work, so it is important to verify that local spending does indeed have an effect on local construction payrolls. As noted above, if the construction sector were large, competitive, and Cobb-Douglas with labor

²²The Local Area Unemployment Series is an exception, where household survey data are used to estimate local unemployment and labor force participation levels by county of residence. My use of these data is discussed below.

share of 0.28 (the observed labor share in the highway construction sector), then, with no crowd-out, a marginal dollar spent on construction should boost local construction payroll by 28 cents. Likewise, a “marginal” million dollars in construction should increase employment by 5.5 jobs. In this no crowd-out case, letting Y_{ct} denote the county-level value of road construction at time t and $W_{ct} \equiv w_{ct} \times L_{ct}$ denote the county-level construction wage bill (itself the product of the local wage and employment level), a linear regression of the construction wage bill (in dollars) on the road construction spending level (in dollars) should yield a coefficient of .28. Thus, the first task is to verify whether such an effect is obtained in practice. I then test for local general equilibrium effects on aggregate employment and income.

To motivate the difference-in-differences specification for the partial-equilibrium analysis, consider a simple two-period scenario in which total non-stimulus road construction spending G_{ct} in county c is supplemented with a additional stimulus amount $S_{c,post}$ in the second (or “post”) period. While the total spending in county c in the post period is $S_{c,post} + G_{c,post}$, employers in outside jurisdictions may bid to work on these projects, and firms in c may be eligible to bid on outside projects. Suppose that $v_{c,t}$ represents other latent drivers of local construction payroll, such as unobserved productivity attributes. Then, letting ρ^{out} be the average share of local construction work done by employees of firms in other jurisdictions, and ρ^{in} be the share of all spending outside of c done by construction employees at firms based in c , we can obtain a regression equation as follows:

$$\Delta W_c = (1 - \rho^{out})\alpha(S_{c,post} + \Delta G_c) + \rho^{in}(S_{-c,post} + \Delta G_{-c}) + \Delta v_c \quad (1)$$

Or, in terms of observable variables:

$$\Delta W_c = \beta S_{c,post} + \Delta \epsilon_c \quad (2)$$

where the Δ operator indicates pre-post differences. Here, $\beta \equiv (1 - \rho^{out})\alpha$ is the effect on local construction payroll accounting for outside competition. The error term $\Delta \epsilon_{c,t} \equiv (1 - \rho^{out})\alpha(\Delta G_{c,t}) + \rho^{in}(S_{-c,post} + \Delta G_{-c,t}) + \Delta v_{c,t}$ reflects unobserved determinants of local construction payroll. If supplemental stimulus spending was randomly assigned across space, independently of baseline spending and other components of the error term so that $E[S_{c,post}\Delta \epsilon_{c,post}] = 0$, then this β would be identified by the difference-in-differences regression in (2).

Empirical Model

I estimate the following dynamic version of the simple equation in 2, normalizing the spending and outcome variables by the 2008 regional population level:

$$\frac{Outcome_{c,t}}{Pop_c^{2008}} = \delta_c + \gamma_{state,t} + \sum_{\tau \neq 2008} \beta^\tau \frac{Spend_c}{Pop_c^{2008}} \times \mathbf{1}\{t = \tau\} + \sum_{\tau \neq 2008} \eta^\tau X_c^{pre} \times \mathbf{1}\{t = \tau\} + \Delta \epsilon_c \quad (3)$$

With the inclusion of the county fixed effect δ_c , each year-specific β can be interpreted as the effect of \$1 of additional stimulus spending on the level of the outcome in year τ , relative to 2008. While the normalization by 2008 population allows for a better fit to the data, the interpretation of β is unchanged from equation 2: if $Outcome_{c,t}$ is average annual employment, then β is the number of job-years added in county c during year t as per Recovery Act dollar spent on a project in that county.

I restrict my focus to within-state variation in some specifications by including a state-by-year fixed effect $\gamma_{state,t}$. When including covariates, I control for differences in trends that vary with the included observable variables (reflected in the interaction of the pre-period control with an outcome-year dummy). Because per-capita spending measures are mechanically correlated with the baseline population level, represented in the denominator, I include a control for the log of the 2008 region population in all specifications. Standard errors are clustered at the county level. While some information is available about the timing of spending authorizations, the data does not provide detailed information about when work took place. Accordingly, I suppress all time variation in the treatment variable and instead study the dynamic effects of the total Recovery Act highway construction spending level, which is defined as a time invariant level.

When examining alternative outcomes and specifications, I focus on a simple two-period differences-in-differences equation using 2008 as the pre-period year and 2010 as the post-period year. Since 2010 is both the year in which the highest level of Recovery Act construction was reported to have occurred and the year in which I find the largest construction-sector effects, it is natural to examine the effects in 2010 in more detail. I estimate the following equation:

$$\frac{Outcome_{c,t}}{Pop_c^{2008}} = \delta_c + \gamma_{state,t} + \beta^{post} \frac{Spend_c}{Pop_c^{2008}} \times \mathbf{1}\{t = 2010\} + \sum_{\tau \neq 2008} \eta^\tau X_c^{pre} \times \mathbf{1}\{t = 2010\} + \Delta\epsilon_c \quad (4)$$

The baseline specification includes state-by-year fixed effects, time-invariant county fixed effects, and a control for year-specific effects of 2008 population. In practice, estimating β^{post} from the simple difference-in-differences specification in 4 will yield nearly identical point estimates to β^{2010} estimated from the dynamic specification in 3.

Threats to Identification

Equation (1) points to several limitations and concerns that arise when estimating (2) in actual data, where $S_{c,post}$ is not fully randomized. First, even if $S_{c,post}$ were randomized the local employment effect β will be less than the full partial equilibrium employment effect α due to a “firm-commuting effect”: If some contractors working on projects in c are based in other locales, and if $S_{c,post}$ is uncorrelated with the amount of stimulus work in other locales that firms based in c are able to win, then the local employment effect will only be a fraction $(1 - \rho^{out})$ of the total partial equilibrium effect.²³ This firm-commuting effect

²³Separately, any partial equilibrium effect will also omit general equilibrium effects, such as macroeconomic multipliers, not captured in (1).

is not an econometric bias but rather an economic effect: As the spatial mobility of firms increases, the ability to target demand to firms in c using local construction in c will diminish. However, if the level of geographic aggregation is too fine relative to meaningful concepts of “local” markets, one may obtain an arbitrarily small employment effect. Thus, I consider robustness to different definitions of local labor markets in the analysis below.

A second concern is that the level of local stimulus spending $S_{c,post}$ may not be orthogonal to other road spending $G_{c,t}$ if the Recovery crowds out other public *or* private infrastructure spending. One well-documented form of crowd-out is “flypaper effects,” as in Knight (2002) and Leduc and Wilson (2017): Because budgetary funds are fungible, state and local governments may use federal construction grants to offset their own planned spending and in turn use the additional budgetary resources for different purposes. At the state level, Leduc and Wilson (2017) find that federal highway awards actually lead to a subsequent *increase* in state-level spending on highway construction. Nonetheless, crowd-out may exist at the local level within states if states systematically direct their own funds away from regions with stimulus-funded projects towards projects in other locales. Unfortunately, there is no nationwide data source on non-stimulus federal, state, or local government road construction spending that is disaggregated below the state level; thus, it is not possible to directly test for local flypaper effects. Yet, if such behaviors were first-order, then one should expect crowd-out to be more severe and effects to be smaller when restricting the analysis to within-state variation across counties (by including state-by-year fixed effects) compared to analyses that allow for comparisons of counties across states. I show below that this is not the case.

Another major concern is that authorities might have systematically targeted funds towards regions experiencing different labor market or industry trends in a manner that violates the parallel trends assumption (i.e., $E[S_{c,post} \times \Delta v_c] \neq 0$). This assumption warrants particular scrutiny in the context of counter-cyclical spending. In particular, if funds are directed towards regions experiencing larger adverse shocks that are transient and mean-reverting, then the more-treated locales would have grown faster in the absence of any intervention. In the case of the Great Recession, one might separately worry that harder-hit counties experienced persistent adverse trend-breaks in 2009 that continued throughout the observation window.

I address this identification concern in several ways. First, I implement a dynamic event study specification that facilitates a direct test of whether differently-treated locales experienced different trends prior to 2009. If the parallel trends assumption is valid, then Recovery Act spending should be uncorrelated with pre-2009 trends in the outcome variables. However the parallel trends assumption may still be violated, even in the absence of differential pre-period trends, if state governments targeted funds on the basis of changes in Δv_c that occurred after 2008. While it is impossible to fully rule out any sorting of this sort, I probe the validity of the no-selection assumption by testing for robustness to varied control sets. If estimates are robust as the choice of X_i varies, any unobservable source of bias would have to be both systematically correlated with the error term ϵ_{ist} and systematically uncorrelated with all of the observable dimensions of heterogeneity. This insight guides my inference below.

Targeting of Stimulus Spending and Covariate Balance

The Recovery Act encouraged state and local governments to create jobs in distressed areas, but where did the funds flow in practice? Panel A of Figure 1 plots per-capita spending against the 2007-2009 change in first-quarter unemployment, netting out state averages and adjusting for population size. Interestingly, spending was no higher in regions with higher unemployment levels. While these figures from the LAUS are measured with considerable noise, they nonetheless represent the best information policymakers could have used on which to base any intentional targeting of spending towards slack markets. Another indicator of local labor market distress prior to the Recovery Act is the decline in the construction employment rate from 2006–2008; however, there is once again no clear relationship between a larger decline (i.e., more negative growth) and local spending.

Figure 2 plots the pairwise standardized correlation coefficients between Recovery Act construction spending and a broader range of covariates, residualized on state-level covariates in order to isolate within-state correlations. There is a clear relationship between spending per-capita and 2008 log population; while this relationship is mechanical in part, it also reflects the substantive fact that rural regions with smaller populations have disproportionately more roads relative to their populations. As I control for 2008 log population in all specifications, all other correlations in Figure 2 are adjusted accordingly. Conditional on population size, there is little relationship between spending and observable indicators of local labor market distress. Rather, the clearest determinant of the spatial allocation of spending appears to be how many frequently-used roads were in each county. Looking within states, primary road lane-miles per capita, average daily vehicle-miles traveled (VMT) per capita, and land area are the strongest correlates of per-capita Recovery Act road construction spending.²⁴ This is unsurprising, given that the majority of Recovery Act road-construction funds went to pavement improvement and road-widening projects—work that naturally occurs in places where more roads already exist.

The stock of roads is by no means randomly assigned; however, it is hard to imagine what latent drivers of the boom, bust, and recovery cycle would be correlated with the baseline stock of roads yet *not* observable indicators of economic distress. Leduc and Wilson (2017) argue that measures of local road mileage and usage (VMT and lane-miles, in particular) are completely orthogonal to latent drivers of short-run growth, and the authors employ these variables to construct instruments for highway spending based on formulary apportionment rule. I do not take a strong stance on this assumption here, but I will show that the results below are robust to conditioning on these variables.

4 County-Level Spending Effect

I first examine estimated effects of project spending on local construction employment. In order to de-

²⁴“Major roads” includes all roads designated by the Federal Highway Administration as “principal arterial” or that are part of the Interstate Highway or National Highway Systems in the Highway Performance Management System data. Lane-miles are a distinct concept from road-miles; one mile of a road with two lanes running in each direction would constitute four lane-miles of road.

termine when any effects occur and whether pre-period trends differ with the treatment variable, I estimate the dynamic specification in equation 3; results are displayed in Table 2. To facilitate inspection of these results, Figure 3 plots the year-specific effects of local per-2008-capita project spending (in millions of dollars) on per-2008-capita construction-sector employment in the recipient county, presented in Column 1 of Table 2. This baseline specification includes state-by-year fixed effects and controls for differential trends across 2008 population levels; the coefficients can be interpreted as the number of construction jobs added per each incremental million dollars spent. Consistent with the parallel trends assumption, construction employment trends evolve similarly on average across all spending levels. While there is little to no effect in 2009—consistent with an implementation lag—there is a significant effect of two additional jobs per million dollars in 2010, relative to the counterfactual.²⁵ This within-state, cross-county effect is nearly identical to the cross-state employment effect on construction employment found in Leduc and Wilson (2017). In 2011, there is a borderline-significant effect of 1.58 jobs per million, which continues to dissipate in subsequent years. The magnitude and timing of this effect are highly plausible: Strikingly, the combined 2009-2013 point estimates suggest a combined *local* effect of 5.9 job-years per million dollars, close to the back-of-the-envelope closed economy calculation above had predicted.

Table 2 also reports year-by-year point estimates from the dynamic specification in 3 for other partial-equilibrium outcomes. Column 2 tests for effects on “heavy and civil construction” sub-industry employment (NAICS 237), in which nearly all general contractors for highway projects would be classified. The time pattern of results is similar to Column 1, though point estimates are less than one-half of the size of Column 1’s estimates. This attenuation is primarily due to censoring in the QCEW; tests for effects in other sub-sectors yield point estimates close to zero in the post-period.²⁶ If construction labor supply is imperfectly elastic, increased construction demand may be reflected in higher wages, in addition to higher employment levels. The results in Column 3 indicate that the average annual construction salary did, in fact, rise in each of 2010 and 2011 by three dollars per construction worker in response to one additional dollar per resident. Column 4 reports the combined effect on construction payrolls, which reflects both wage and employment increases. The point estimates imply that for each additional dollar of spending, local construction payroll rose by 10 additional cents from 2008 to 2010. Summing the annual effects over each of the post-period years 2009-2013, the estimates in Column 4 imply that thirty percent of road expenditures ultimately wound up in the pockets of construction workers at local establishments.

To probe the robustness of this result, Table 3 presents the 2010 point estimate under a range of alternative specifications. Columns 1, 2, and 3 report results with no region-year fixed effects, state-year fixed effects, and commuting zone-year fixed effects. The magnitude of the effect is invariant to the inclusion

²⁵This suggests that \$27 billion in total spending could have supported over 50,000 jobs in 2010. Direct surveys of contractors involved in Recovery Act construction found that number of full-time-equivalent employees at work on Recovery Act highway projects peaked at 40,000 in September 2010 (CBO, 2012).

²⁶One may be concerned that the construction effect reflects a boost in residential or nonresidential building construction that is not plausibly due to Recovery Act road spending. However, in Appendix Table A.3 I find this is not the case, and that the combined 2009-2013 effect on employment in these sub-sectors is approximately zero. Nor is the effect driven by the remaining construction subsectors, where I find a combined 2009-2013 effect of 0.197 (1.79%). These results are all consistent with the censoring explanation.

of state fixed effects. This suggests that within-state crowd-out factors do not have first-order effects on the results. More surprisingly, the effect is only slightly reduced by the inclusion of commuting zone fixed effects. This latter finding implies that the local employment effects identified in the county-level analysis are quite local indeed, and it implies that demand spillovers to neighboring counties are minimal.

In Columns 4 and 5, I test for confounding selection on observables by including state-year fixed effects, and I add a wide array of covariates including county size and density; the 2007 and 2009 Q1 unemployment rates; the number of housing permits issued in 2003, 2006, and 2008; the housing density change between 2000 and 2010; and the 2008 size and 2006–2008 growth of the local construction sector. The point estimates are largely robust to these inclusions, consistent with a causal effect. Additionally, Column 6 tests whether the effects are driven by the extensive margin (the comparison between places with positive spending and those with none) or the intensive margin (the comparison among places with different levels of positive spending) by restricting the baseline specification of Column 2 to the subsample of places that received positive spending. The point estimates are essentially the same, suggesting that the main effects occur primarily on the intensive margin.

Having established a clear partial equilibrium effect, I next test whether the gains—or averted losses—in the construction sector are reflected in countywide industry aggregates. The results are ambiguous. Table 4 summarizes the results, and Figure 4 plots the year-specific effects of local per-2008-capita project spending (in millions of dollars) on *total* per-2008-capita private sector employment in the recipient county analogous to Figure 3. In both the short and long runs, there is no clear effect on aggregate employment. Yet, although the combined 2009–2013 effect is less than two job-years per million dollars, the confidence interval implied by the standard error is large, and is consistent with employment effects of anywhere between a loss of twenty-five jobs to a gain of over thirty jobs per million dollars. Table 4 presents results from tests under a number of alternative specifications. The 2010 point estimate is not stable to specification, nor is the estimate economically large in any instance.

The simplest explanation for this small, imprecise finding is a lack of statistical power. Stimulus highway spending may have had a first-order impact on the construction sector, yet the intervention was likely too small relative to the fluctuations in total employment during the Great Recession to detect any plausibly-sized impact on local employment. Although the small effect is consistent with crowding-out of non-construction spending by stimulus projects, large effects on the order of those in Suarez Serrato and Wingender (2014) are also not ruled out by the obtained confidence intervals.

5 Regional Spillovers and Spatial Heterogeneity

This section examines how effects might propagate through space and how those effects may differ across settings. To the extent that firms and workers are mobile across space, there may be regional labor market impacts of highway construction projects that extend beyond county borders. As a result, one may expect to find larger effects when conducting the analysis at levels of aggregation reflecting more meaningful labor markets, as there will be less “leakage” of labor demand into other jurisdictions. (In the

context of Equation 1, this corresponds to a higher level of ρ^{out} .) I first test for effects at more aggregated geographies, as well as for localized spillovers using observed commuting flows in the spirit of Dupor and McRory (2017). I find no evidence of significant regional demand spillovers. I then test for heterogeneity in the effects and, importantly, I only find significant effects on construction employment in places that are smaller and less conducive to commuting. This is consistent with effects being detectable only in settings where an intervention was sufficiently large relative to a local economy.

Tests for Regional Effects

To study how the labor-market impacts of local spending vary at different geographic scales, Table 5 presents results at the county, CZ, MSA, and state levels. While urban agglomerations represented by MSAs offer the most standard notion of a labor market, an important part of the identifying variation comes from regions outside urban MSAs, as discussed below. Therefore, CZs offer a comparable alternative that includes all counties in the data. Since the largest effects occurred in 2010, I use a version of the baseline difference-in-differences specification in Equation 4 in order to yield a clear basis for comparison with the baseline results.²⁷

At broader levels of aggregation, the results are less precise due to the small number of observations. Although the baseline point estimates of the per-dollar effects of spending on construction employment and payroll effects are indeed larger at the CZ level, these results are less robust to the choice of regression specification. This supports the hypothesis that stimulus construction work increases demand at construction firms outside the immediate locale of the project, such that the total regional labor market effect may exceed those found above. At the MSA and state level,s these results become increasingly imprecise but remain mostly consistent with that hypothesis. The effects on total employment remain small and highly imprecise at all levels of aggregation. There is no increase in power, as any increases in potential effect size at higher levels of aggregation are offset by lost precision due to the smaller sample size.

A more direct test for geographically diffuse demand spillovers would be to examine whether high levels of spending in proximate areas have effects on local employment. To develop a economically meaningful index of exposure, I use data on 2006-2010 county-to-county commuting flows from the American Commuting Survey (ACS) to measure how much stimulus road spending occurred in the places where people living in the observation counties go to work. Although the primary outcome data is based on place of work rather than place of residence, worker flows may still provide a useful proxy for firms' ability to send their employees to work in projects in another given locale.²⁸ I construct exposure using two alternative methods: First, I measure average per-capita spending in the places to which residents of the observation county commute outside of the observation county itself. For an observation county o and each destination county d , I calculate λ_{od}^{out} , the number of residents in o who commute to county d as a share of all residents

²⁷The specifications includes a control for 2008 population, but exclude state fixed effects. At broader levels of aggregation, the arguments that justified the simple differences-in-differences design for the county-level analysis above may be less valid. Accordingly, these estimates are presented for comparison.

²⁸While the data is less reliable at fine scales, I also examine employment outcomes for local residents based in treated locales obtained from the ACS and LAUS in the Appendix

in o who commute outside of o (so that $\sum_{d \neq o} \lambda_{od}^{out} = 1 \forall o$). I then use these shares to calculate the average treatment intensity that commuters from o are exposed to, $Exp_o^{out} \equiv \sum_{d \neq o} \lambda_{od}^{out} \times \frac{Spend_d}{Pop_{2008,d}}$. Second, I measure how much spending in other counties d would go to local firms in o if the share of projects in d worked on by firms in o were the same as the number of workers in d who reside in o and commute. To do this, I denote the share of workers in d who commute from o as λ_{do}^{in} and calculate the expected outside revenues (per capita) in o as $Exp_o^{in} \equiv \frac{1}{Pop_{2008,o}} \sum_{d \neq o} \lambda_{do}^{in} \times Spend_d$. I then test whether these measures predict employment outcomes when included in regressions along with the main own-county spending measure.

Results are presented in Table 6. For both measures, there is no evidence of a significant positive construction demand spillover across space. This is consistent with the aggregated results and the finding in Section 4 that county-level results are insensitive even to the inclusion CZ-level fixed effects. In particular, any positive “partial equilibrium” effects on construction are either in the project within the locale or are so highly dispersed across space relative to standard commuting patterns that it is hard to detect under any standard definition of a broader labor market. However, theory is ambiguous as to whether cross-regional labor demand spillovers should be positive on net. An alternative explanation may be that spending in a locale can have *negative* partial-equilibrium spillovers onto construction employment in other locales if construction workers leave establishments in neighboring regions in order to work on local projects. Thus, the lack of observed spillovers may be because spillover effects are actually zero on net, rather than because they are positive but difficult to identify in the data.²⁹ As with the main own-county spending variable, effects on total employment are too imprecise to make any definitive conclusion.

Given the high incidence of vendors located outside of project counties, one can also test for effects in locales where *vendors* were located. Thus, rather than defining the county-level variable as total per-capita spending on projects within a county, I define the “vendor treatment” variable as the total per-capita receipts by vendor firms based in a county. While it is straightforward to use this variable as the treatment in the differences-in-differences, the no-selection-on-unobservables assumption is much less plausible, and one should be hesitant to interpret any results as having a causal interpretation. In particular, firms and regions with better latent productivity trends may be in a stronger position to win competitive bids, which could bias results upwards. At the same time, firms that work on stimulus projects may be systematically less likely to work on non-stimulus projects, leading to a possible downwards bias. Appendix Figure A.4 plots the annual coefficients estimated from the dynamic specification in 3, using per-capita vendor payments as the treatment and per-capita construction employment as the outcome variable. While there is no clear post-2008 effect, there is a pronounced increase in construction employment in 2008—before the Recovery Act was passed—that is associated with higher levels of receipts, indicating confounding selection. Thus, only specifications controls for pre-period evolutions in sectoral employment can plausibly be interpreted as causal; that said, even with controls, any causal interpretation is suspect.³⁰

²⁹Regardless, these results do not constitute evidence of net negative spillovers.

³⁰There are nonetheless sources of variation in vendors’ receipts of stimulus funds that would satisfy a parallel trends assumption. For example, if new road construction opportunities were scarce, so that firms that lost stimulus bids had little to do otherwise, and if the factors that determined which firms won bids were mostly arbitrary, then assignment of contracts could be plausibly exogenous.

Table 7 displays differences-in-differences estimates of the 2010 effects of both local project spending per capita and local vendor receipts per capita when both are included jointly in regressions. Interestingly, inclusion of the vendor variable has little to no impact on the project spending variable relative to Table 3; this may be unsurprising, given the lower correlation of the two variables. This implies that the county-level construction employment effect does not operate through the primary vendor, suggesting that many construction workers on-site employed at establishments other than the prime vendor, either at regional branch offices of the vendor firms, or at separate subcontracting firms.³¹ While the size of the estimated effect of local project spending is mostly robust to controls, the effect of vendor receipts is more sensitive. With the inclusion of controls, the vendor receipts effect becomes both positive and similar in size to the local spending effect. Taking the local results at face value, these effects imply that there is a 2010 effect of roughly two jobs per million dollars in *each* of the project and vendor counties, suggesting that there may in fact be an aggregate effect that is larger than the initial effects had implied. The state is the only level where vendors and projects are sufficiently co-located that the effect only loads onto the spending variable. Vendor spending is also associated with larger total employment effects conditional on controls; however, while suggestive, these results are tenuous at best, given the sensitivity to specification.

Treatment Effect Heterogeneity and Nonlinear Effects

Since the effects of spending may not be constant in all places, I next examine heterogeneity in the treatment effect. In particular, Equation 1 highlights the possibility that effects will be larger in places where local projects are more likely to be worked on by employees of local establishments. Accordingly, I test whether the effect is different in counties that are more or less open to commuting, measured by the share in the ACS who commute from out of state.³² In addition, it is possible that the effects may differ across counties with larger and smaller populations, as a given change in spending per capita could represent a bigger proportional shock to demand aggregates in smaller locales.³³ Thus, I also test for heterogeneity by 2008 population as well.

The results are presented in Table 8. Focusing on the core effect on construction employment in 2010, it appears that there is no effect in counties that are above the median population in 2008 or that have an above-median share of workers who commute outside the county. By contrast, the baseline effects appear to be driven by a more pronounced effect in smaller counties that are less open to commuting. While interactions with continuous measures of openness and population yield less precise results, the signs of the estimate are consistent with a story in which the main effects were driven by smaller, more isolated counties. These findings are consistent with claims that it is more difficult to detect the impacts of an

³¹Since data on subcontractors is not available, it is not possible to determine for certain whether this is the case

³²To test for heterogeneity in effects, I estimated interacted difference-in-differences specifications of the form $\frac{Outcome_{c,t}}{Pop_c^{2008}} =$

$$\delta_c + \gamma_{state,t} + \beta^{open,post} \frac{Spend_c}{Pop_c^{2008}} \times \mathbf{1}\{t = 2010\} \times \mathbf{1}\{HiCommute\} + \beta^{closed,post} \frac{Spend_c}{Pop_c^{2008}} \times \mathbf{1}\{t = 2010\} \times \mathbf{1}\{LoCommute\} \\ + \gamma^{open} \mathbf{1}\{t = 2010\} \times \mathbf{1}\{HiCommute\} + \sum_{\tau \neq 2008} \eta^\tau X_c^{pre} \times \mathbf{1}\{t = 2010\} + \Delta \epsilon_c$$

³³Put differently, the variation generated by Recovery Act road spending may be too small to significantly impact larger economies.

intervention due to low power in larger local economies and in economies that are more interconnected with broader regions.

Table 9 and Figure 5 examine heterogeneity across commuting openness levels for a broader range of outcomes. I find the effects on the construction sector (and on the heavy and civil construction sub-sector in particular) are much more pronounced in areas with less commuting both in the short run and in the long run. This is intuitive in the context of the framework laid out above and summarized in Equation 1: When commuting mobility is lower (corresponding to a lower level of ρ^{out}) there will be less “leakage” of labor demand into other regions and a more concentrated effect in the target area. This commuting effect—which is really about the ability of firms to send workers to other locales—is somewhat different from the effects of commuting discussed in Monte et al. (2016). That study considers commuting flows of residents in one locale (o) to firms in different locales (d), presuming the demand level of firms in locale d is known. As d becomes more open to commuting, the employment impact of a *given* local labor demand boost in d will be greater. Since the local labor supply elasticity is higher when it is more open to commuting, a given change in labor demand affects the quantity margin (employment) more and the price margin (wages) less. By contrast, in the scenario studied here, the first-order question is whether local spending has a detectable effect on local labor demand in the same region; this is the question formalized in the framework above.

One should expect effects to be easier to detect in the upper regions of the treatment distribution if the absence of an effect in bigger regions is due to of lack of power because the treatment variation is too small. To test whether effects are more pronounced as the intervention grows, I also include a quadratic term into the baseline specification in Table 8. Indeed, it appears that the main effect loads onto the quadratic term, suggestion that even when restricting the focus to construction employment, the spending intervention was only large enough to have a noticeable impact at very high levels. Meanwhile, when looking at broader employment outcomes, even the upper ranges of per-capita spending fail to identify a precise effect. Similarly, although Figure 5 provides suggestive evidence that there may be larger payroll multipliers in more self-contained regions, the results are too imprecise to rule out a multiplier of zero. Given the massive background fluctuations in regional economies and the relatively small size of the Recovery Act highway construction program, it is likely that these tests lack the statistical power to detect any plausibly-sized multiplier effect with a high degree of confidence.

6 Conclusion

In the years following 2009, many highway, tunnel, and bridge projects across the United States were accompanied by signs bearing the slogan “Putting America to Work.” This paper examined whether this slogan held true at the local level, and whether construction projects are ultimately an effective way to stimulate *local* labor markets during slack periods. Should local governments prioritize highway construction as a way to boost local employment? I used a unique dataset containing detailed information on all projects funded by the Recovery Act and their contractors to test whether places that had more construction work

had better local labor market conditions during the recovery period 2009–2013. The siting of these projects appears to have been uncorrelated with any observable indicators of economic performance, but I examine robustness of results to a wide array of selection-on-observables of assumptions, as well as an instrumental variable approach.

I find that highway impacted construction employment at the county level: a dollar of additional Recovery Act spending on local construction increased local construction payrolls by thirty cents over 2009–2013. The magnitude of this effect matches the national labor share of construction revenues, suggesting that targeted spending did not crowd out other local construction. These effects are most pronounced among counties with smaller populations and smaller fractions of residents who commute to outside counties for work. However, when I test for general equilibrium effects on local employment and payroll aggregates, I find effects close to zero with very wide confidence intervals across all specifications. Although the Recovery Act was a significant enough intervention to have a sizable impact on the construction sector in counties with low mobility, these findings suggest that the local variation in highway spending was too small relative to baseline regional volatility to detect a local employment “multiplier.”

My findings suggest the credibility of larger estimates rests on the plausibility of the local transmission mechanism: if one finds large employment or income multiplier effects, one should also observe first-order effects on the direct recipients of government funds. I argue that while it is highly plausible that that construction spending has a detectable effect on construction employment yet does not have a detectable effect on aggregate employment, it is less plausible that there could be a true effect on aggregate employment without an effect on construction payrolls. If this link is missing, this casts doubt on the validity of the source of variation used to estimate the multiplier.

Importantly, though the magnitude of the effects on the construction sector is what one might expect in the absence of crowd-out, the implied cost-per-job I estimate is high relative to other estimates in the literature. My estimate of five construction job-years per \$1 million implies that one construction job-year requires \$200,000 in highway spending to be sustained. While I find a larger impact on construction employment compared to the state-level findings than in Leduc and Wilson (2017), the cost per construction job-year is much more expensive than the roughly \$30,000 cost per *total* job year in Suarez Serrato and Wingender (2014) and Buchheim and Watzinger (2017). In order to reconcile the construction employment effects I find—which imply a cost per construction job of \$200,000—with a \$30,000 cost per total job, a construction job would have to result in the creation of *five* additional non-construction jobs. While not ruled out by my estimated effects on total employment, this is not plausible based on prior studies of local employment multipliers, which typically find that an exogenous one job-year increase in low-skilled work supports only one additional job in the same metropolitan area (Moretti, 2010; van Dijk, 2016).

Thus, the primary value of infrastructure improvement work in a community arises in the value of the infrastructure itself. To this end, the goals of policymakers should be to channel funds toward projects with high infrastructure value. It is possible that improved highway infrastructure supports robust local employment growth in the much longer run, although I do not find clear evidence of any effects as of 2013.

In a classic survey, Gramlich (1994) notes that the returns on investment for road construction work are potentially quite high, but that they vary significantly depending on the type of project selected: Spending to build new rural roads has little value, while repair of heavily-used highways can have a rate of return as high as 35 percent. Thus, while it could be easy to justify repairing aging highway arteries with stimulus funds, there appears to be little justification for widening sparsely used roads in a rural county hit hard by the housing bust in hopes of boosting local employment. One should note that concentration of Recovery Act construction in paving activity rather than in new construction is not inherently problematic in this view, as the social value of resurfacing crucial roads may exceed that of building a new “bridge to nowhere.”

While most analyses of stimulus spending have studied interventions like the Recovery Act from a macroeconomic angle, this paper has taken a more microeconomic approach in its investigation of the spatial nature of the transmission mechanism. Further analysis of microeconomic effects of stimulus spending using disaggregated data in future work and will help supplement theoretical micro-foundations for macroeconomic models with better empirical content. Even if highway spending did not have effects on specific local labor markets, it may nonetheless have proved effective in providing the means for individual firms facing constraints to stay afloat during the construction bust and in keeping individual workers attached to the labor force during the downturn. If such effects exist, better evidence would point to more cost-effective counter cyclical policy measures. The data assembled for this paper should be of use to those pursuing work along these lines.

References

- Olivier Blanchard and Lawrence F Katz. Regional Evolutions. 23(1):1–76, 1992.
- Lukas Buchheim and Martin Watzinger. The Employment Effects of Countercyclical Infrastructure Investments. November 2017.
- Matias Busso, Jesse Gregory, and Patrick Kline. Assessing the Incidence and Efficiency of a Prominent Place Based Policy. *American Economic Review*, 103(2):897–947, April 2013.
- (United States Congressional Budget Office) CBO. Estimated Impact of the American Recovery and Reinvestment Act on Employment and Economic Output from October 2011 Through December 2011. February 2012.
- Minnesota Population Center. *National Historical Geographic Information System: Version 2.0*. University of Minnesota, 2011.
- Gabriel Chodorow-Reich. Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned? *American Economic Journal: Economic Policy*.
- Gabriel Chodorow-Reich, Laura Feiveson, Zachary Liscow, and William Gui Woolston. 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, 2012.
- Timothy G. Conley and Bill Dupor. The American Recovery and Reinvestment Act: Solely a Government Jobs Program? *Journal of Monetary Economics*, 60(5):535–549, 2013.
- Arindrajit Dube, Ethan Kaplan, and Ben Zipperer. Excess Capacity and Heterogeneity in the Fiscal Multiplier: Evidence from the Obama Stimulus Package (. August 2014.
- Bill Dupor and Peter McRory. A Cup Runneth Over: Fiscal Policy Spillovers from the 2009 Recovery Act. *Economic Journal*, 2017.
- Emmanuel Farhi and IvÃ¡n Werning. Fiscal Unions. 2012.
- Martin Feldstein and Marian V Wrobel. Can State Taxes Redistribute Income? *Journal of Public Economics*, 68(3):369–396, June 1998.
- James Feyrer and Bruce Sacerdote. Did the Stimulus Stimulate? Real Time Estimates of the Effects of the American Recovery and Reinvestment Act. February 2011.
- (United States Government Accountability Office) GAO. Funding Used for Transportation Infrastructure Projects, but Some Requirements Proved Challenging. *GAO-11-600*, June 2011.

- Edward M Gramlich. Infrastructure Investment: A Review Essay. *Journal of Economic Literature*, 32(3): 1176–1196, 1994.
- Rohan Kekre. Unemployment Insurance in Macroeconomic Stabilization. 2016.
- Patrick Kline and Enrico Moretti. Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority. *Quarterly Journal of Economics*, 129(1): 275–331, 2014a.
- Patrick Kline and Enrico Moretti. People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, 6(1):629–662, 2014b.
- Brian Knight. Endogenous federal grants and crowd-out of state government spending: Theory and evidence from the federal highway aid program. *American Economic Review*, 92(1):72–91, 2002.
- Sylvain Leduc and Daniel Wilson. Are State Governments Roadblocks to Federal Stimulus? Evidence on the Flypaper Effect of Highway Grants in the 2009 Recovery Act. *American Economic Journal: Economic Policy*, 9(2):253–292, 2017.
- Pascal Michaillat and Emmanuel Saez. The Optimal Use of Government Purchases for Macroeconomic Stabilization. 2015.
- Ferdinando Monte, Stephen J. Redding, and Esteban Rossi-Hansberg. Commuting, Migration and Local Employment Elasticities. November 2016.
- Enrico Moretti. Local Multipliers. *American Economic Review: Papers & Proceedings*, (100):1–7, May 2010.
- Emi Nakamura and Jon Steinsson. Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review*, 104(3):753–792, 2014.
- Juan Carlos Suarez Serrato and Philippe Wingender. Estimating Local Fiscal Multipliers. March 2014.
- Charles M Tolbert and Molly Sizer. U.S. Commuting Zones and Labor Market Areas: A 1990 Update. September 1996.
- (United States Government Printing Office) US GPO. Final Report on the WPA Program 1935-1943. December 1946.
- (United States Government Printing Office) US GPO. 111-H.R.1 *The American Recovery and Reinvestment Act*. 2009.
- Jasper Jacob van Dijk. Local employment multipliers in US cities. *Journal of Economic Geography*, 17(2): 465–487, 2016.

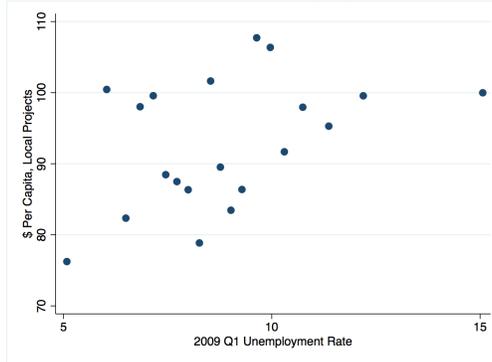
Daniel Wilson. Fiscal Spending Jobs Multipliers: Evidence from the 2009 American Recovery and Reinvestment Act. *American Economic Journal: Economic Policy*, 4(3):251–282, August 2012.

Danny Yagan. The Enduring Employment Impact of Your Great Recession Location. April 2016.

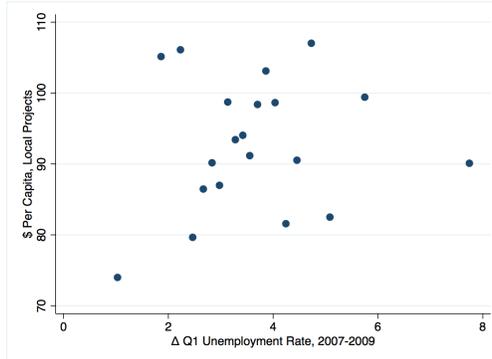
Tables and Figures

Figure 1: Relationship between Covariates and ARRA Treatment Variables

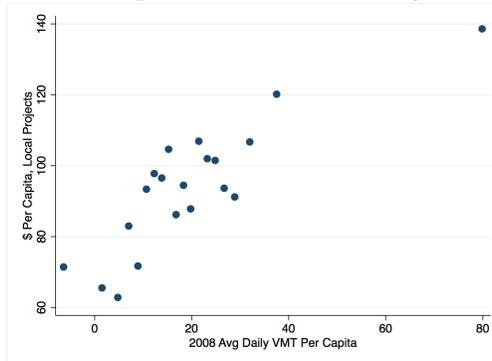
Panel A: 2009 First Quarter Unemployment Rate (LAUS)



Panel B: Change in First Quarter Unemployment Rate, 2007-2009 (LAUS)



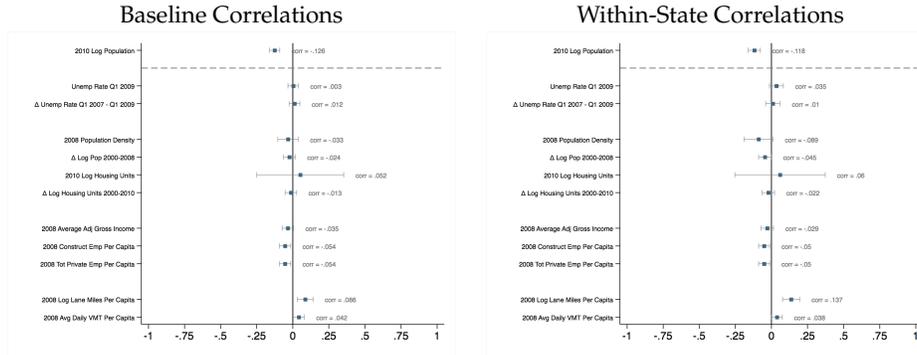
Panel C: 2008 Per Capita Lane Miles of Primary Roads (HPMS)



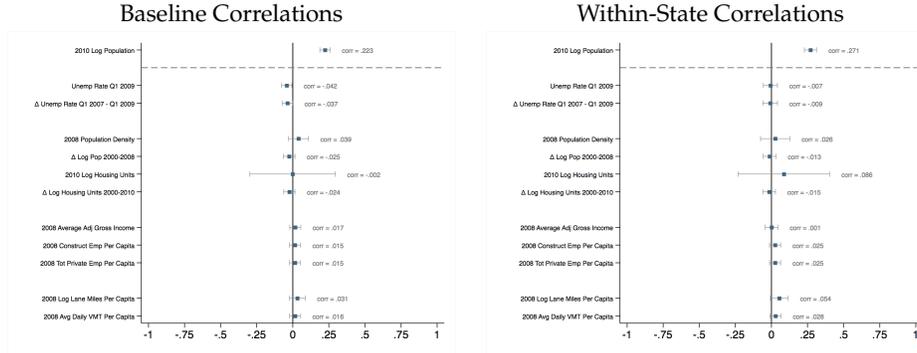
Notes: Each diagram is a binned scatter plot: the sample is divided into twenty equally-sized bins corresponding to vintiles of the x-axis variable, and the average per capita ARRA road spending level within each bin is calculated and plotted against the y-axis. N = 2,922 counties included in primary analysis sample.

Figure 2: Pairwise Correlations of Spending Variables with Covariates

Panel A: Local Project Construction, \$ Per Capita

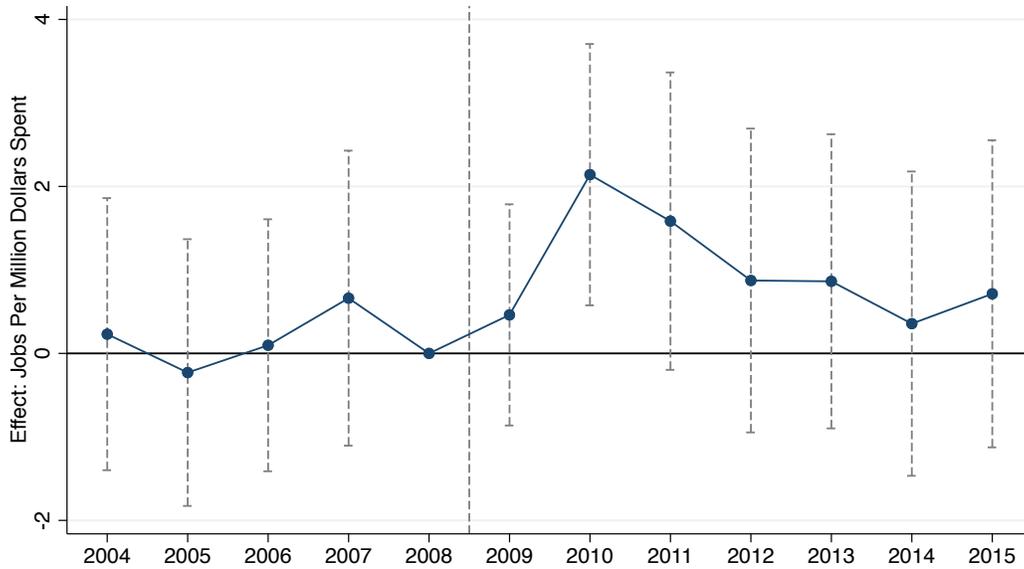


Panel B: Local Vendor Receipts, \$ Per Capita



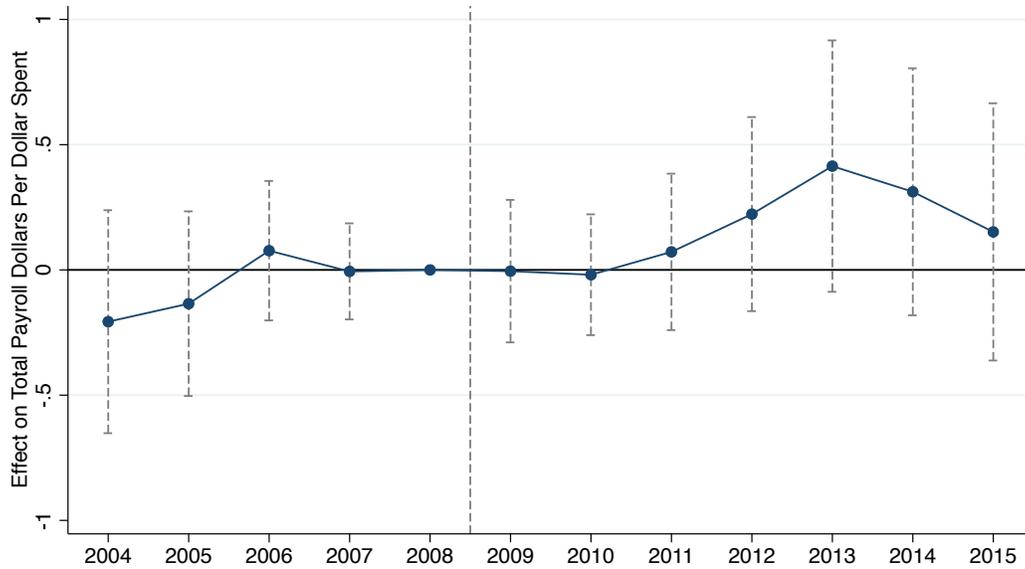
Notes: Each point estimate is the standardized coefficient of correlation between the Recovery Act exposure measure and the stated covariate. 95% confidence intervals implied by the robust standard errors for each point estimate are plotted around each correlation coefficient. Except for the first correlation in each diagram (with 2008 log population), all correlation coefficients are calculated conditional on 2008 log population by including a control in a standardized regression. "Within-state correlations" isolate within-state variation by including state-level fixed effects in these regressions. N = 2,922 counties included in primary analysis sample.

Figure 3: Event Study: Dynamic Effects of Local Construction Spending on Construction Employment



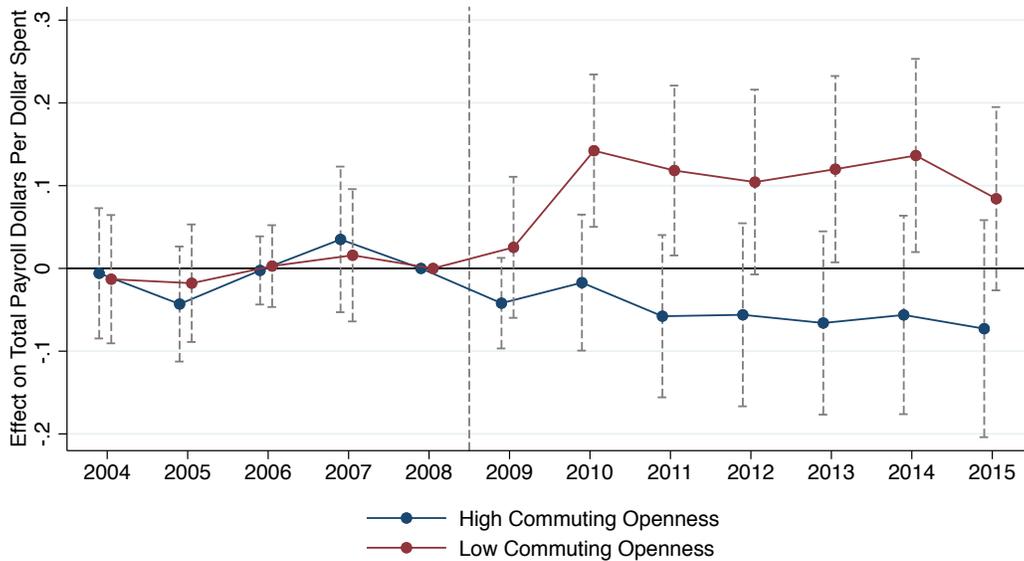
Notes: Figure plots year-specific β^t coefficients estimated jointly from the dynamic difference-in-differences specification in equation 3, also presented in Column 1 of Table 2. The outcome is the county-level annual average construction employment level from the QCEW, both the treatment and the outcome variables are scaled by 2008 population. Each point estimate is the the county-level effect of \$1 Million per capita of Recovery Act road construction spending on construction employment per-capita in the specified year, relative to the 2008 level of the outcome variable. Regression includes state-by-year fixed effects and year-specific controls for 2008 log population. The treatment variable is not time-varying, rather the same treatment variable is interacted with dummies for each outcome year. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. Standard errors are clustered at the county level, and the implied 95% confidence intervals are plotted around each point estimate. N = 35,052, reflecting 2,922 counties included in primary analysis sample.

Figure 4: Event Study: Dynamic Effects of Local Construction Spending on Total Employment

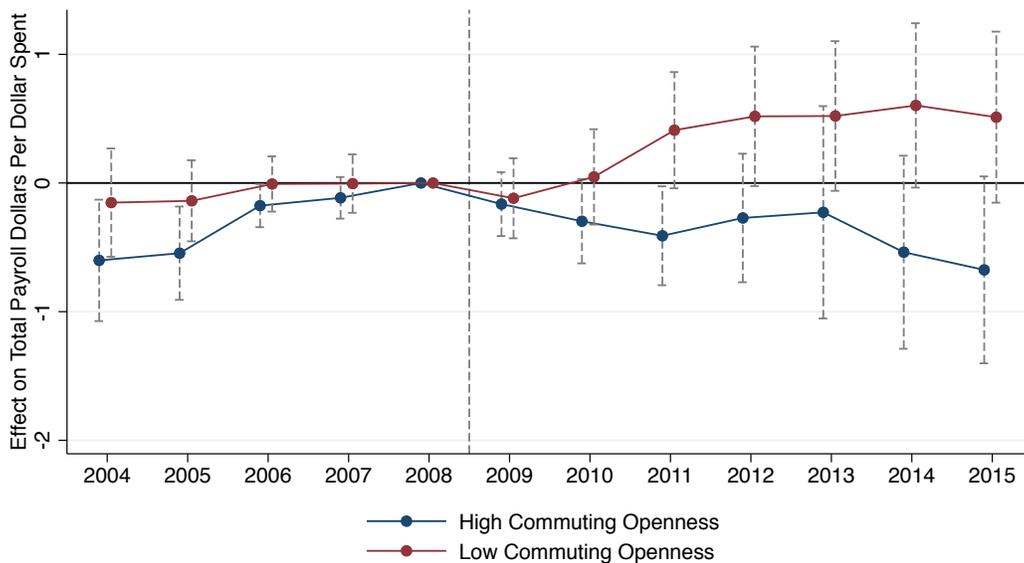


Notes: Figure plots year-specific β^t coefficients estimated jointly from the dynamic difference-in-differences specification in equation 3. The sum of the 2009-2013 effects is displayed in Table 4. The outcome is the county-level annual average total employment level from the QCEW, both the treatment and the outcome variables are scaled by 2008 population. Each point estimate is the the county-level effect of \$1 Million per capita of Recovery Act road construction spending on construction employment per-capita in the specified year, relative to the 2008 level of the outcome variable. Regression includes state-by-year fixed effects and year-specific controls for 2008 log population. The treatment variable is not time-varying, rather the same treatment variable is interacted with dummies for each outcome year. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. Standard errors are clustered at the county level, and the implied 95% confidence intervals are plotted around each point estimate. N = 35,052, reflecting 2,922 counties included in primary analysis sample.

Figure 5: Heterogeneity by Commuting Openness
Outcome: Per Capita Construction Payroll



Outcome: Per Capita Total Payroll



Notes: Figure plots separate year-specific $\beta^{t,group}$ coefficients estimated jointly from a variant of the dynamic difference-in-differences specification in equation () that includes two sets of coefficients, one set interacted with an indicator of whether the county has a higher-than-median share that commutes outside the county (“High Commuting Openness”), and the other set interacted with an indicator of the opposite (“Low Commuting Openness”). Each set of coefficients are plotted as a distinct series indicated in the legend, but are estimated jointly in a regression that includes controls for the indicator variable interacted with year dummies. The sums of the 2009-2013 effects are displayed in Table 9. The outcomes are the county-level annual average per-capita payroll level, both in the construction sector and overall, from the QCEW. Each point estimate is the county-level effect of \$1 Million per capita of Recovery Act road construction spending on payroll dollars per capita in the specified year, relative to the 2008 level of the outcome variable. Dollar figures are in constant 2009 dollars. Regression includes state-by-year fixed effects and year-specific controls for 2008 log population. The treatment variable is not time-varying, rather the same treatment variable is interacted with dummies for each outcome year. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. Standard errors are clustered at the county level, and the implied 95% confidence intervals are plotted around each point estimate. N = 35,052, reflecting 2,922 counties included in primary analysis sample.

Table 1: Summary Statistics

Counties (N= 2,922)	Sample Total Sum	Mean	Median	Mean if > 0	Median if > 0	Standard Dev.	95th Percentile	Nonzero Obs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>All Projects</i>								
Project Expenditure (\$Millions)	23,210	7.94	1.96	9.52	2.99	23.30	32.90	2,437
Project \$ Per Capita	83.82	180.97	53.95	216.99	73.43	602.52	603.42	2,437
<i>Projects With Vendor Information</i>								
Project Expenditure (\$Millions)	16,780	5.70	1.17	7.72	2.53	14.90	32.60	2,174
Project \$ Per Capita	60.60	147.26	35.17	197.92	65.42	560.03	528.55	2,174

Notes: Table summarizes county-level Recovery Act highway spending exposure for 2,922 counties in primary analysis sample.

Table 2: Dynamic County-Levels Effects of Spending on Construction Employment and Pay

	Effect of \$1 Million Per Capita on Constr. Jobs Per Capita:		Effect of \$1 Per Capita on:	
	All Construction	Heavy and Civil Constr.	Per Capita Constr. Payroll	Average Annual Constr. Salary
	(1)	(2)	(3)	(4)
<i>Effect in:</i>				
2004	0.231 (0.831)	0.113 (0.267)	0.015 (0.037)	0.262 (1.529)
2005	-0.229 (0.815)	0.202 (0.247)	-0.005 (0.036)	0.430 (1.522)
2006	0.098 (0.770)	0.017 (0.237)	0.030 (0.034)	1.089 (1.445)
2007	0.662 (0.901)	0.147 (0.210)	0.055 (0.040)	0.858 (1.491)
2009	0.461 (0.676)	0.083 (0.176)	0.025 (0.030)	0.766 (1.234)
2010	2.141*** (0.799)	0.533** (0.229)	0.093*** (0.035)	2.973** (1.396)
2011	1.584* (0.909)	0.593** (0.280)	0.079** (0.039)	2.539 (1.606)
2012	0.873 (0.928)	0.461 (0.286)	0.050 (0.040)	-0.977 (1.712)
2013	0.864 (0.899)	0.594** (0.294)	0.054 (0.040)	-0.842 (1.628)
2014	0.357 (0.930)	0.173 (0.313)	0.050 (0.042)	0.832 (1.708)
2015	0.714 (0.938)	0.534* (0.307)	0.050 (0.043)	0.610 (1.726)
Sum 2009-2013:	5.924* (3.462)	2.265** (1.078)	0.301** (0.151)	
N Counties	2921	2921	2921	2921
Observations	35,052	35,052	35,052	35,052
R-squared	0.864	0.780	0.864	0.768

Notes: Table displays year-specific β^t coefficients estimated jointly from the dynamic difference-in-differences specification in equation (1), each column reports the output from a separate regression. In columns 1 and 2, the outcomes are the per-2008-capita county-level annual average construction (NAICS 23) employment level and “Heavy and Civil Construction” (NAICS 237) employment level, respectively. The outcome in Column 3 is the per-2008-capita construction total payroll level. The outcome in Column 4 is the average annual salary in the construction sector, this variable is the only outcome not in per capita units. Dollar figures are in constant 2009 dollars. Each point estimate is the the county-level effect of \$1 Million per capita of Recovery Act road construction spending on the outcome variable in the specified year, relative to the 2008 level of the outcome variable. Regressions include state-by-year fixed effects and year-specific controls for 2008 log population. The treatment variable is not time-varying, rather the same treatment variable is interacted with dummies for each outcome year. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. The sum of the coefficients for years 2009-2013 are reported with standard errors in the “Sum 2009-2013” row. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10% , * indicates two-sided p-value less than 5%.

Table 3: Robustness of 2010 Construction Employment Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Effect of \$1 Million per Capita on Construction Jobs per Capita	2.054** (0.752)	2.141** (0.798)	1.874** (0.896)	2.441** (0.865)	1.852** (0.672)	1.854** (0.669)	1.608** (0.673)
<i>R-squared</i>	0.935	0.939	0.952	0.945	0.954	0.954	0.954
Effect of \$1 per Capita on Construction Payroll per Capita	0.103** (0.034)	0.094** (0.035)	0.085** (0.039)	0.103** (0.039)	0.072** (0.033)	0.070** (0.033)	0.057* (0.034)
<i>R-squared</i>	0.935	0.938	0.952	0.947	0.952	0.952	0.952
2008 Population Control	x	x	x	x	x	x	x
State x Year FE		x		x	x	x	x
CZ x Year FE			x				
Intensive Margin Only				x			
Pre-Period Industry and Demographic Controls					x	x	x
Pre-Period Unemployment Controls						x	x
Pre-Period Road Controls							x
N obs	5844	5842	5736	4344	5,056	5,056	4,796
N counties	2922	2921	2868	2172	2528	2528	2398

Notes: Table displays 2008-2010 difference-in-differences coefficients estimated from the two-period specification in equation 4. Each point estimate is obtained from a separate regression. In all columns the outcome is the per-2008-capita county-level annual average construction employment level. Each point estimate is the the county-level effect of \$1 Million per capita of Recovery Act road construction spending on the outcome variable in 2010, relative to the 2008 level of the outcome variable. "Intensive Margin Only" restricts the sample to counties with positive spending levels. "Pre-period Industry and Demographic Controls" include 2006 and 2008 log employment and log payroll, both total and in the construction sector, as well as a control for population density and 2000-2008 log changed in population. Dollar figures are in constant 2009 dollars. "Unemployment Controls" include 2007 and 2009 Q1 average unemployment from the LAUS. "Road Controls" include 2008 primary road lane miles per capita and 2008 average daily vehicle miles travelled per capita from the HPMS. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. The size varies across specifications due to missing covariates. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10% , * indicates two-sided p-value less than 5%.

Table 4: County-Levels Effects of Spending on Total Employment and Pay

	2010 Effect		Total 2009-2013 Effect	
	(1)	(2)	(3)	(4)
Effect of \$1 Million per Capita on Total Jobs per Capita	-2.257 (2.708)	-0.579 (2.787)	1.997 (14.60)	-2.456 (15.39)
Effect of \$1 per Capita on Total Payroll per Capita	-0.019 (0.123)	-0.097 (0.130)	0.685 (0.671)	0.289 (0.707)
<i>Controls:</i>				
2008 Population	x	x	x	x
Full Controls		x		x
State x Year FE	x	x	x	x
N obs	5844	5842	5736	4344
N counties	2922	2921	2868	2172

Notes: Each point estimate is obtained from a separate regression. In all columns the outcome is the per-2008-capita county-level annual average total employment level across all sectors. "2010 Effect" indicates estimate is 2008-2010 difference-in-differences coefficient estimated from the two-period specification in equation 4. "Total 2009-2013 Effect" is the sum of the 2009, 2010, 2011, 2012, and 2013 coefficients from the dynamic specification in equation 3, corresponding to the "Sum 2009-2013" estimates in Table 2. "Full Controls" includes "Pre-period Industry and Demographic Controls," "Unemployment Controls," and "Road Controls" as defined in Table 3. The size varies across specifications due to missing covariates. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10%, ** indicates two-sided p-value less than 5%.

Table 5: Results at Varying Levels of Spatial Aggregation

Aggregation Level	Effect of \$1 Million Per Capita on Jobs Per Capita:				Effect of \$1 Per Capita on Total Payroll \$ Per Capita:			
	Construction		All Employment		Construction		All Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
County	2.054** (0.752)	2.185** (0.673)	-2.710 (2.757)	-2.392 (2.971)	0.102** (0.034)	0.105*** (0.034)	-0.088 (0.128)	-0.129 (0.139)
<i>N</i>	2922	2522	2922	2522	2922	2522	2922	2522
Commuting Zone	3.336** (1.511)	2.235 (1.529)	1.152 (5.886)	-4.390 (5.273)	0.204** (0.086)	0.168* (0.093)	0.120 (0.290)	-0.101 (0.284)
<i>N</i>	690	668	690	668	690	668	690	668
Metropolitan Area	0.344 (2.680)	-0.310 (2.152)	-2.607 (9.258)	-9.339 (9.520)	0.196 (0.147)	0.041 (0.125)	-0.307 (0.497)	-0.499 (0.502)
<i>N</i>	336	322	336	322	336	322	336	322
State	8.062 (10.614)		52.019 (39.712)		0.386 (0.462)		0.531 (3.381)	
<i>N</i>	94		94		94		94	
<i>Controls</i>								
2008 Population	x	x	x	x	x	x	x	x
Additional covariates		x		x		x		x

Notes: Table displays 2008-2010 difference-in-differences coefficients estimated from the two-period specification in equation 4 estimated at differing levels of aggregation. Each point estimate is obtained from a separate regression. Outcomes are as defined in Tables 2-4. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. . "Additional Covariates" includes "Pre-period Industry and Demographic Controls" as defined in Table 3. Regressions do *not* include state-by-year fixed effects. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10% , * indicates two-sided p-value less than 5%.

Table 6: Effects of Exposure to Spending in Commuting-Proximate Counties

	Effect of \$1 Million Per Capita on Jobs Per Capita:				Effect of \$1 Per Capita on Total Payroll \$ Per Capita:			
	Construction		All Employment		Construction		All Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Exp^{out}	1.307 (1.416)	1.452 (1.359)	1.987 (5.803)	-0.023 (5.950)	0.077 (0.063)	0.098 (0.064)	0.241 (0.287)	0.265 (0.299)
Exp^{in}	5.610 (3.787)	4.253 (3.801)	21.840 (15.449)	-12.713 (15.524)	0.353** (0.170)	0.233 (0.178)	0.117 (0.755)	-0.271 (0.760)
<i>Controls</i>								
2008 Population	x	x	x	x	x	x	x	x
Additional covariates		x		x		x		x

Notes: Table displays 2008-2010 difference-in-differences coefficients estimated from the two-period specification in equation 4. Exp^{out} is the weighted average per-capita spending level in counties that residents of county c commute to to work, weighted by how many residents commute to each outside county and excluding c 's own exposure. That is, $Exp_c^{out} \equiv \sum_{d \neq c} \lambda_{cd}^{out} \times \frac{Spend_d}{Pop_{2008}_d}$. Exp^{in} measures how much outside construction work would be done by local employees if the probability that a project in outside county d were constructed by a worker based in observation county c were the same as the probability (the observed share) of workers in d who commute from c ; this amount is then scaled by local resident population in c . That is, $Exp_c^{in} \equiv \frac{1}{Pop_{2008}_c} \sum_{d \neq c} \lambda_{dc}^{in} \times Spend_d$. The inbound and outbound commuting shares are obtained from the 2006-2010 American Community Surveys. Each point estimate is obtained from a separate regression. Specifications are as in the first row of Table 5, see notes for details. "Additional Covariates" includes "Pre-period Industry and Demographic Controls" as defined in Table 3. Regressions do *not* include state-by-year fixed effects. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10% , * indicates two-sided p-value less than 5%.

Table 7: Vendor Location versus Project Location: Difference in Difference Estimates

Panel A: 2010 Effect on Construction Jobs Per Capita							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Spending on Projects in Region (\$Millions per Capita)	2.116*** (0.754)	2.077*** (0.675)	2.874* (1.533)	1.601 (1.575)	2.702 (3.251)	-0.261 (2.054)	8.069 (16.849)
Payments to Vendors in Region (\$Millions per Capita)	-1.319 (1.403)	2.270** (1.023)	1.033 (1.542)	1.610 (1.451)	3.250* (1.710)	2.977* (1.791)	-0.010 (17.672)
Region Level	County	County	CZ	CZ	MSA	MSA	State
<i>Controls</i>							
2008 Population	x	x	x	x	x	x	x
Additional covariates		x		x		x	
N obs	5,844	5,044	1,380	1,336	672	644	94
N counties	2922	2522	690	668	336	322	47

Panel B: 2010 Effect on Construction Jobs Per Capita							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Spending on Projects in Region (\$Millions per Capita)	-2.596 (2.773)	-2.776 (2.975)	0.920 (5.930)	-5.025 (5.286)	-2.711 (9.301)	-9.724 (9.577)	-0.014 (56.867)
Payments to Vendors in Region (\$Millions per Capita)	-2.430 (4.751)	8.028* (4.773)	2.260 (5.882)	5.267 (6.047)	1.227 (10.801)	4.802 (9.882)	74.259 (45.874)
Region Level	County	County	CZ	CZ	MSA	MSA	State
<i>Controls</i>							
2008 Population	x	x	x	x	x	x	x
Additional covariates		x		x		x	
N obs	5,844	5,044	1,380	1,336	672	644	94
N counties	2922	2522	690	668	336	322	47

Notes: Table displays 2008-2010 difference-in-differences coefficients estimated from the two-period specification in equation 4 estimated at differing levels of aggregation, including both treatment variables in each equation. Coefficients in each column of each panel are estimated jointly. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. "Additional Covariates" includes "Pre-period Industry and Demographic Controls" as defined in Table 3. Regressions do *not* include state-by-year fixed effects. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10% , * indicates two-sided p-value less than 5%

Table 8: Heterogeneity and Nonlinearities in Effects on Construction Employment

	(1)	(2)	(3)	(4)	(5)	(6)
Per Capita Spending	1.852** (0.672)	-2.811 (2.030)	12.274** (5.577)		5.103** (1.595)	
(Per Capita Spending) ²		10.066** (4.324)				
Per Capita Spending x Log (2008 Pop)			-1.069* (0.550)			
Per Capita Spending x (Pop > Median)				-0.116 (0.838)		
Per Capita Spending x (Pop ≤ Median)				2.803** (0.898)		
Per Capita Spending x % Commute					-9.632** (3.877)	
Per Capita Spending x (% Commute > Median)						-0.104 (0.807)
Per Capita Spending x (% Commute ≤ Median)						3.423** (0.983)
<i>Controls:</i>						
2008 Population	x	x	x	x	x	x
Additional Covariates	x	x	x	x	x	x
Interaction Main Effect	x	x	x	x	x	x
State x Year FE	x	x	x	x	x	x
N obs	5056	5056	5056	5056	5056	5056
N counties	2528	2528	2528	2528	2528	2528

Notes: Table displays 2008-2010 difference-in-differences coefficients estimated from the two-period specification in equation 4, including interactions with the treatment effect. All coefficients in a column are estimated jointly. Whenever a variable is interacted with the treatment variable (times the Post indicator), I also include a control for the interaction of that covariate and the Post indicator. (*Pop > Median*) is an indicator for whether the observation county had a 2008 resident population greater than median. Commuting shares are obtained from the 2006-2010 American Community Surveys. "Additional Covariates" includes "Pre-period Industry and Demographic Controls" as defined in Table 3. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10% , * indicates two-sided p-value less than 5%.

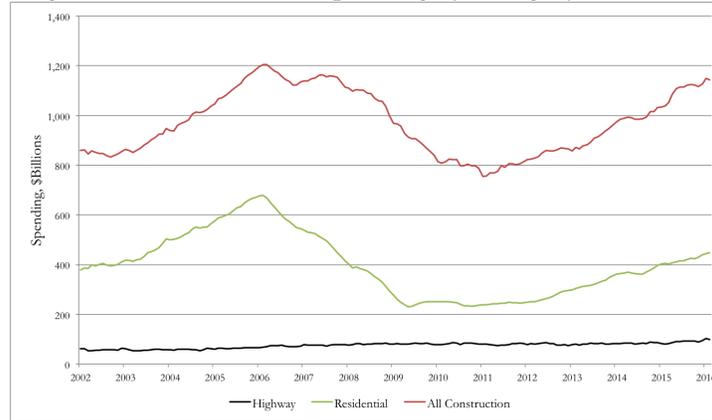
Table 9: Effect Heterogeneity across High and Low Commuting Share Counties

	Effect of \$1 Million Per Capita on Jobs Per Capita:				Effect of \$1 Per Capita on Total Payroll \$ Per Capita:			
	Construction		All Employment		Construction		All Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Per Capita Spending x (% Commute > Median)	-0.104 (0.807)	-4.241 (3.982)	-5.495 (3.666)	-32.44 (20.10)	0.066 (0.056)	-0.239 (0.193)	-0.297* (0.167)	-1.371 (0.951)
Per Capita Spending x (% Commute ≤ Median)	3.423** (0.983)	10.67** (4.375)	1.469 (4.258)	17.24 (21.94)	0.116*** (0.042)	0.510** (0.212)	0.048 (0.189)	1.378 (0.951)
P Value, Coefficients Equal	0.004	0.009	0.189	0.021	0.01	0.008	0.157	0.027
2010 Effect	x		x		x		x	
Sum 2009-2013 Effect		x		x		x		x

Notes: Table displays effects of the primary spending exposure variable interacted with indicator for whether or not the observation county had an outside-commuting share that exceeds the sample median. The outcomes are payroll dollars per 2008 capita in constant 2009 dollars. All coefficients in a column are estimated jointly. "Total 2009-2013 Effect" is the sum of the 2009, 2010, 2011, 2012, and 2013 coefficients on the corresponding interaction term from an interacted version of the dynamic specification in equation 3, see Figure 5 for more details. "2010 Effect" indicates the estimates are 2008-2010 difference-in-differences coefficient estimated from the two-period specification in equation 4, corresponding to the specification in Column 6 of Table 8. Whenever a variable is interacted with the treatment variable and any year interactions, I also include a controls for that covariate interacted with the corresponding year indicators. Commuting shares are obtained from the 2006-2010 American Community Surveys. All specifications include control for 2008 log population and "Pre-period Industry and Demographic Controls" as defined in Table 3. All regressions include state-by-year fixed effects. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10%, ** indicates two-sided p-value less than 5%.

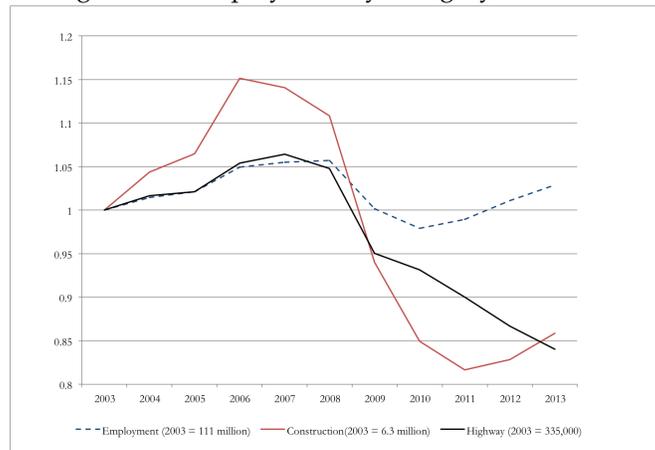
A Appendix: Supplemental Figures

Figure A.1: Construction Spending by Category, 2002-2013



Notes: Figure plots estimates of construction expenditures in the United States reported by the Census Bureau. Spending figures are reported monthly at a seasonally adjusted annual rate. Source: Census (FRED 2016)

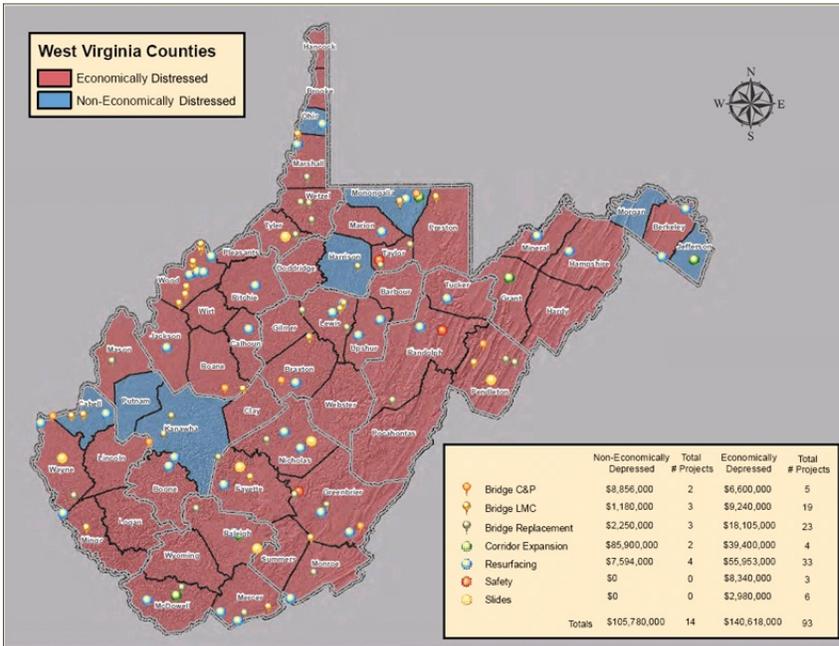
Figure A.2: Employment by Category, 2003-2013



Notes: Figure plots evolution of US total employment, construction employment, and highway/road construction employment (NAICS 2373) tabulated from payroll tax returns in the County Business Patterns Data. Each series is normalized to the 2003 value of that series by dividing by the 2003 levels listed in the legend.

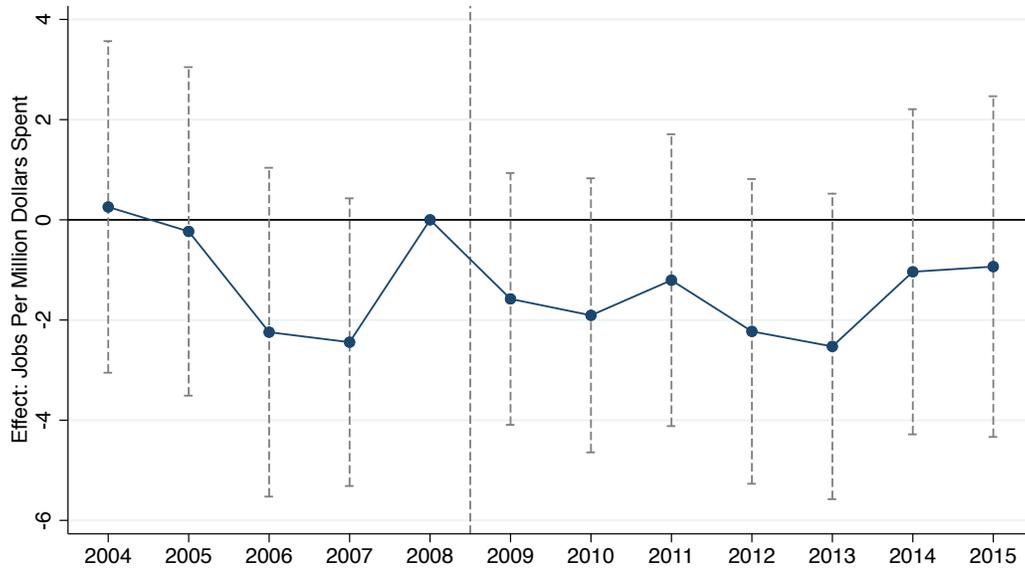
Figure A.3: Federal Highway Administration Example of “Economically Distressed Areas”

Example Map with Projects Added



Notes: Figure displays counties in West Virginia designated as “economically distressed” (in red), and “not economically distressed” (in blue), and was posted on the website of the Federal Highway Administration 2010 as an example of how States might make this designation. Eighty-five percent of counties were designated as “economically distressed”.

Figure A.4: Event Study: Dynamic Effects of Local Vendor Receipts on Construction Employment



Notes: Figure plots year-specific β^t coefficients estimated jointly from the dynamic difference-in-differences specification in equation 3, where the treatment is defined as the per-capita amount of Recovery Act highway construction dollars (in \$ Millions per capita) that went to firms located in the observation county according to the vendor records. The outcome is the county-level annual average construction employment level from the QCEW, both the treatment and the outcome variables are scaled by 2008 population. Each point estimate is the the county-level effect of \$1 Million per capita of Recovery Act road construction receipts on construction employment per-capita in the specified year, relative to the 2008 level of the outcome variable. Regression includes state-by-year fixed effects and year-specific controls for 2008 log population. The treatment variable is not time-varying, rather the same treatment variable is interacted with dummies for each outcome year. Year dummies and interactions for 2008 are omitted due to inclusion of identify county-level fixed effects. Standard errors are clustered at the county level, and the implied 95% confidence intervals are plotted around each point estimate. N = 35,052, reflecting 2,922 counties included in primary analysis sample.

B Appendix Tables

Table A.1: Summary Statistics of Recovery Act Road Construction Projects

	All	Pavement/Resurfacing	Bridge and Road Construction/ Reconstruction
<i>Number of Projects</i>	6,993	2,452	903
Project Expenditure (\$):			
<i>Mean</i>	1,666,175	1,734,447	1,880,179
<i>Median</i>	569,490	679,317	666,600
<i>SD</i>	4,013,480	3,302,186	3,690,526
<i>Max</i>	105,000,000	51,200,000	46,900,000
<i>Number of Contracts Awarded</i>	13,581	3,936	2,386
Contract Value (\$):			
<i>Mean</i>	805,596	1,043,307	655,777
<i>Median</i>	165,285	275,162	157,154
<i>SD</i>	2,578,052	2,519,135	1,979,237
<i>Max</i>	86,800,000	51,200,000	42,900,000
Distance of Vendor from Project (Mi):			
<i>Mean</i>	77	70	85
<i>Median</i>	42	43	46
<i>SD</i>	133	112	148
<i>Max</i>	2,295	1,827	2,295
Projects by Number of Vendors			
<i>1</i>	4,963	1,840	436
<i>2</i>	1,099	417	148
<i>3 or More</i>	931	195	319

Notes: Table summarizes all individual Recovery Act infrastructure projects (sub-awards) administered by the Federal Highway Administration (FHWA) in the 48 contiguous states which include information on at least one vendor. These figures are aggregated to the county level to generate the primary treatment variable summarized in Table 1. "Pavement and Resurfacing Projects" restrict to the sample of projects with terms in the text field containing the project description that corresponding to resurfacing, paving, widening, lane addition, or curve improvement. "Bridge and Road Construction" restricts to projects with descriptions that include words related to building, constructing, bridge work, and tunnel work.

Table A.2: Robustness to Other Definitions of Recovery Act Road Construction Spending

<i>Spending Variable</i>	Effect of \$1 Million Per Capita on Jobs Per Capita:				Effect of \$1 Per Capita on Total Payroll \$ Per Capita:			
	Construction		All Employment		Construction		All Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Primary Spending Definition	2.141*** (0.798)	5.924* 3.462	-2.257 (2.708)	1.997 14.60	0.093*** (0.035)	0.301** 0.151	-0.019 (0.123)	0.685 0.671
All Projects inc. Missing Vendors	1.894*** (0.669)	5.715* (2.939)	-1.001 (2.640)	3.228 (13.50)	0.073** (0.030)	0.236 (0.131)*	0.117 -0.755	0.685 (0.658)
Paving / Resurfacing Only	2.401* (1.258)	8.682 5.727	1.799 (4.967)	21.77 24.91	0.083 (0.057)	0.346 0.261	-0.130 (0.363)	1.042 1.097
Bridge and Other Major Construction Only	3.054 (2.488)	4.547 11.01	-7.041 (7.489)	-5.281 43.46	0.194* (0.112)	0.407 0.493	0.113 (0.290)	0.236 2.126
2010 Effect	x		x		x		x	
Sum 2009-2013 Effects		x		x		x		x

Notes: Table shows robustness of main estimates to other definitions of Recovery Act road construction expenditure. Sample is N = 2921 counties in primary sample. All spending variables are in per-2008-capita units. "All projects including missing vendors" corresponds to the variable summarized in the top rows of Table 1. "Paving/resurfacing only" and "Bridge and other major construction only" are variables aggregated from the corresponding projects in Table A.1, and restrict to projects with vendor information as in the primary treatment variable definition. Each point estimate is obtained from a separate regression. "2010 Effect" indicates estimate is 2008-2010 difference-in-differences coefficient estimated from the two-period specification in equation 4. "Total 2009-2013 Effect" is the sum of the 2009, 2010, 2011, 2012, and 2013 coefficients from the dynamic specification in equation 3, see notes to Table 4 for more details. All specifications include state-by-year fixed effects and year-specific controls for 2008 log population. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10%, ** indicates two-sided p-value less than 5%.

Table A.3: Effects on Other Outcomes

	Effect of \$1 Million Per Capita on		Effect of \$1,000 Per Capita on Unemployment Rate (pp)			
	Building Construction Jobs Per Capita:		LAUS		ACS	
	(1)	(2)	(3)	(4)	(5)	(6)
Spending on Projects in County (\$Millions per Capita)	0.140 (0.284)	0.160 (0.332)	0.194 (0.203)	0.282 (0.227)	-0.219 (1.189)	-0.837 (1.139)
<i>Controls</i>						
2008 Population	x	x	x	x	x	x
Additional covariates		x		x	x	
N obs	5,842	5,056	5,842	5,056	1,440	1,418
N counties	2921	2528	2921	2528	720	709

Notes: Table displays 2008-2010 difference-in-differences coefficients estimated from the two-period specification in equation 4. "Additional Covariates" includes "Pre-period Industry and Demographic Controls" as defined in Table 3. All regressions include state-by-year fixed effects. Sample size varies across columns due to missing values of covariate and outcome variables. Standard errors are clustered at the county level. * indicates two-sided p-value less than 10%, ** indicates two-sided p-value less than 5%.