HARVARD Kennedy School
TAUBMAN CENTER
for State and Local Government

Putting America to Work, Where? The Limits of Infrastructure Construction as a Locally-Targeted Employment Policy

by

Andy Garin

Taubman Center for State and Local Government

May 2016

Taubman Center Working Paper

WP – 2016 – 01

Putting America to Work, Where? The Limits of Infrastructure Construction as a Locally-Targeted Employment Policy

Andy Garin*

This Version: May 11, 2016

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 Recovery Act to study the relationship between construction work and local employment growth. I show that the method for allocating funds across space facilitates a plausible selection-on-observables strategy. However, I find a precisely-estimated zero effect of spending on road construction employment–or other employment–in the locale of the construction site. Reported data on vendors reveal this is because the majority of contractors–selected by competitive bidding–commute from other local labor markets. I also find no robust effect in the locales of the contractors' offices, but argue that this source of variation does not capture an economically meaningful local demand shock. I conclude that infrastructure construction is not effective as a way to stimulate local labor markets in the short-run so long as projects are allocated by competitive bidding.

*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

Between 2007 and 2009, the United States economy shed over eight million non-farm jobs including nearly two million in construction. Yet 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). If the decline in employment varies from place to place, can locally targeted policy spur job growth in specific locales?

This paper studies a frequently proposed policy intended to boost local construction employment and overall economic health: spending on "shovel-ready" public infrastructure construction projects. Such projects are intuitively appealing as way to boost employment in distressed areas because they can be geographically targeted in a highly visible way. These projects create a need for construction laborers to go to work in precise locations, which may later benefit from workers' spending their earnings in turn. Moreover, when labor markets are slack, infrastructure projects can be of substantial value and the costs of construction may be low. 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 begin construction promptly, with priority given to economically distressed areas. This policy intervention provides a unique opportunity to study the local employment effects of infrastructure spending.

I test for local employment effects using a unique dataset with detailed information both about each road construction project funded by the Recovery Act and about the vendors who worked on those projects. The national coverage and rich geographic detail of these data enable me to test how road spending plausibly effects local employment. First, I test for such effects in the local highway construction sector for each project, as any broader economic effects should stem from the project's direct effect on employment. To the extent that construction employees might have been doing construction work elsewhere in the absence of stimulus spending, it does not suffice to count bodies at project sites–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. While one might be concerned that funds were systematically targeted to places with unobservably worse downturns, I surprisingly find little evidence of *any* targeting based on observable employment trends.

Across a range of increasingly demanding specifications, I consistently obtain the same result: Highway spending has little to *no* impact on local highway construction employment in the county or broader commuting zone of construction projects. Given this finding, one should not expect any broader local employment effects—nor do I find evidence of any. Although these findings are consistent with a crowd-out story, the unique nature of the data enables me to offer a simpler explanation. Contract firms are selected by competitive bidding, and firms will bid for projects within a large radius if they can offer a lower price than competitors; therefore, most construction laborers working on a given project are employed by firms in different labor markets. Thus, the first step in the stimulus transmission mechanism is already highly geographically diffuse, limiting the ability of policymakers to use construction spending to boost local labor demand. Unsurprisingly, Recovery Act project vendor firms were disproportionately based in locales with large highway construction sectors. It is plausible that the funds flowing to these "employer" counties led to better economic outcomes than if the Recovery Act had not been enacted. However, because the ultimate "destination" of the federal expenditures was selected by competitive bidding and not project site-selection concerns, and because these winning firms would have provided heightened competition on non-Recovery-Act projects in counterfactual scenarios, it is difficult to make conclusive claims about counterfactuals using cross-sectional variation in the data.

While a long line of research has attempted to assess the aggregate employment effects of government purchases, this paper stands apart it in its direct focus on local employment effects. What warrants such explicit focus on local labor market effects? First, there are substantive reasons: place-based policies are commonly employed as strategies to benefit local residents and workers[1]; hence, is important to know whether or not the impacts of location-based policies effectively target local individuals. 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 heterogenous nature of the recession—for instance, Yagan 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 under-employment. In theory, this could create an opportunity for a spatially targeted counter-cyclical employment policy to improve on UI extensions or Keynsian policies with no spatial aspect. 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 assess the plausibility of using targeted infrastructure construction to boost local employment given the mobility of the agents involved.

Another rationale for a local focus is the lack of valid counterfactuals for entire macroeconomies—to the extent that regions such as states or counties form local "macroeconomies," one might hope to learn about the effects of macroeconomic policy by studying variation in regional outcomes[2]. Most relevant

[1]In this regard, all policies enacted by local governments are "place-based"

[2]A long literature has studied fiscal multipliers using macroeconomic time series data, Ramey (2011), Barro and Redlick

to this paper, work by Leduc and Wilson (2012; 2013) have estimated state-level GDP and employment multiplier effects of highway spending both during and prior to the Recovery Act, exploiting variation from institutional rules governing the allocation of funds. Also closely related is recent work by Dube and Kaplan and coauthors that draws from the same underlying data on Recovery Act awards to study county-level macroeconomic multipliers. These papers examine all types of awards, not just highway construction grants, and show: first, that the allocation of across-the-board spending is uncorrelated with observable indicators of economic distress (Boone et al., 2014), as I find in the case of highway spending; and second, that the local level multipliers are small and only different than zero once one stratifies the sample based on severity of downturns after 2009 (Dube et al., 2014)[3]. Similarly, recent work by 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 find a cost per job of $30,000 per year. In addition, work by Chodorow-Reich et al. (2012), Feyrer and Sacerdote (2011), Wilson (2012), and Nakamura and Steinsson (2014) has used state-level variation to study multipliers.

Such approaches come at a cost, however—while more disaggregated datasets offer richer variation, the relationship between regional and aggregate outcomes becomes less clear at finer levels of aggregation. Nakamura and Steinsson (2014) demonstrate that the relationship between regional multipliers and aggregate multipliers is ambitious in a models with immobile agents, price rigidities, and non-tradable goods. In their model, local spending can drive up short-run prices for local nontradables; 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. 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[4]. 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 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 ARRA and about the overall performance of the highway-construction sector in the context of the broader construction boom, bust, and recovery. Section 3

(2011), Pereira (2000), and Blanchard and Perotti (2002) are notable examples.

[3]This paper differs from their work in important ways. In general, the site of an "award" is not informative about the location of the economic activity—for example, purchases of new buses will be listed as an award taking place at the a 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 are 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.

[4]This point is made formally in Kline and Moretti (2014a)

describes the data sources and sample construction. Section 4 examines the effects of road construction work on local labor markets and argues for the plausibility of the selection-on-observables assumption. Section 5 considers the zero result and provides evidence that Recovery Act contractors were generally not located in the same local labor market as their projects. Section 6 proposes tests for effects in the locale of the Recovery Act *contractors* but ultimately argues that the variation in the data does not identify local labor demand shocks. Section 7 concludes.

2 Background

2.1 Highway Spending in the Recovery Act

On January 6, 2009, the United States Congress passed the American Recovery and Reinvestment Act, a stimulus bill that has authorized approximately $831 billion in supplemental spending to date(CBO, 2012). In practice, less than $100 billion of the Recovery Act authorizations went towards infrastructure spending, and approximately $43 billion went to transportation projects. The $27 billion that the Act authorized the Federal Highway Administration (FHWA) to disburse for highway, bridge, and tunnel[5] improvement projects may appear small relative to the overall bill. Yet the $27 billion in supplemental spending was nontrivial relative to the baseline annual appropriation level of approximately $40 billion[6]—in fact, the entire value of road construction work completed in 2007 was $102 billion.

Although the bill text reads over 400 pages long, the Recovery Act itself provided only basic guidance as to how FHWA funds should be used; the Act's text does not specify specific projects. The distribution of Recovery Act highway funds across states was wholly determined by the preexisting federal apportionment formula for distributing annual surface transportation appropriations across states[7]. The law left it to individual states to propose which projects to undertake with their supplemental funds, subject to broad guidelines and FHWA approval, but no record exists of any project's being rejected.

Specifically, the law specified two criteria for consideration of road, bridge, and tunnel projects: "[P]riority shall be given to projects that are projected for completion within a 3-year time frame, and are located in economically distressed areas" (GPO 2009). The first criteria essentially established a "shovel-readiness" requirement—this was a stimulus bill, so it prioritized projects that would begin promptly. This was partially enforceable; all projects had to be obligated (that is, to have had construction contracts signed) within a year, or else the funds would be retracted. This generally ruled out the use of funds for projects that required extensive new planning or design. 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 time frame. Appendix Figure 1 shows that over two-thirds of spending fell into this category.

The latter directive was to prioritize spending in "economically distressed areas". This might give rise

[5] Throughout the paper, I refer to highway, road bridge, and road tunnel construction work collectively as "highway" construction.
[6] https://www.fhwa.dot.gov/policy/olsp/financingfederalaid/approp.cfm
[7] See Leduc and Wilson (2013)for detailed discussion of these formulas

to concern that variation in funding levels was systematically correlated with mean-reverting shocks. In practice, 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.[8]As shown below, there is no evidence this provision had any effect.

While local authorities could choose project sites, they were required to select vendors through standard competitive bidding processes. That is, firms were chosen strictly on competitive cost-bid bases, not by discretion.[9]The Recovery Act gave federal authorities no special authority to interfere with the procurement process. FHWA had to approve all *projects*, but it did not have any say over which *firms* should be put to work on specific projects.

2.2 The Highway Construction Industry in the Great Recession

Before turning to the primary analysis, it is worth placing this policy intervention in context of the broader construction boom and bust that occurred over the past decade. Figure 1 plots the annual value of highway construction spending relative to overall construction spending during the boom, bust, and recovery. Three facts about highway construction spending during this period deserve particular attention. First, 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. Second, as shown in Figure 1A, highway spending rose dramatically during the housing boom, and it actually continued to grow through 2008—before the Recovery Act was enacted—despite the dramatic decline in residential construction. Yet the Recovery Act did not notably increase the total level of highway construction spending, nor did highway spending fall during this period. Third, this $27 billion of Recovery Act highway spending paled in comparison to the pre-recession level of construction spending—which had exceeded $1 trillion dollars annually—and certainly in comparison to the $400 billion decline in *annual* residential construction from 2006 to 2009. Thus, it is *apriori* implausible that Recovery Act highway spending could have absorbed much of the slack in overall construction demand, although it might certainly have staved off more rapid, dramatic declines in demand for highway construction.

Likewise, employment in the highway construction sector[10] is a small share of total construction employment and of private employment in the United States. At the start of the housing boom in 2003, there were

[8]State 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.

[9]Title 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.

[10]I refer to NAICS code 2373 (highway, bridge, and tunnel construction) as the "highway construction" sector. Firms in this sector account for nearly all construction contracts funded by the Recovery Act.

approximately 111 million private-sector jobs in the U.S. economy, 6.3 million of which were construction jobs and 335,000 of which were highway-construction jobs. Figure 2 plots the evolution of employment in these 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) through 2006 and then declined dramatically—by two million jobs—come 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 through 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. Once again, the small size of the sector and the net loss of jobs together suggest that highway construction work would not have been able to absorb any additional slack in construction labor markets during the downturn.

While publicly-funded highway construction evokes the memory of millions of Works Progress Administration (WPA) workers building roads in the 1930s, modern highway construction is quite different. The WPA was essentially a program that directly hired workers and worked with local governments to assign them to jobs—typically labor-intensive work like clearing simple dirt roads and covering them with gravel using shovels. Compensation of workers accounted for a full 69 percent of spending on WPA projects (US GPO, 1946). By contrast, modern private road construction firms are relatively high-tech, relying heavily on expensive, specialized machinery. In 2012, compensation of employees represented only 28 percent of the cost of construction[11]. Highway construction is also a relatively skill-intensive sector, and compensation is therefore high for blue-collar work: The average annual earnings for all employees in 2012 was $56,276, while the average annual earnings for a typical blue-collar laborer was $51,824. Thus, if demand were perfectly inelastic, and supply were Cobb-Douglas in labor and other factors, an extra million dollars in demand should result in an extra $280,000 in labor compensation, accounting for about 5.5 job-years. Put differently, absent crowd-out, the cost of one highway job-year through highway construction spending should be about $182,000. If supplemental federal spending crowds-out local government or private construction spending (either by raising prices or by "flypaper effects" which do not increase nearly-satiated demand, but simply change who foots the bill), then this figure is a lower bound for the cost of a job-year. One should note that even at this bound, the consumption generated by a single highway-construction job would need to support at least five additional jobs to yield a multiplier as large as the ones estimated in Suarez-Serrato and Wingender (2014).

Given $27 billion in Recovery Act funds for highway construction, 5.5 job-years per one million dollars suggests these projects would have supported 150,000 construction job-years in gross terms (prior to accounting for crowd-out). Direct surveys of contractors involved in Recovery Act construction found that number

[11]This 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 subsector. Figures are taken from the 2012 Census of Construction.

of full-time-equivalent employees at work on Recovery Act highway projects peaked at 40,000 in September 2010 and declined thereafter (CBO, 2012). Interestingly, this figure seems lower than one might expect by the above calculation. However, these survey-based figures are not necessarily reliable or consistently reported, nor do they speak to the counterfactual employment levels in the absence of stimulus spending.

3 Data Sources and Sample Construction

The data on Recovery Act expenditures and vendor payments used in this paper were obtained from the website of the Recovery Act Accountability and Transparency Board at www.recovery.gov. In general, the federal government does not systematically report all awards made to state and local governments funded by federal appropriations. However, to ensure transparency, the Recovery Act required public disclosure of all awards authorized by that law—in particular, awardees were required to disclose the ZIP codes where and the purposes for which the funds were used, so that 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[12].

In addition, local governments were requested to report information about the vendors contracted to work on each project. For most states, the reported data on payments to vendors accounts for the majority of funds awarded. For these states, I am therefore able to determine which establishments won bids to be vendors for most projects.[13] In these cases, I can determine how much was paid to each contracting establishment, as well as the name, industry, and address of the establishment[14]. 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 at a national level. To my knowledge, this is the first such dataset with national coverage

As the location of vendors plays a central role in my analysis below, I focus primarily on projects with at least some information on vendor payments available in the "continental" 48 states. Table I details the construction of the analysis sample. Comprehensive data on projects are available for 47 states[15], accounting for $24 billion in expenditures. However, only $15.4 billion of these funds can be attributed to

[12]A "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.

[13]The 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.

[14]The 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.

[15]Michigan is excluded from these data, as most projects were erroneously reported as located in the State Capitol complex.

vendors—this shortfall is due in large part to several states that appear to have not systematically reported vendor information. I therefore restrict my analysis to the 39 states with comprehensive reporting about highway-construction vendors. I further drop all projects with *no* information on vendors from my primary analysis; however, I examine robustness of my results to the inclusion of these projects in the Appendix. This limits my sample to $16 billion of spending of projects, for which $14.6 billion can be attributed to payments to vendors. Panel A of Table I summarizes the projects comprising the primary analysis sample.

To conduct local labor-market-level analyses, I aggregate these data to the county level. I separately calculate the amount spent on *projects* located within each county and the amount paid to *vendors* whose offices are listed as being in that county. As will be noted below, these two amounts can differ substantially. In some analyses, I aggregate the data further into "commuting zones" (CZs), groups of counties meant to represent self-contained commuting regions (akin to Metropolitan Statistical Areas, but with full national coverage). The geographies were originally constructed by Tolbert and Sizer (1996); I use the county groupings defined in Autor and Dorn (2013). The county- and CZ-level treatment variables are summarized in Panel B of Table I.

The primary outcomes are private-sector employment levels within a locality, broken out by industry. Importantly, the employment concept I employ is based on place of work, not place of resident. Any mentions of "local" employment growth or decline that follow refer to the number of jobs reported by employers as taking place in that locale, *not* the number of local residents at work. To the extent that individuals commute beyond counties and CZs, the level of employment at local establishments may differ substantially from the level of employment along local residents. Both concepts are meaningful measures of local economic health. However, there is only scant data available on employment by place-of-residence, primarily collected through household surveys that are too small to use for local-level analyses.[16]

County-level employment figures are taken from two distinct datasets: the County Business Patterns (CBP), released by the Census Bureau, and the Quarterly Census of Employment and Wages (QCEW), released by the Bureau of Labor Statistics. Both datasets report total employment and payroll by industry-county pair, where the location of the employee is on an establishment basis. While these datasets are meant to measure similar objects (employment levels in counties), they are based on different data sources and thus reflect slightly different employment concepts. The CBP draws from annual payroll-tax records and therefore reports the total counts of distinct employees and wages paid by establishments at any point during the year—in these data, twelve distinct workers with monthlong contracts would count as twelve employees. The QCEW data draws on administrative records from unemployment insurance system to tabulate monthly employment and payroll levels at establishments. Annual employment figures are *averages* of monthly levels—in these data, twelve distinct workers with monthlong contracts would count as one annual employee. While the QCEW offers higher-quality data at higher frequencies, the CBP has higher-quality information as finer levels of detail, due to the nature of reporting rules. Thus, I use CBP outcomes in my primary

[16]The 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.

analysis, though I show robustness to use of either data source. Both datasets measure employment at the location of the employing establishment, not at the workers' places of residence.

Additionally, I supplement these data with additional control variables and 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 surveys 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 DOT. Throughout, I scale variables by 2010 population from the full-county Decennial Census, which I refer to as "per capita" units.

4 The Local Employment Impacts of Highway Construction Projects

4.1 Empirical Methodology

To assess whether roadwork boosts employment in the vicinity of project sites, I adopt a straightforward difference-in-differences design, allowing for heterogeneity in treatment intensity. The "treatment," in this case, is the amount of Recovery Act funds spent per-capita on projects contained within a locale. My baseline specification tests how differences in employment levels between the two years preceding the passage of the Recovery Act (2007–2008) and the three years during which the Recovery Act funds were intended to be spent (2009–2011) varied across locales that received different amounts of spending after the passage of the Recovery Act.

The baseline estimating equation is

$$Y_{ist} = \alpha_{st} + \delta_i + \beta \text{Spend}_i * \text{Post}_t + \gamma X_i * t + \epsilon_{ist} \tag{1}$$

In this specification, Y_{ist} is the (per capita) sectoral employment level in locale i in state s during year t, and the coefficient of interested is β, the average effect of one million dollars per capita on the outcome during each post-passage year. Because both the dependent and independent variables are scaled by the same quantity in all cases, β can be interpreted as the *number of jobs added per million dollars spent in each post-period year*[17]. To the extent that the spending may not have been smooth during the post period, this specification may create attenuation bias in the estimates. However, the data do not specify the dates during

[17] To convert to job-years, this figure should be multiplied by 3.

which construction took place, only the data of the approval of the award[18] and the quarter of completion[19].

I include a locality-level fixed effect in all regressions (which absorbs the "Spend" main effect) to net out permanent level differences across locales, as well as year fixed effects, which net out aggregate trends. In most specifications, I let the year fixed effect vary by state—with state-by-year fixed effects, β is only identified off of within-state variation. Finally, while fixed attributes are absorbed by the locale fixed effects, I allow localities to follow differential trends on the basis of pre-period covariates[20].

The standard identifying assumption required to estimate a causal effect in this specification is that locales with different levels of spending *would* have experienced identical trends in the absence of any Recovery Act spending. 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. This is a direct violation of the parallel-trends assumption, and it would lead to a spurious positive estimate of the effect of spending. However, in the case of the Great Recession, one might separately worry that hard-hit counties experienced *persistent* adverse trend-breaks in 2009 that continued throughout the observation window. If harder-hit regions received more funds in this scenario, it might be the case that places with more spending would have done comparatively worse in the absence of intervention. Again, the parallel-trends assumption is violated, but the violation leads to a downward bias. The violation of the identifying assumptions in these cases arises not from the presence of differential trends across counties *per se*, but rather systematic selection into treatment-intensity related to these latent trends.

If there was systematic selection into treatment in the case of Recovery Act spending, it is likely that locales that received higher levels of spending were *observably* different. In particular, if regions were targeted on the basis of economic distress, one should observe funds flowing to locales that experienced greater distress in the period leading up to the passage of the Recovery Act. If spending is correlated with observable variables that in turn may correlate with latent trends, one can address this bias by controlling for local trends that vary with these observables, as in Equation (1). With the inclusion of the $\gamma X_i * t$, identification requires a weaker selection-on-observables assumption, $E[\epsilon_{ist} * Spend_i * Post_t | X_i * t] = 0$.

A simple method of probing the validity of the selection-on-observables assumption is to test for robustness to varied control sets, in the spirit of Altonji, Elder, and Taber (2005). If estimates are robust as the choice of X_i varies, then any unobservable source of bias would have to be systematically correlated with the error term ϵ_{ist}, but systematically *uncorrelated* with all of the observable dimensions of heterogeneity. This insight guides my inference below.

[18] Approval of the award generally took place before the bidding process and therefore does not line up with timing of construction

[19] Few projects ended prior to 2011, thus this creates little variation in the timing. Specifications run using quarterly data that distributed funds proportionately across all quarters between the award date and completion quarter yield similar results to those reported in Table II.

[20] In previous versions, these trends were allowed to vary for each year to year interval by including a set of terms $\sum_{\tau=2008}^{2011} \gamma^{\tau} X_i * \mathbf{1}[\tau = t]$; this did not affect the results in any substantial way.

4.2 Spatial Allocation of Recovery Act Infrastructure Spending

The Recovery Act encouraged state and local governments to create jobs in distressed areas, but where did the funds flow in practice? Column A of Figure 3 examines how county-level per-capita spending varied with local economic conditions nationwide. The first presents a somewhat striking fact: spending levels were no higher in regions with higher unemployment levels—if anything, counties with the lowest unemployment rates had the highest rates of per-capita construction spending. While the local unemployment data are measured with considerable noise, they represent the best information policymakers could have used to base any intentional targeting of spending to slack markets. This appears to present *prima facie* evidence that funds were not targeted to observably distressed labor 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 (more negative growth) and local spending.

Figure 4 plots the pairwise correlations between Recovery Act construction spending and a broader range of covariates, as well as correlations of the residuals of the same variables regressed on state-level indicator variables. While smaller, less affluent counties that experienced less population growth during the prior decade received slightly more funds for local highway construction projects, there is indication that places were targeted more if they had worse downturns, whether measured by change in unemployment and changes in broader employment rates. Similar findings hold when one limits the analysis to within-state variation—there is no relationship between spending and observable indicators of local labor market distress. If counties were in fact systematically targeted due to unobserved adverse shocks, these unobserved shocks must have had no correlation with unemployment levels, changes in unemployment levels during the downturn, or changes in the construction sector and private sector employment rates. Is is hard to imagine any dimension of short-run labor market distress that would not covary with any of these variables.

If the selection of projects was not driven by local economic conditions, then how was the geographic distribution of construction determined? Figures 3 and 4 offer at least one explanation: places with more lane-miles of major roads per resident tended to receive more funding for road construction[21]. Looking within states, lane-miles of primary roads are the single strongest correlate of Recovery Act construction spending. This is unsurprising: since the majority of Recovery Act road-construction funds went to pavement improvement and road-widening projects, work naturally occurred in places where more roads already existed. The other covariates that have notable correlations with Recovery Act highway spending all covary with the per-capita stock of roads: average daily vehicle-miles traveled (VMT) on local roads per resident, lower population densities, and more Interstate Highway miles per capita.

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 base-line stock of roads but *not* observable

[21]"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. Note that a mile of a road with two lanes running in each direction constitutes four lane-miles of road.

indicators of economic distress. Leduc and Wilson (2013) argue that measures of local road milage and usage (VMT and lane-miles, in particular), are completely orthogonal to latent drivers of short-run growth, and they 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 all results estimated below are robust to *either* conditioning on these variables or using them as instruments for local project spending.

4.3 Results

I first examine estimates of effects on local highway construction employment, as highway construction is a key first step in any plausible local stimulus transmission mechanism. Table II presents difference-in-difference estimates of β in increasingly demand specifications of equation (1). The first column presents the bivariate regression. The second column adds state-by-year fixed effects, isolating the within-state variation. The third column adds basic controls for population and housing size and density. The fourth column adds a much richer set of occupation-by-gender composition variables from the 2000 Census, basic demographic variables from the 2005–2009 ACS, and 2008 income level controls. The fifth column includes additional controls concerning the highway stock. The sixth column additionally controls for fourth-order polynomials in 2003 and 2006 per capita highway employment, construction employment, and private employment, capturing the pre-trends in outcome variables prior to the observation window. Since Leduc and Wilson (2012; 2013) consider the pre-existing highway stock to be a source of identifying variation rather than confounding variation, column 7 employs lane miles per capita as an instrument for local spending, including county and state-by-year fixed effects but no other controls[22]. All standard errors are clustered at the state level, and observations are weighted by log 2010 population to reduce noisy outcomes in sparsely-populated counties.

The results are consistent across specifications: There appears to be *no* association between higher local construction spending levels on and highway employment at the county level. Not only is there no association the the bivariate regression (whether or not one focuses on within-state variation), the estimates are also not affected by the inclusion of basic controls for region size, demographics, attributes of the 2008 highway stock, or the 2003 and 2006 employment rates in the highway construction sector, the broader construction sector, or total private sector employment. The standard errors on these estimates are tight in nearly every specification, and the implied confidence intervals rule out local highway construction employment effects larger than approximately 1 local job-year per million dollars spent.

Likewise, one finds no clear association between local construction spending and *total* private employment levels across all specifications. Given the lack of an effect on local highway employment, it is not clear why one should expect a broader employment effect. The standard errors are wider in this case—given the small size of the intervention relative to the dramatic variation in employment trends, these tests offer less power to reject a zero effect. Nonetheless, the confidence intervals in the central specifications again rule out impacts of more than six jobs per year (18 job-years) per million dollars. These results thus appear to contrast starkly with prior estimates in the literature, which find effect sizes as large as thirty job-years per million.

[22] A positive first-stage is evident in Figure 4 as discussed above.

To probe this finding further, I estimate a variant of (1) that allows one to plot an "event study" around the enactment of the Recovery Act:

$$Y_{ist} = \alpha_{st} + \delta_i + \sum_{\tau=2006}^{2013} \beta^{\tau} \text{Spend}_i * \mathbf{1}[t{=}\,\tau] + \sum_{\tau=2006}^{2013} \gamma^{\tau} X_i * \mathbf{1}[t{=}\,\tau] + \epsilon_{ist} \tag{2}$$

I estimate equation (2) with per-capita highway construction employment and per-capita private employment as outcomes, using the control set from column 6 in Table II[23]. I plot the estimated β^{τ} coefficients year by year in Figure 5. Rather than observing no pre-period effect and then a significant effect during the immediate aftermath of the Recovery Act, one finds no significant change in the coefficients between 2008 and 2009–2011. If anything, a small positive pre-trend appears to be present between 2006 and 2008 for both outcomes, which then levels off at the start of the policy. In addition to ruling out short-run effects on local labor markets due to construction activity, these findings cast doubt on the possibility that the improved infrastructure directly boosted employment in the medium-run, as counties with more project spending saw no differential growth as late as 2013, well after most work had been completed.

The null result is not an artifact of the outcomes chosen. Table III displays estimates of the preferred specification of (1) for a wide range of outcome variables. These included the per capita total wage bill in highway construction, total construction, and total private employment; employment and wage bill levels in the construction sector and the broader private sector as reported in the QCEW; employment in retail employment (a standard example of a non-tradable sector) and residential construction (which serves as a placebo, since there is no reason to expect an effect); and LAUS-based local unemployment and labor force participation rates. In nearly every case, the effects are small, and large effects are ruled out in every case.

One exception are the LAUS outcomes. The results are consistent with positive effects on heightened labor force participation and lower unemployment. A key difference between these outcomes and the others is that they measure employment outcomes *for individuals living in the same locale as the projects,* while the CPB and QCEW outcomes measure employment *by employers in the same locale as the project.* Therefore, it is possible that construction work benefits people who live near project sites, without boosting employment figures at firms near project sites. However, these data are largely imputed from sparse data and are measured with considerable noise—while the point estimates are large, the standard errors are wide. While I cannot rule out beneficial impacts on local populations, no clear inference can be drawn.

While it is possible that these results are still downward-biased due to selection on unobservable factors correlated with discrete negative trend breaks in 2009, these latent confounders would have to have nearly implausible properties. In particular, given the robustness of the zero effect to all specifications, any latent confounding factors would have to be completely uncorrelated with any of the controls for observable characteristics of population and housing trends, demographic composition, employment trends by sector, and road stock attributes. It is difficult to imagine what latent driver of local labor market outcomes could

[23]Because this control set includes 2006 outcome levels, the estimates for 2006 are mechanically zero

satisfy these conditions and still also have economically large effects.

5 Zero Effect: Crowd-Out or Commuting?

What explains the zero effect? One standard story is crowd-out—federal spending in a region may discourage other highway spending by lower levels of government or private citizens so the aggregate spending level is unchanged. However, in this setting, local crowd-out would not necessarily imply aggregate crowd-out. Consider the following hypothetical: after the passage of the Recovery Act, a state government plans one large project in every county, and plans to pay for these projects with a combination of Recovery Act and locally-sourced funds; once the total pot of funds is fixed, designating a project to be the nominal "recipient" of Recovery Act funds *mechanically* crowds-out non-federal spending. Nonetheless, *complete* crowd-out along these lines would still be extreme.

Yet there is a much more straight-forward explanation: more often than not, the firms selected via competitive bidding to work on a given project are based in a different locale than the project. If a construction firm's office is in a different locale than the Recovery Act project for which they are vendors, and one uses employer-reported outcome data where locations refer to the permanent office of employment and not the daily site of work, then there is no reason to expect to find any highway-employment effect near the project. Moreover, unless workers systematically live (and consume) closer to Recovery Act projects than their employing offices, there is no reason to expect any local labor impacts of any type near the project site. In fact, in the analysis sample, fully *78 percent* of vendors have offices in different counties than where the project sites are[24]. Even more strikingly, the coefficient of correlation between per-capita spending on projects in a county and per-capita payments to firms for Recovery Act roadbuilding work to firms sited in the same county is only 0.039.

This finding suggests the county level is too narrow of a definition of a labor market in which to expect to find local results[25]. Thus, one might look for effects at higher levels of aggregation. Accordingly, Panel B of Table II and column 2 of Table III present results from the same specifications shown above estimated at the Commuting Zone level. By construction, these regions are designed to represent self-contained local

[24] This 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.

[25] In light of this finding, the earlier finding of positive effects in the county-level LAUS data despite the precise zero effects in the CBP and QCEW outcomes may make sense in the case that the vendors offices were located elsewhere, but local residents were hired by these distant firms to work on these projects. This is highly plausible. In this scenario, local construction might provide work for local residents, even if their paychecks come from distant firms. Unfortunately, there are no publicly-available sources of data that measure labor market outcomes by place of residence with precision and at high frequencies—the primary source of information on LFP and unemployment are the Current Population Surveys, which are the primary source data for the LAUS data used here. Thus, while I cannot rule out beneficial impacts on local workers, neither can I establish positive effects with confidence.

labor markets. Panel B of Figure 5 plots CZ-level event studies corresponding to equation (2). One should note that although local effects may be more plausible at more aggregated levels, the number of observations falls (by a factor of four in the case of CZs). This widens confidence intervals in specifications with richer covariate sets.

I find slightly larger effects on highway construction employment; however, the effect size is sensitive to controls and not significant with larger control sets. The event-study plot in Figure 5 shows point estimates that are still noisy, though consistent with a heightened level of local construction employment in 2009 and 2010. Nonetheless, the estimates are still economically small, on the order of one job-year per million dollars.

The estimates of effects on total employment are also noisy but larger. The magnitude of these point estimates, centered around thirty jobs per million dollars, are particularly questionable given the small effect on highway-construction employment. Taking these results seriously would imply that a policy that directly raises construction employment by one job in turn begets an additional twenty to thirty jobs, which appears implausible. In fact, the plot in Figure 5 complicates this finding, as a notable pre-trend is present. However, the results in Table III show no clear effect as one examines different QCEW-based measurements of the same outcome. Having obtained such wide standard errors, it is difficult to draw any clear inference about total employment effects at the CZ level. Given the smaller number of observations and the small size of the treatment relative to the size of variation in CZ-level local labor market outcomes, it appears these specifications are underpowered.

A bigger issue is that even large commuting zones may not be sufficiently self-contained—in fact, *55 percent* of vendors have offices *in different commuting zones* than the project sites in my data. Even at this broader level, the coefficient of correlation between per capita spending on projects in a county and per capita payments to firms for Recovery Act roadbuilding work to firms sited 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. While these distances seem large, they are actually consistent with standard bidding behavior for large projects. Construction offices routinely bid for projects at considerable difference, particularly when the demand for construction work is lower. It is not uncommon for construction workers to drive 90 minutes to a project site on a daily basis[26]. This creates considerable difficulty in defining a sub-state local labor market notion in which employers and worksites are co-located[27].

None of this suggests highway spending has no aggregate effect. Rather, these findings offer new evidence that the transmission mechanism of public-infrastructure construction spending is highly geographically diffuse in nature. While public authorities could carefully target project sites, the competitive bidding

[26] 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.

[27] The state may be the lowest level of geography where this is the case—90 percent of vendors are in the same state as the project sites in my data.

processes mandated by law will likely allocate the labor demand to firms and workers in different local labor markets. If the first-step of the transmission mechanism is so diffuse, then large local multipliers are implausible. Indeed higher-order induced spending by workers and capital owners is likely to be even more geographically dispersed (in terms of where value-added takes place[28]). The zero result is not surprising in this light.

6 Effects in Locales of Recovery Act Contractors

6.1 The Locations of Winning Bidders

An implication of these findings is that the spatial allocation of *employment* by Recovery Act vendors was only loosely related to the spatial allocation of construction projects. One might therefore wonder whether highway-construction spending affected employment in the local labor markets where Recovery Act vendors' offices were located. Causal inference on this subject would require exogenous variation in the level of funds assigned to firms located in different regions.

But to the extent that siting of construction projects cannot explain the spatial allocation of payment to vendors, this allocation is essentially entirely determined by the competitive bidding process. Given assignment by competitive bidding, any residual variation in local firm receipts, conditional on observables, *must* be driven by latent advantages that enable firms to offer lower costs. Any valid source of identifying variation must arise from some competitive advantage for Recovery Act projects that would not have had any counterfactual impact on employment outcomes. I briefly assess the plausibility or implausibility of such a strategy in this section.

If the employers who won bids were typically not in the locale of the project, where were the employers that Recovery Funds went to? While spending on construction did not flow to particularly hard-hit regions, were the employers involved in the projects in places that had lost jobs and therefore had more slack? Or were they in places that were weathering the recession well? The right-hand columns in Figures 3 and 4 show how the per-capita level of Recovery Act payments received by highway construction establishments located in a county covaried with the same sets of observable attributes examined in section 3.2. These plots show that the locations of the employers involved in Recovery Act highway construction work were no more distressed than the counties where the projects took place—as before, there is little, if any, relationship between the level of funds received and indicators of economic slack, such as unemployment or changes in employment levels.

Rather, the best predictor of payments to local firms is the 2008 size of the highway construction sector[29]. Counties with a larger number of individuals employed by local firms in highway construction sector in the year before the Recovery Act were more likely to be where offices of firms selected to work on Recovery-Act-

[28]Spending at a local retailer may not boost local employment if most of the cost of goods sold is value-added produced elsewhere

[29]Notably, highway construction firms are concentrated in a small number of counties, as visible in the maps in Appendix Figure 3 and the high number of zero values for the "payments to local firms" variable in Table I.

funded projects were located. As noted earlier, highway construction is relatively intensive in specialized equipment and skills compared to other forms of construction. It is unlikely that firms would have been able to make expensive additions to their capital stock in a rapid manner during the credit crunch; therefore, any firm able to make competitive bids for Recovery Act projects would have almost certainly had preexisting capacity to take on similar projects. Interestingly, growth or decline in highway-construction employment is not a strong predictor of receiving funds, implying that counties where Recovery Act vendors were located were consistently larger in the years leading up to the Great Recession. The correlations with all other observable economic indicators are close to zero.

6.2 Testing for Differential Outcomes Near Recovery Act Vendors

In light of this fact and of the earlier observation that highway-construction employment was in net decline during the years after the Recovery Act's passage, the likely effect of the Recovery Act was to provide work opportunities and cash flow to pre-existing highway construction employers who otherwise may have been forced to downsize or exit the market outright. Yet directly evaluating the extent to which Recovery Act payments prevented layoffs—and the broader local economy based around those jobs—is difficult, because firms were selected to do Recovery Act construction due to their endogenous productivity attributes. While the ability of local firms to win bids is highly endogenous to later outcomes, it is unclear whether effects estimated in a difference-in-difference regression of (highway) employment growth on payments to local firms would be upwards- or downwards-biased. On one hand, firms with the ability to post winning Recovery Act bids may have been able to be more competitive for other non-Recovery Act projects in counterfactual scenarios and would have fared better regardless of the passage of the law, biasing estimates upwards. On the other hand, the firms that were most competitive in 2009 may have been suffering larger demand shortfalls in their other operations and were competitive due to high levels of idle capacity. These firms may have experience large layoffs, but those layoffs would have only been worse if not for the Recovery Act. In this case, estimates could possibly be downwards-biased.

Yet there is another obstacle to using local variation in Recovery Act payments to determine the effects on local labor market demand. In the absence of the Recovery Act, the vendors in my data would have competed for other projects in a wide geographic radius—projects that nominally "untreated" firms had posted winning bids in the post-Recovery Act period—but they might have lost those bids if facing tougher competition. In this sense, all highway construction firms were "treated" by being given the opportunity to bid on Recovery Act projects. What ultimately matters for economic outcomes is not whether one's portfolio of projects was slated for Recovery Act funding but rather the total demand for a firm's services. Even if the Recovery Act increased construction demand, it is not clear whether the relevant demand shock had any meaningful geographic dimension–precisely because firms bid for projects at great distance. In this sense, the locations of the Recovery Act vendors is arbitrary.

Having acknowledged these limitations, I estimate a variant of equation (1), but using per-capita payments to vendors reported as located within a county as the treatment variable, rather than per-capita spending on projects located within a county

$$Y_{ist} = \alpha_{st} + \delta_i + \beta \text{Recipts}_i * \text{Post}_t + \gamma X_i * t + \epsilon_{ist} \tag{3}$$

For comparability, Table IV reports estimates of equation (3) under the same set of specifications as in Table II[30].

In nearly every specification, the results show no relationship between payments to local firms and relative highway employment growth in the same county. The logic laid out in the previous subsection explains this, and it suggests that the treatment variable does not actually induce a meaningful local demand shock. Note that while these "more-treated" locales do not experience differential *growth*, they have persistently higher highway employment levels, both before and after the passage of the Recovery Act. The final specification, which controls for the 2006 lagged outcome, shows a somewhat positive effect, but this may be driven by mean reversion. If the treatment variable is correlated with 2008 highway employment, and if highway employment is serially but imperfectly autocorrelated, then less of the variation in the outcome will load onto the treatment variable as one looks at years closer to 2006 for mechanical reasons. The event studies in Appendix Figure 4, which are analogous to those in Figure 5, are consistent with this; one observes more loading onto the treatment variable further from 2006, even *within* the pre-period[31] event. The results from regressions where total private per-capita employment is the outcome behave similarly, as do Commuting-Zone-level specifications.

These findings do not offer evidence for or against the effectiveness of Recovery Act spending in promoting local or aggregate employment growth so much as they suggest a methodological limitation of this setting. Estimating local labor market effects requires the isolating shocks to local construction demand. However, I have documented that competition for highway construction projects extends well beyond the locale of the project site and that Recovery Act spending comprised only a minority of highway spending during the recovery period. Thus, the regions that received larger amounts of stimulus payments did not necessarily experience higher demand than highway construction work—it is more accurate to say that the work that these firms did happened to be on Recovery Act projects rather than on projects with other funding sources. One would need data on *all* private and public (and local and federal) construction contracts to fully assess local demand conditions; however, the Recovery Act data on their own do not capture meaningful variation in demand for local firms' services. Put differently: because competitive bidding generally directs work for construction projects to firms located in different locales than the projects at hand, the process does not offer a useful setting to study the effects of local demand shocks, regardless of how one defines the treatment variable.

[30]Since per capita lane-miles of primary roads are not correlated with the treatment variable, I do not report the IV results from column 7 in Table II.

[31]One observes similar patterns using lagged outcomes from alternative years, suggesting that this effect is in fact mechanical

7 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 a good 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. 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.

In all specification, I found that there is little to *no* county-level impact of highway spending on local employment outcomes reported by employers. There appears to be no effect on local highway-construction employment, overall construction employment, or total private-sector employment. Moreover, my estimates are sufficiently precise and robust to rule out large effects on any of these outcomes, in contrast to earlier work. I propose a simple explanation for the ineffectiveness of construction as a place-based employment policy: Contracting firms were selected by competitive bidding, and firms tend to bid for projects in a large radius if they can offer a low price. As a result, most construction laborers working on a given project are employed by firms in different labor markets[32]. I find that this remains the case even at higher levels of aggregation—over 55 percent of vendor offices were located in different *commuting zones* then the projects they worked on, and one finds little evidence of employment impacts in these broader geographies. Thus, the first step in the stimulus transmission mechanism is already highly geographically diffuse, limiting the ability of policymakers to use construction spending to boost local labor demand. This is further compounded by the likelihood that workers may commute from homes that are neither in the locale of the employer or the locale of the project, though I cannot assess this possibility directly in the available data. I also examined outcomes in the locales of the vendors' offices, rather than the locales of the project sites; however, I offered reasons to think that this latter source of variation does not capture an economically meaningful local demand shock, complicating causal inference.

I therefore conclude that infrastructure construction is not effective as a way to stimulate local labor markets in the short run, so long as projects are allocated by competitive bidding. This finding relates directly to the analysis in Monte et al. (2016), who show that the local employment elasticity in response to local expenditures may vary significantly in different settings and depends crucially on commuting behavior. Competitive bidding for highway construction work induces fairly extreme commuting behavior, which in turn diminishes the local employment elasticity to zero. Policymakers looking to increase employment

[32]However, this does not rule out the possibility that these distant firms in turn hired workers who resided closer to project sites. As discussed above, there is no high-quality data on employment outcomes by place of residence. When I use the best available indicators of outcomes by place of residence from the LAUS, I find noisy results consistent with lower unemployment and higher labor force participation rates, but the standard errors were wide.

in a geographically delimited region using construction projects would need to change procurement rules by, for example, instituting a "hire local" policy. While such provisions would likely increase the costs of construction[33], they deserve consideration—if boosting employment is an explicit goal of a policy in addition to the value of infrastructure improvement itself.

These results also indicate that the partial-equilibrium effects of government spending depend crucially on the transmission mechanism, and different forms of government spending will likely affect local economies and labor market segments quite differently. I find substantially smaller effects than other work that examines local multipliers. This may be because other forms of government spending better target local residents. However, my findings suggest the credibility of larger estimates rests on the plausibility of the *local* transmission mechanism—if one finds large employment/income multipliers, one should also observe first-order effects on the direct recipients of government funds. If this link is missing, this casts doubt on the validity of the source of variation used to estimate the multiplier.

However, these results do not imply the Recovery Act highway-construction program was misguided; I can only claim that construction work was not effective in increasing jobs in the short run for employees in the locale of the projects. These finds are not inconsistent with significant state- or national-level macroeconomic multipliers arising from highway spending—such as those estimated by Leduc and Wilson (2013).

Given the lack of a clear short-run effect on local labor markets, 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 to 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 on road construction work are potentially quite high, but they vary significantly depending on the type of project selected—spending on building 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 hard-hit 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 focusing on the spatial nature of the transmission mechanism. Further analysis of microeconomic effects of stimulus spending using disaggregated data is a promising direction for future work and will help supplement theoretical microfoundations 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 at-

[33]In the context of Monte et al. (2016), a "hire local" policy might be recast as a commuting tax.

tached 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

Joseph G Altonji, Todd E Elder, and Christopher R. Taber. Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools. *Journal of Political Economy*, 113(1):151–184, February 2005.

David H. Autor and David Dorn. The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597, August 2013.

Robert Barro and Charles Redlick. Macroeconomic Effects From Government Purchases and Taxes. *Quarterly Journal of Economics*, 126(1):51–102, 2011.

Olivier Blanchard and Lawrence F Katz. Regional Evolutions. 23(1):1–76, 1992.

Olivier Blanchard and Roberto Perotti. An Empirical Characterization of the Dynamic Effects of Changes in Government Spending. *Quarterly Journal of Economics*, 117(4):1329–1368, 2002.

Christopher Boone, Arindrajit Dube, and Ethan Kaplan. The Political Economy of Discretionary Spending: Evidence from the American Recovery and Reinvestment Act,. *Brookings Papers on Economic Activity*, 48(1):375–441, 2014.

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, 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.

Arindrajit Dube, Ethan Kaplan, and Ben Zipperer. Excess Capacity and Heterogeneity in the Fiscal Multiplier: Evidence from the Obama Stimulus Package (. August 2014.

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.

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.

Sylvain Leduc and Daniel Wilson. Roads to Prosperity or Bridges to Nowhere? Theory and Evidence on the Impact of Public Infrastructure Investment. *NBER Macroeconomics Annual*, 27:89–142, 2012.

Sylvain Leduc and Daniel Wilson. Are state governments roadblocks to federal stimulus? Evidence from highway grants in the 2009 Recovery Act. (16), 2013.

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.

Emi Nakamura and Jon Steinsson. Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review*, 104(3):753–792, 2014.

Alfredo Pereira. Is All Public Capital Created Equal? *Review of Economics and Statistics*, 82(3):513–518, 2000.

Valerie Ramey. Identifying Government Spending Shocks: It's All in the Timing. *Quarterly Journal of Economics*, 126(1):1–50, 2011.

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.

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: Construction Spending by Category, 2002-2013

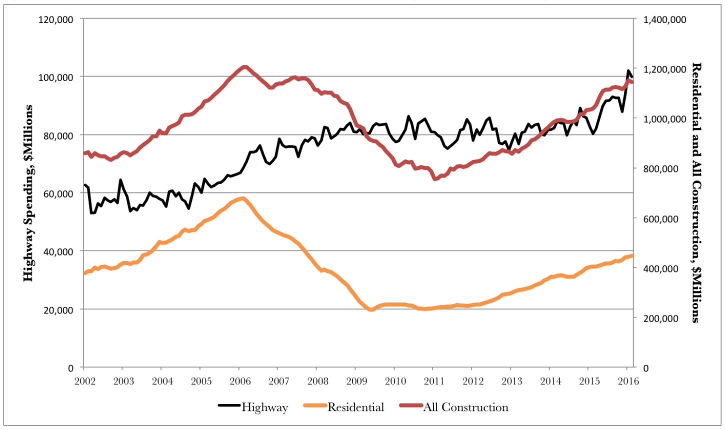

Source: Census (FRED 2016)

Figure 2: Construction Spending by Category, 2003-2013

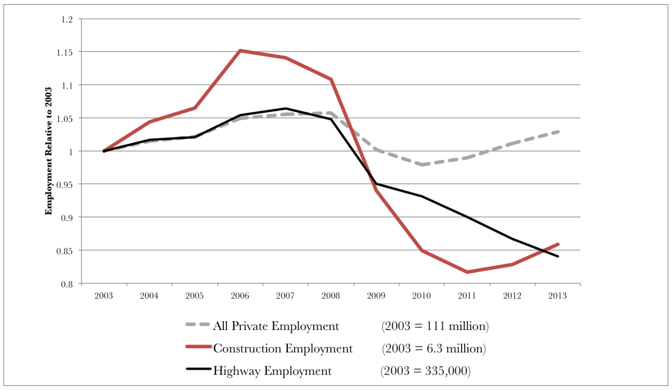

Source: CBP

Figure 3: Relationship between Covariates and ARRA Treatment Variables

Column A: Local Project Construction | Column B: Local Firm Contracts

Covariate: 2009 First Quarter Unemployment Rate (LAUS)

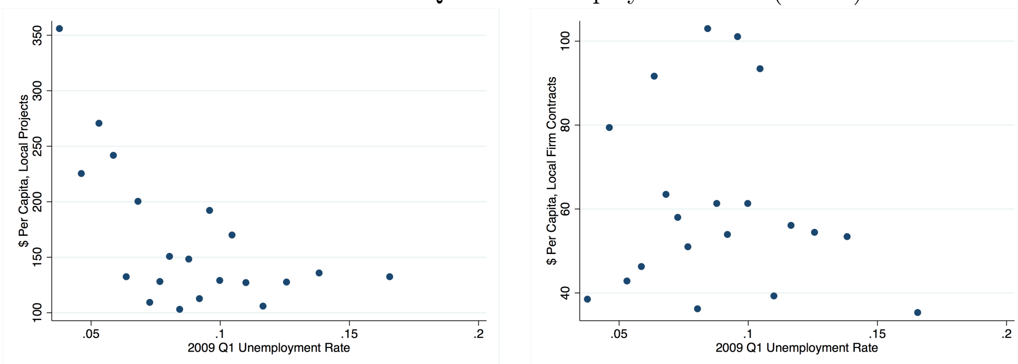

Covariate: 2006 - 2008 Change in Per Capita Total Construction Employment (CBP)

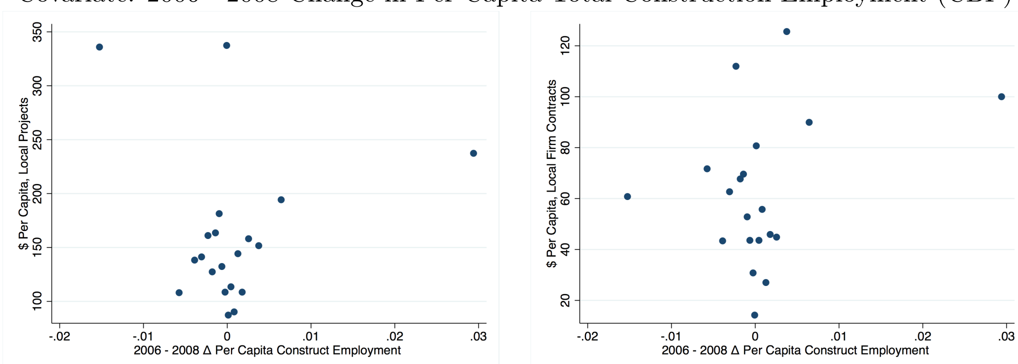

Covariate: 2008 Per Capita Lane Miles of Primary Roads (HPMS)

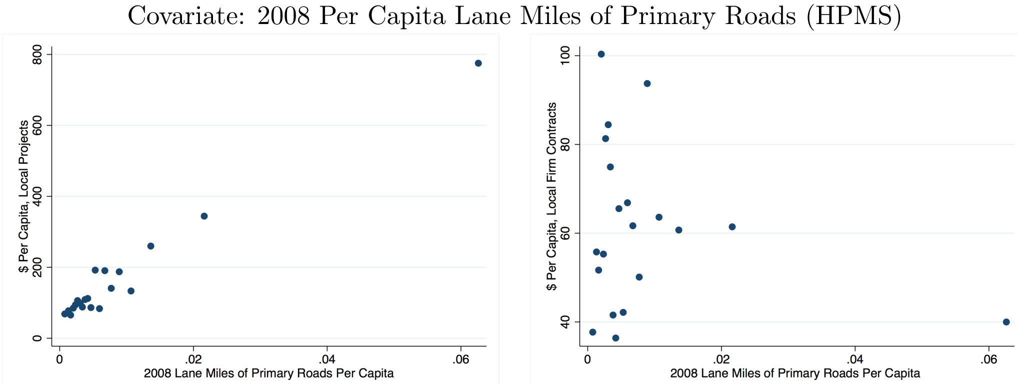

Covariate: 2008 Per Capital Highway Construction Employment Level (CBP)

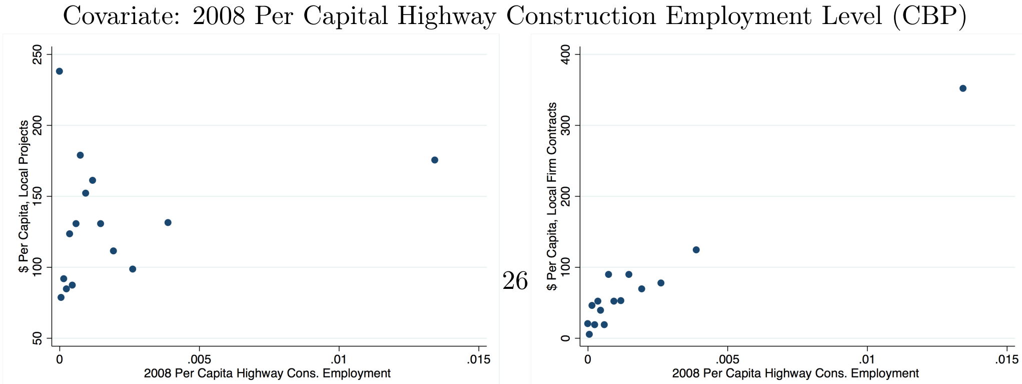

Notes: Each diagram is a binned scatter plot, displaying mean levels of the specified treatment variable within quantile bins of the covariate.

Figure 4: Pairwise Correlations of Spending Variables with Covariates

Panel A: County-Level Correlations

$ Per Capita, Local Project Construction

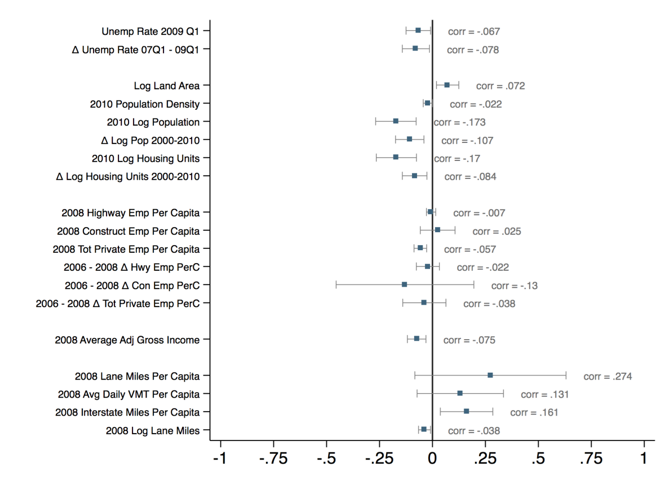

$ Per Capita, Local Firm Contracts

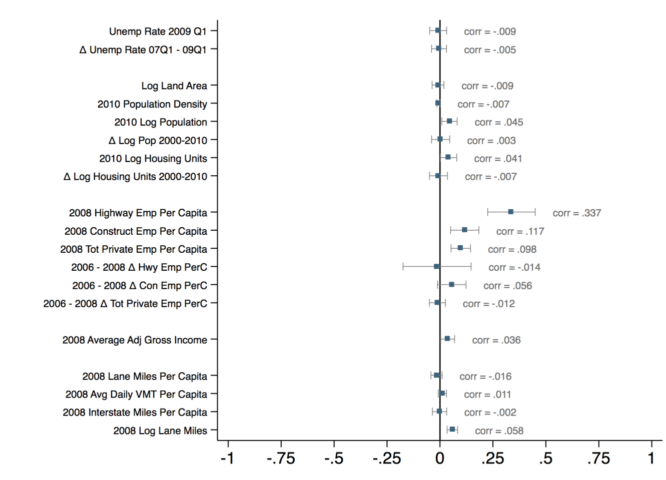

Panel B: Correlations of Residuals, State Fixed Effects

$ Per Capita, Local Project Construction

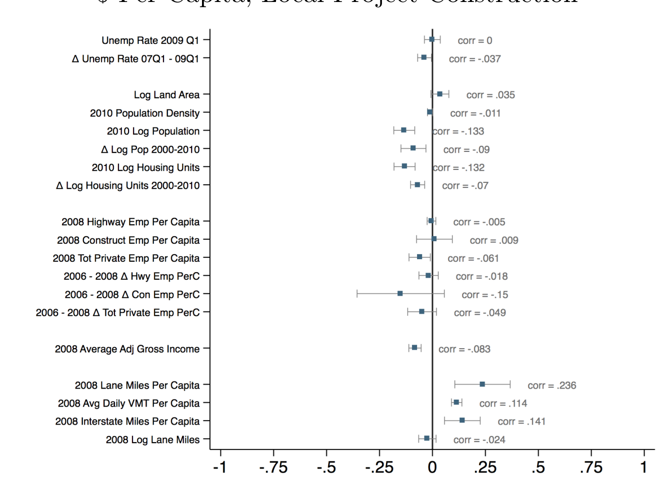

$ Per Capita, Local Firm Contracts

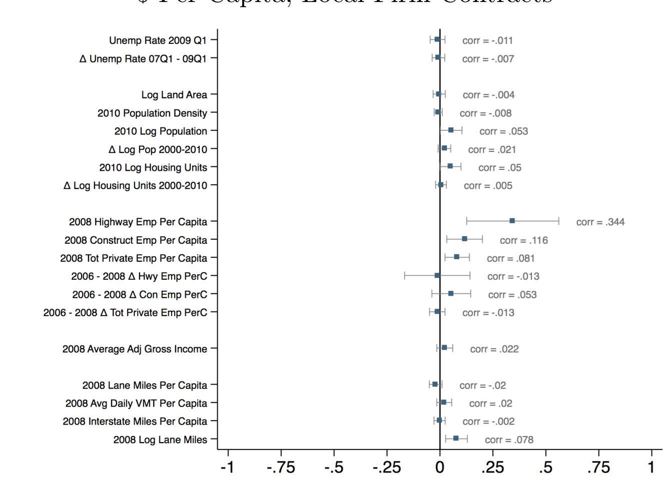

Notes: Point estimates are correlation coefficients, standard errors are clustered at the state level.

Panel B plots correlation coefficients of residuals from regressions of each variable on state fixed effects.

Figure 5: Event Studies of Local Construction Impacts

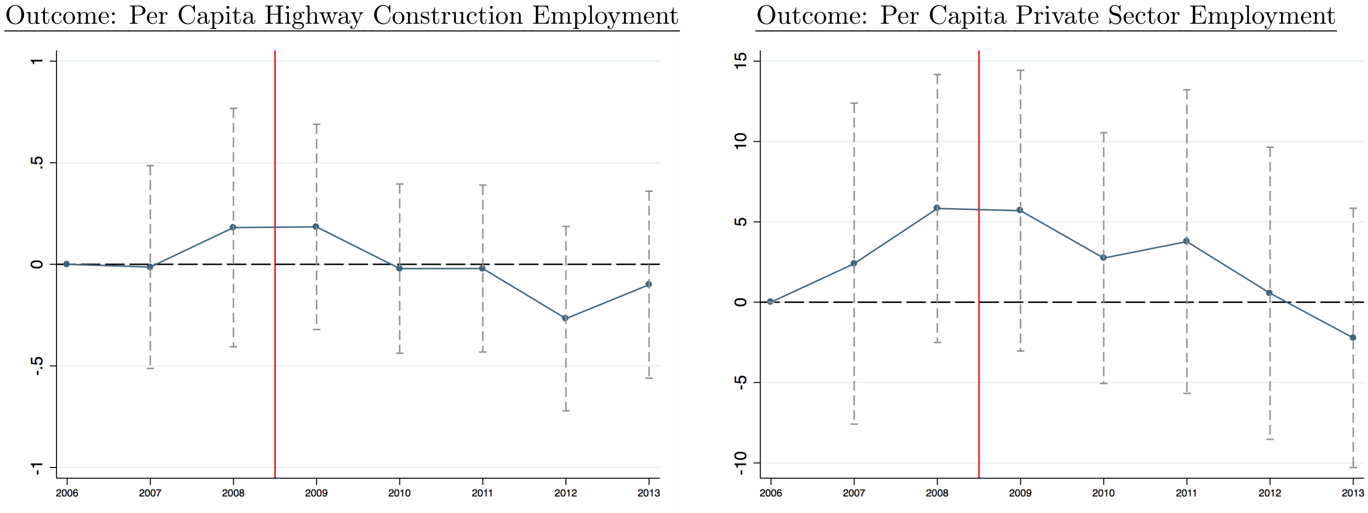

Notes: Figure displays county-level estimates of β^{τ} coefficients in equation (2). Specification of controls and fixed effects is the same as in Column (6) in Table II. See notes to Table II for further details.

Table I: Summary Statistics

Panel A: Coverage of Recovery Act Data

	Full Data	39-State Sample
N counties	3,023	2,494
Total Project Spending ($Mil)	24,060,000,000	18,360,000,000
Projects with Vendor Info ($Mil)	16,780,000,000	16,010,000,000
Total Vendor Receipts ($Mil)	15,430,000,000	14,610,000,000
Projects With Vendor Info:		
Number of Projects	7,436	7,162
Spending per Project ($)		
Mean	2,281,178	2,254,672
Median	668,536	667,892
SD	6,619,615	6,650,084
Vendors Per Projects		
Median	1	1
p90	3	3
p99	22	23

Panel B :County-Level Ohio, Analysis Sample Summary Satitistics

	Nonzero Values	Mean	Median	SD	98th pctile
Project Spending ($)	2,020	6,417,509	1,597,157	15,700,000	50,900,000
Per Capita ($)	2,020	165	45	599	1,216
Vendor Receipts ($)	1,170	5,858,132	0	21,800,000	67,000,000
Per Capita ($)	1,170	61	0	236	701

Notes: "Full Data" include data from 47 states, excluding Alaska, Hawaii, the District of Columbia, and other terriroties of the United States. Michigan is also excluded from these data, as most projects were erroneously reported as located in the State Capitol complex. "Projects" are subawards correspoinding to work conducted at a specified ZIP code. Vendors and contract payments are always associated with a subaward. "39-State Sample: excludes eight states with absent or sparse vendor reporting: Arizona, California, Illinois, Montana, North Dakota, Ohio, Virginia, and Wyoming. The "Primary Analysis Sample" drops all projects within the remaining states for which no vendor is listed; county-level aggregates are presented in Panel B.

Table II: Estimates of Effects of Local Construction Spending

Panel A: County Level Results

	1	2	3	4	5	6	7
	Outcome: Highway Constrution Employment (Per Capita)						
Per Capita Spending on Local Projects ($Mil)	-0.007	-0.052	-0.057	-0.102	-0.030	-0.090	0.508
s.e.	(0.187)	(0.232)	(0.230)	(0.226)	(0.217)	(0.202)	(1.562)
N	*12,470*	*12,470*	*12,470*	*12,470*	*11,705*	*11,705*	*11,705*
R-squared	*0.808*	*0.811*	*0.811*	*0.813*	*0.807*	*0.831*	*0.805*
	Outcome: Total Private Employment (Per Capita)						
Per Capita Spending on Local Projects ($Mil)	2.100	0.520	-0.545	-0.747	1.321	1.204	54.699
s.e.	(3.781)	(2.507)	(2.352)	(2.207)	(2.447)	(2.595)	(109.327)
N	*12,470*	*12,470*	*12,470*	*12,470*	*11,705*	*11,705*	*11,705*
R-squared	*0.978*	*0.979*	*0.979*	*0.980*	*0.981*	*0.982*	*0.977*

Panel B: Commuting Zone Level Results

	1	2	3	4	5	6	7
	Outcome: Highway Constrution Employment (Per Capita)						
Per Capita Spending on Local Projects ($Mil)	0.529**	0.592***	0.518**	0.361	0.368	0.285	-0.838
s.e.	(0.201)	(0.210)	(0.257)	(0.234)	(0.234)	(0.198)	(2.113)
N	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*	*2,995*
R-squared	*0.837*	*0.849*	*0.850*	*0.852*	*0.853*	*0.868*	*0.846*
	Outcome: Total Private Employment (Per Capita)						
Per Capita Spending on Local Projects ($Mil)	15.311	15.170**	12.943*	7.185	7.263	6.133	-38.505
s.e.	(11.617)	(7.169)	(7.433)	(7.195)	(7.368)	(5.061)	(134.819)
N	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*	*2,995*
R-squared	*0.957*	*0.964*	*0.965*	*0.969*	*0.969*	*0.974*	*0.962*
Clusters	39	39	39	39	39	39	39
Locality Fixed Effects	x	x	x	x	x	x	x
State x Year FE		x	x	x	x	x	
Population and Housing Controls			x	x	x	x	
Detailed Demographic Controls				x	x	x	
Highway Controls					x	x	
2003/2006 Hwy, Con, Priv Emp Controls						x	
Lane Miles IV							x

Notes:*** p<0.01, ** p<0.05, * p<0.1. Analysis sample is as defined in notes to Table I. Sample includes years 2007-2011. The displayed estimates are the OLS coefficients on the treatment intesity variable interacted with an indicator of whether the year is 2009 or after as in equation (1). All standard errors are clustered at the state level. Both the dependent and independent variables are scaled by 2010 residential population as measured in the Decennial Census. Employment outcomes are from the County Business Patterns data, and correspond to employment by employers based in the reference location. The dependent variable is winsorized at the 98th percentile. See Appendix for detailed information about control variable definitions and sources. All observations are weighted by log 2010 population.

Table III: Estimates of Effects of Local Construction Spending on Other Outcomes

	County - Level	**CZ - Level**	*Source*
Highway Employment	-0.030	0.368	CBP
	(0.217)	(0.234)	
Highway Wage Bill ($Thous)	-0.455	1.597	CBP
	(2.885)	(6.021)	
All Construction Employment	-0.746	-0.701	CBP
	(1.038)	(1.374)	
All Construction Employment	1.673	1.039	QCEW
	(1.086)	(1.481)	
All Construction Wage Bill ($Thous)	11.538	-25.432	CBP
	(21.610)	(46.952)	
All Construction Wage Bill ($Thous)	108.858	35.278	QCEW
	(79.481)	(82.028)	
Total Private Sector Employment	1.321	7.263	CBP
	(2.447)	(7.368)	
Total Private Sector Employment	1.452	-2.080	QCEW
	(1.866)	(3.627)	
Total Private Sector Wage Bill ($Thous)	33.828	118.957	CBP
	(75.994)	(206.731)	
Total Private Sector Wage Bill ($Thous)	-46.825	0.824	QCEW
	(112.800)	(168.214)	
Retail Employment	0.193	-0.991	CBP
	(0.599)	(2.071)	
Residential Construction Employment	-0.020	0.047	CBP
	(0.143)	(0.431)	
Unemployment Rate SDs (Residents)	0.057	-0.188	LAUS
	(0.355)	(0.195)	
Labor Force Participation Rate SDs (Residents)	0.504	0.133	LAUS
	(1.203)	(0.639)	
N	11,705	3,005	

Notes:*** $p<0.01$, ** $p<0.05$, * $p<0.1$. Analysis sample is as defined in notes to Table I. Treatment variable is Spending on Local Projects, in millions of dollars per capita. All outcome variables are in per capita units. The displayed estimates are the OLS coefficients of the outcome on the treatment intesity variable interacted with an indicator of whether the year is 2009 or after as in equation (1); each estimate from separate regressions. All regressions correspond to the specification in column (6) of Table II. All standard errors are clustered at the state level. Both the dependent and independent variables are scaled by 2010 residential population as measured in the Decennial Census, with the exception of the LAUS variables (unemployment and LFP rates) which are county level rates normalized by the sample standard deviation. All observations are weighted by log 2010 population.

Table IV: Estimates of Effects of Payments to Local Firms

	1	2	3	4	5	6
Panel A: County Level Results						
	Highway Constrution Employment (Per Capita)					
Per Capita Payments to Locally-Based Firms ($Mil)	-0.210	-0.280	-0.339	-0.206	-0.125	1.459
s.e.	(0.811)	(0.817)	(0.829)	(0.846)	(0.921)	(0.941)
N	*12,470*	*12,470*	*12,470*	*12,470*	*11,705*	*11,705*
R-squared	*0.808*	*0.811*	*0.811*	*0.813*	*0.807*	*0.832*
	Total Private Employment (Per Capita)					
Per Capita Payments to Locally-Based Firms ($Mil)	-8.230*	-6.980	-5.652	-4.172	-3.905	0.264
s.e.	(4.493)	(4.390)	(4.224)	(3.567)	(3.528)	(3.889)
N	*12,470*	*12,470*	*12,470*	*12,470*	*11,705*	*11,705*
R-squared	*0.978*	*0.979*	*0.979*	*0.980*	*0.981*	*0.982*
Panel B: Commuting Zone Level Results						
	1	*2*	*3*	*4*	*5*	*6*
	Highway Constrution Employment (Per Capita)					
Per Capita Payments to Locally-Based Firms ($Mil)	-0.005	-0.165	-0.222	-0.414	-0.397	0.855
s.e.	(0.791)	(0.958)	(0.958)	(0.970)	(0.966)	(0.946)
N	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*
R-squared	*0.837*	*0.848*	*0.850*	*0.852*	*0.853*	*0.868*
	Total Private Employment (Per Capita)					
Per Capita Payments to Locally-Based Firms ($Mil)	-4.383	2.762	-0.688	-2.071	-1.413	3.789
s.e.	(11.645)	(10.175)	(8.451)	(9.142)	(9.145)	(6.451)
N	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*	*3,005*
R-squared	*0.957*	*0.964*	*0.965*	*0.969*	*0.969*	*0.974*
Clusters	39	39	39	39	39	39
Locality Fixed Effects	x	x	x	x	x	x
State x Year FE		x	x	x	x	x
Population and Housing Controls			x	x	x	x
Detailed Demographic Controls				x	x	x
Highway Controls					x	x
2003/2006 Hwy, Con, Priv Emp Controls						x

Notes:*** p<0.01, ** p<0.05, * p<0.1. Analysis sample is as defined in notes to Table I. Sample includes years 2007-2011. The displayed estimates are the OLS coefficients on the treatment intesity variable interacted with an indicator of whether the year is 2009 or after as in equation (1). All standard errors are clustered at the state level. Both the dependent and independent variables are scaled by 2010 residential population as measured in the Decennial Census. Employment outcomes are from the County Business Patterns data, and correspond to employment by employers based in the reference location. The dependent variable is winsorized at the 98th perentile. See Appendix for detailed information about control variable definitions and sources. All observations are weighted by log 2010 population.

Appendix 1: Detailed List of Covariates

Variables	Source
Basic Population and Housing Controls	
2000/2010 Housing Units, Population, Density	Decennial Census (NHGIS)
2009 Q1 Unemployment Estimates	LAUS
Detailed Demographic and Income Controls	
2008 Per Capita AGI and Taxable Wages	IRS Statistics of Income County Series
2000 Industrial Employment by Gender as Share of Population	Decennial Census (NHGIS)
2005-2009 Race and Educational Attainment	2005-2009 ACS (NHGIS)
Highway Controls	
Per Capita Primary Road Lane Miles and Vehicle Miles Traveled; Interstate Miles	FHWA Highway Performance Management System
2003/2006 Hwy, Con, Priv Emp Controls	
Quartics in 2003 and 2006 Per Capita Highway, All Construction, and All Private Employment Per Capita	County Business Patterns

Appendix 2: Supplemental Figures

Appendix Figure 1: The Use of Recovery Act Highway Improvement Funds

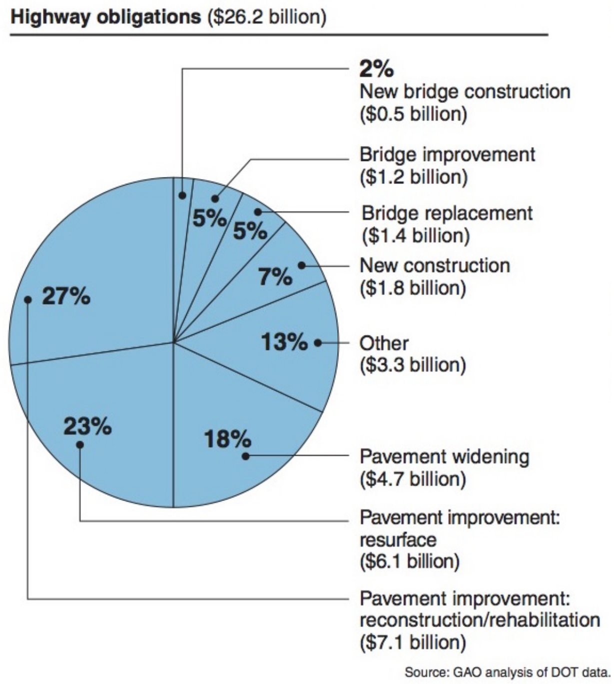

Source: CBO (2012)

Appendix Figure 2: FHWA Example of "Economically Distressed Areas"

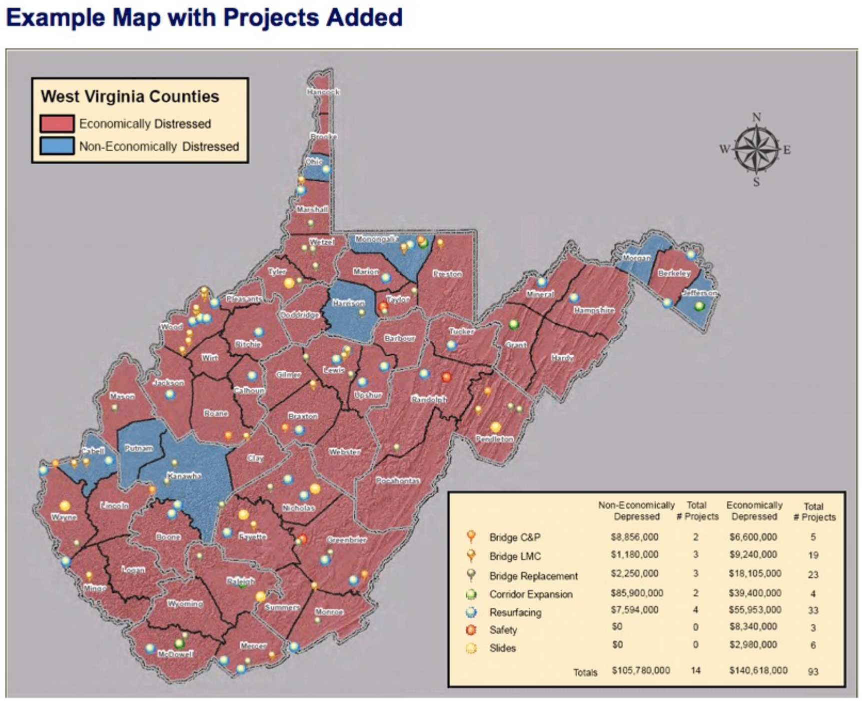

Source: FHWA

Appendix Figure 3: Geographic Distribution of Funds Relative to Population

A. Per Capita Spending on Local Projects

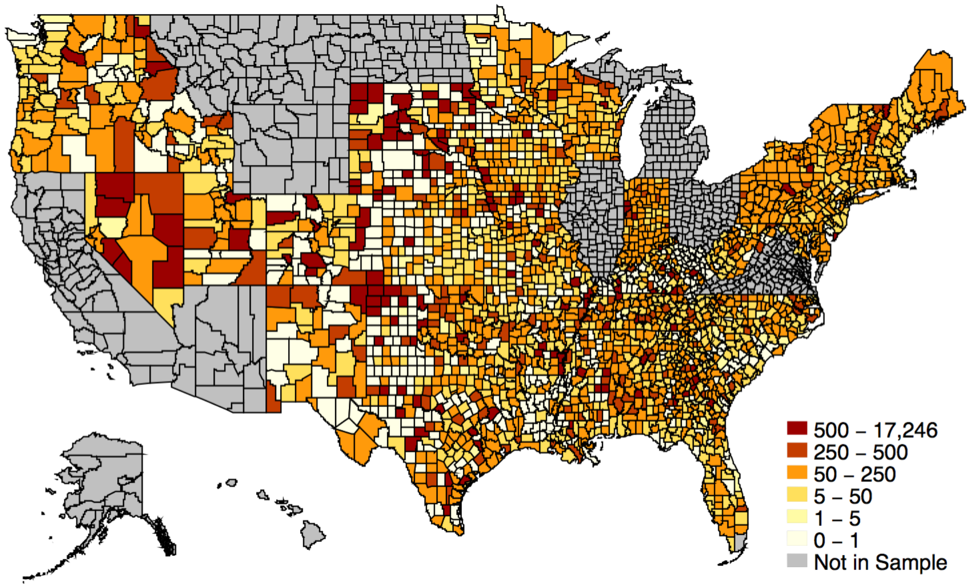

B. Per Capita Payments to Local Construction Firms

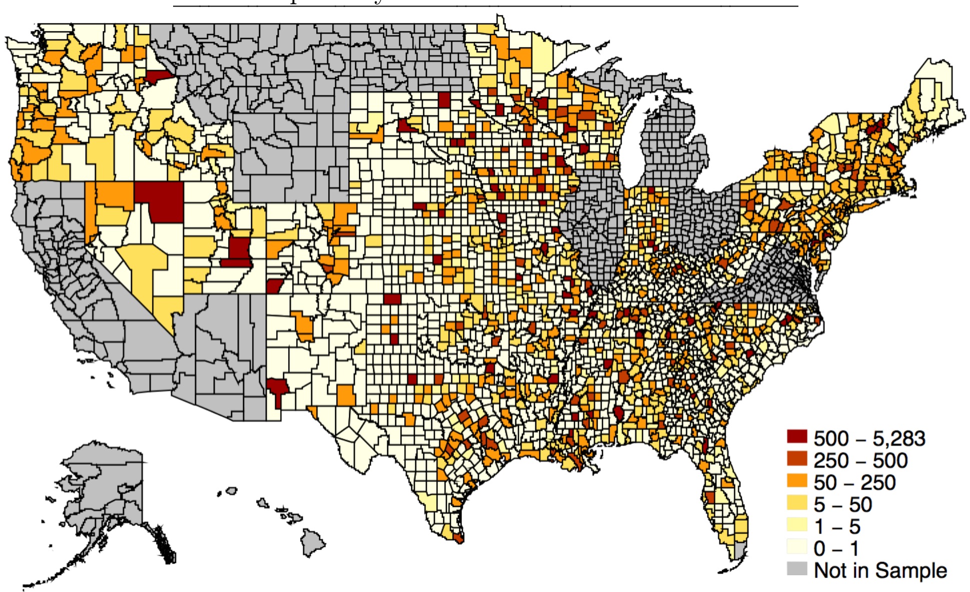

Appendix Figure 4: Event Studies of Payments to Local Vendors

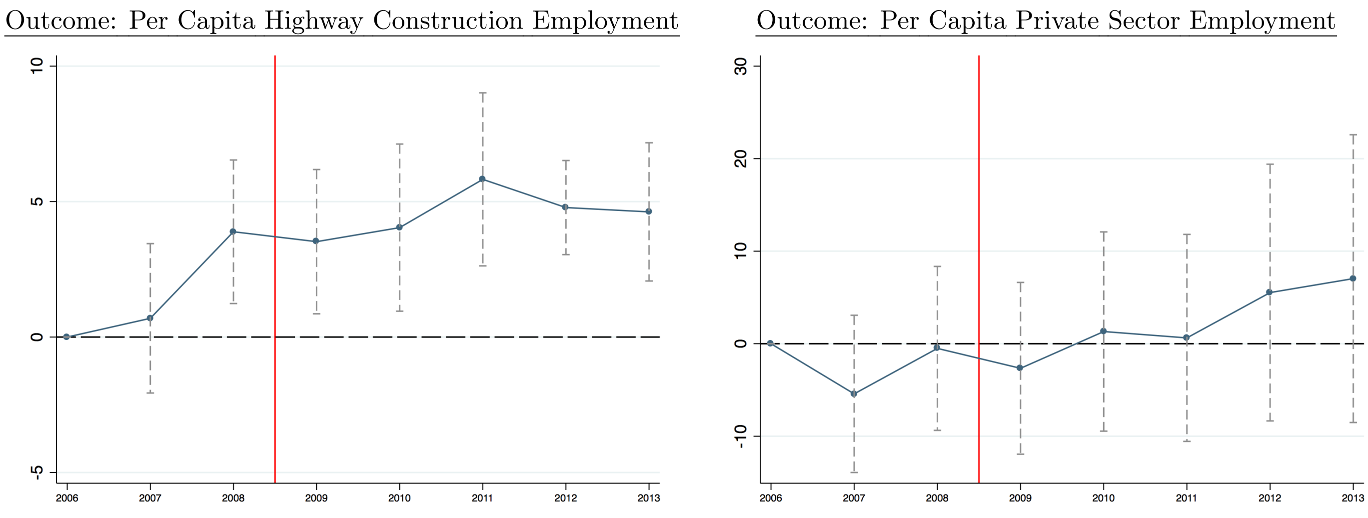

Notes: Figure displays county-level estimates of β^{τ} coefficients in equation (2). Specification of controls and fixed effects is the same as in Column (6) in Table IV. See notes to Table IV for further details.

Appendix Table I: Estimates of Effects of Local Construction Spending, Sample with All Project Data

Appendix Table I: Estimates of Effects of Local Construction Spending

	1	*2*	*3*	*4*	*5*	*6*	*7*
			Outcome: Highway Constrution Employment (Per Capita)				
Per Capita Spending on Local Projects ($Mil)	-0.043	-0.048	-0.055	-0.074	0.023	-0.042	0.713
s.e.	(0.143)	(0.154)	(0.154)	(0.144)	(0.135)	(0.117)	(1.140)
N	*15,115*	*15,115*	*15,115*	*15,115*	*14,185*	*14,185*	*14,185*
R-squared	*0.812*	*0.815*	*0.816*	*0.817*	*0.812*	*0.834*	*0.810*
			Outcome: Total Private Employment (Per Capita)				
Per Capita Spending on Local Projects ($Mil)	4.206	1.471	0.150	-0.571	1.542	1.640	1.471
s.e.	(3.251)	(2.095)	(1.980)	(1.951)	(2.271)	(2.459)	(1.859)
N	*15,115*	*15,115*	*15,115*	*15,115*	*14,185*	*14,185*	*15,115*
R-squared	*0.977*	*0.979*	*0.979*	*0.980*	*0.982*	*0.983*	*0.979*
Clusters	47	47	47	47	47	47	47
Locality Fixed Effects	x	x	x	x	x	x	x
State x Year FE		x	x	x	x	x	
Population and Housing Controls			x	x	x	x	
Detailed Demographic Controls				x	x	x	
Highway Controls					x	x	
2003/2006 Hwy, Con, Priv Emp Controls						x	
Lane Miles IV							x

Notes:*** p<0.01, ** p<0.05, * p<0.1. Analysis sample includes 47-state sample as described in Table I instead of the 37-state sample used in the primary analysis. Projects with no information on vendors are included in county-level totals. Otherwise, all notes in Table II apply.