ESSAYS ON THE VALUE OF LOCAL PUBLIC GOODS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Rhiannon Leigh Jerch

May 2019

ESSAYS ON THE VALUE OF LOCAL PUBLIC GOODS

Rhiannon Leigh Jerch, Ph.D.

Cornell University 2019

This dissertation consists of three chapters studying the effects of public goods and public infrastructure investments on urban growth and local government finances.

The first chapter estimates, first, how local governments finance federal mandates and, second, how much value local residents place on mandated local spending using a change in federal rules on municipal infrastructure following the 1972 Clean Water Act (CWA). I leverage the role of river networks in distributing pollutants across cities, combined with pre-CWA state regulatory intensity, to account for the endogeneity of municipal infrastructure adoption decisions, and to predict ex ante compliance with the CWA infrastructure mandate. Cities that were under the burden of compliance experienced substantial improvements to local ambient water quality as well as a three-fold increase in resident fees. Public spending on non-mandated items did not change, indicating that mandates are unlikely to displace local funding of other goods and services. The simultaneous increases to water quality and local costs resulted in taste-based sorting. However, I find that resident value of the mandated infrastructure depends upon the complementarity of surface water quality to pre-existing local features, as well as exposure to upstream polluters. These results imply that mandates may reduce inefficiencies to local public goods provision and provide positive benefits that are valued no less than their costs to local residents.

The second chapter, joint with Matthew E. Kahn and Shanjun Li, considers the

efficiency of local public service provision. A key challenge in quantifying the efficiency of the public sector stems from limited "apples to apples" comparisons of service functions offered by both public and private sectors, as well as the high correlation between local demand, demographic composition, and the local government's ability to deliver quality services. This paper posits a solution to this empirical challenge in two ways. First, we focus on public bus transit which is a relatively undifferentiated service across US municipalities. Second, we apply a regression discontinuity design using local mayoral elections as a source of random variation that predicts privatization levels in order to estimate causal effects of privatization on service efficiency. We find that privately operated firms provide bus transit at significantly lower costs per mile, largely due to their ability to circumvent public sector unions. We estimate that privatizing bus transit - a service used disproportionately by lower income groups - would lower the average bus fare by $1 per trip and create over 26,000 bus operator jobs nationally. However, these cost savings do not necessarily outweigh benefits of providing high-paying public sector employment opportunities.

The third chapter, joint with Panle Barwick, Shanjun Li, and Jing Wu, applies predictions of the Alonso-Muth-Mills model of urban land use to the context of Beijing's 2008 road rationing policy to identify how such policy instruments impact the spatial distribution of wealth within cities. We find that Beijing's rationing policy significantly increased the demand for housing near subway stations as well as central business districts. Further, we find the composition of individuals living proximate to subway stations as well as proximate to Beijing's central business districts shifted toward wealthier households. Our findings are consistent with theoretical predictions of the monocentric city model with income-stratified transit modes. These results provide suggestive evidence that city-wide road rationing policies can have the unintended consequence of limiting access to public transit for lower income individuals.

Dedicated to Janie and Richard Jerch

ACKNOWLEDGEMENTS

I would like to express my gratitude toward my dissertation committee, Professors Shanjun Li, Matt Kahn, Nancy Brooks, Mike Lovenheim, and Nick Sanders. I am indebted to each of you for your invaluable guidance and enthusiasm, and am immensely grateful to have learned from each of you. Prof. Li, I am particularly grateful for your unwavering patience and support over the past six years. I also owe thanks to Steve Coate, Fernando Ferreira, and Nina Harari who were particularly instrumental in the final stages of my dissertation.

Several of the most valuable lessons I learned during my PhD were taught to me by an exceptionally supportive set of peers and friends. Jianwei Xing, Caroline Walker, Angela Cools, Jennie Huang, Caitlin Gorback, and Collin Sullivan, this dissertation would not have been possible without each of your opinions, criticisms, and moral support. Jorgen Harris, I cannot express enough thanks to you. I am immensely grateful to have you as a colleague and friend.

I finally wish to thank my parents, Richard and Janie Jerch for their enthusiasm and love, and my best friend and husband, James Pikul, for his resolute faith in me. You inspire me every day.

Thank you.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

THE LOCAL CONSEQUENCES OF FEDERAL MANDATES: EVIDENCE FROM THE CLEAN WATER ACT

1.1 Introduction

Federal spending mandates are a controversial component of US fiscal federalism. In recent years, local governments have allocated over 13% of their annual expenditures to federally-mandated programs.[1] These include compliance with surface water pollution control under the Clean Water Act, improvements to drinking water quality under the Safe Drinking Water Act, lead-based paint abatement, and vehicle emissions control under the Clean Air Act (PriceWaterhouse 1994; Conlan 1994). Collectively, the average local expenditure share allocated to mandate compliance is equivalent to the average local expenditure share on public safety (US Census Bureau 2015). This figure has tripled since the early 1970's (Conlan 1994). How have these federal mandates impacted local governments? Opponents—both at state and local levels—argue that mandates infringe upon local sovereignty, inhibit the ability of cities to tailor their spending to preferences of local taxpayers, and place substantial cost burdens on local governments unable to flexibly raise revenues to meet mandated expenditure requirements.[2] Despite the largely negative popular opinion,

[1] Based on author's calculations using aggregate local cost needs reported in Conlan (1994) Table 2-2 and mean municipal annual expenditures sourced from US Census Bureau (2015). Conlan (1994) p.13 also cites national survey studies that find mandate compliance costs comprise between 11 and 12% of locally raised revenues.

[2] As recently as July 2018, members of Congress drafted a bill to limit federal mandates, citing its adverse impacts on local businesses and government budgets (Kasperowicz 2018). In 2017, the

empirical evidence on local budgetary responses to federal mandates is quite limited. Yet, understanding how local governments respond to federal mandates reveals important information on factors that govern their budgetary process, and ultimately, the effectiveness of federal policies that are implemented at the local level. In this article, I provide new empirical evidence on how cities fund compliance with federal mandates as well as the effectiveness of mandates at meeting their national goal using federal dictates on infrastructure spending following the 1972 Clean Water Act (CWA).

Is federally mandated spending necessarily mis-aligned with local taxpayer preferences? If local governments provide public goods at an efficient level, federal orders on local spending should, indeed, reduce the social surplus of existing residents (Samuelson 1954). However, there are several reasons why local public goods may be undersupplied relative to a locally efficient level. Large fixed costs, economies of scale, and credit constraints may prevent municipalities from investing in valued infrastructure projects (Fisher 2015). Additionally, failures of coordination, in which local provision decisions partially depend upon the decisions of other local governments, may lead to underprovision of public goods, particularly if those public goods generate spatial spillovers (Cooper and John 1988; Fisher 2015). This is a noteworthy concern with surface water quality, which was the focus of the 1972 CWA.

Mandated programs potentially address these inefficiencies. Several of the largest mandated programs, including the CWA and the Safe Drinking Water Act, are ac-

state of New York passed mandate-relief legislation in an effort to alleviate fiscal burdens on New York school districts (Seward 2017).

companied by federal grants and subsidies (Conlan 1994). If local residents are infra-marginal consumers of the mandated public good, federal mandates may provide local benefits to the extent that they relax fiscal pressure on, or credit constraints to, public goods provision. Additionally, the national scale and uniformity of these programs may induce strategic complementarities across local governments. For example, the marginal benefits to a city of abating surface water pollution may increase if all other cities sharing the same river also abate. To the extent that inefficiencies induce local governments to undersupply public goods with high fixed costs or diffuse benefits, federal mandates may generate substantial local benefits in the long run, even if they are politically unfavorable in the short run.

In this article, I further analyze the effects of mandate compliance on local government by asking: is local spending on mandated public goods valued above its costs to local residents? I measure the net value of mandated spending by assessing changes to house prices, population, and skill composition of local residents following compliance with federal requirements on infrastructure. These hedonic and sorting responses provide an indication as to whether individuals value mandated public goods provision more than the local cost.

Despite the growing prominence of federal mandates as a share of municipal budgets, very little is known about their consequences for the fiscal or economic well-being of local governments. A major impediment to conducting this type of research has, historically, been the lack of national-scale data on the local provision of a federally mandated public good. To remedy this problem, I obtained new data from the Environmental Protection Agency (EPA) on the census of municipal wastewater treatment plants—a public good which was federally mandated under the 1972 CWA. Importantly, I obtained these data from an early survey that predates enforcement of the CWA regulations, thus allowing me to estimate mandate compliance effects by comparing municipal outcomes before versus after federal enforcement of the mandate. The 1972 CWA aimed to improve the environmental health of US rivers and lakes by requiring that polluters of surface waters, including any municipal government operating a public sewerage system, treat their wastewater with a minimum level of pollution abatement technology. Cities with *ex ante* noncompliant wastewater treatment technology were under regulatory pressure to invest in more effective abatement technology following the 1972 legislation.

A second major challenge to understanding the local impacts of mandate compliance is that of causal identification. Specifically, local governments differ widely with respect to the preferences of their taxpayers and their ability to provide certain public goods, including wastewater treatment infrastructure. Such underlying differences across cities in their pre-policy provision of wastewater treatment infrastructure are likely deterministic of differences across these cities in their local fiscal conditions and ability to attract taxpayers.

To solve this endogeneity problem, I construct an instrument that predicts wastewater treatment adoption using city-level variation in riparian exposure to downstream populations and state-level variation in pre-policy water pollution regulation. The intuition behind my identification strategy is that cities with historically large population centers downstream were more likely to be pressured by their downstream neighbors to adopt stringent wastewater treatment, long before the CWA became legislation, in order to reduce conveyance of harmful pollutants to downstream drinking water sources. Furthermore, this inter-jurisdictional pressure was more likely to be enforced for cities within states with more regulation of surface water pollution. Such early-adopter cities were unaffected by the CWA infrastructure standard when the law passed in 1972 because they had already adopted secondary treatment technology. By leveraging variation in infrastructure adoption driven by forces external to the city, this instrument provides variation in *ex ante* CWA compliance that is plausibly exogenous to local spending decisions or growth.

In contrast to the main popular criticism against federal mandates, I do not find evidence that cities were forced to displace funding of other goods and services in order to fund compliance with the CWA mandate. While the infrastructure requirements caused local governments to more than double their expenditures on wastewater from 6% prior to the Act to over 14% of their total budgets, cities primarily funded these compulsory expenditures through federal grants and by tripling fees on residents. I find that the mandated expenditures on wastewater treatment led to economically and statistically significant improvements in water quality: cities under the burden of compliance with the CWA experienced a 18% improvement to surface water quality, on average, as measured by dissolved oxygen concentration. My results are consistent with recent work by Keiser and Shapiro (2018) who find significant positive effects of CWA federal grants on dissolved oxygen levels. However, this study advances prior work concerned with the CWA by providing a new estimate on the effectiveness of the CWA technology standard—the primary regulatory instrument of the 1972 legislation—at meeting the Act's main goal of improving surface water quality. The findings of this paper demonstrate that the CWA minimum technology standard for wastewater treatment was effective at improving the environmental health of rivers and lakes across the US. More broadly, these findings suggest that the local implementation of a federal pollution control policy was effective at improving US surface waters.

When considering the average city, I find that local improvements to surface water quality had positive impacts on population growth and housing prices, suggesting that the mandated infrastructure was at least valued at its marginal cost to local residents. However, the aggregate effects mask important sources of heterogeneity in city responses. First, per capita compliance costs were 20% higher among smaller cities unable to exploit scale economies in infrastructure realized by larger cities. Second, I show that the value of mandate compliance to local residents was greatest among cities with warmer summer climates, closer proximity to large waterbodies, and greater exposure to upstream abatement. I interpret these findings as demonstrating that federal mandates can correct for inefficiencies to local public goods provision in the presence of market failures. I further interpret these findings as suggestive that both efficiency and equity could have improved under the CWA federal aid program if grant allocation followed city-specific abilities to benefit from scale economies as well as improved surface water quality.

Prior work concerned with the local impacts of federal spending mandates is quite limited. Miceli and Segerson (1999) develop a theoretic framework to compare the relative efficiency of fully, partially, or unfunded federal mandates. Crémer and Palfrey (2000) also derive a theoretical framework to understand welfare gains from various federalist systems of public goods provision, but does not consider welfare gains in the case when local decisions on public goods provision are strategically linked across jurisdictions - a scenario that is highly relevant to local water pollution abatement. Empirical work on federal mandates mainly rely on isolated case studies (Weiland 1998 and Hanford and Sokolow 1987) or focus on education outcomes following No Child Left Behind (Imazeki and Reschovsky 2004; Reback et al. 2014; Deming et al. 2016) without consideration for local government budgetary responses to mandate compliance.[3] Work by Katherine Baicker (Baicker 2001 and Baicker and Gordon 2006) provides the closest parallel to my study. She explores state-level budgetary responses to federally mandated changes to healthcare spending and finds significant evidence of crowd out in welfare payments. One plausible reason my findings of no crowd-out contrast from hers is that wastewater treatment has few substitutes. A city cannot readily reduce surface water pollution through means other than by adopting treatment technology. Medicaid spending, in contrast, may be a substitute service for welfare payments, since both mainly serve health and quality of life among lower-income populations.

My paper makes three important contributions to the public finance and envi-

[3]National survey studies by Conlan (1994), NLC (2017), EPA (1988) and Lake et al. (1979) provide evidence that mandates displace local funding of other public goods, however these studies are descriptive and based on self-reported responses to surveys explicitly concerned with federal mandates, which may invite biased responses from local officials.

ronmental economics literature: First, I provide the first test of local fiscal responses to federal mandates. Specifically, I show that the unfunded local costs of CWA mandate compliance were financed through increasing revenues, rather than re-allocating spending away from other public goods and services. This finding contrasts with most existing work on local budgetary responses following shocks to taxpayer wealth (Skidmore and Scorsone 2011; Lutz et al. 2011; Lutz 2008; Alm et al. 2011; Feler and Senses 2017; Melnik 2017; Cromwell et al. 2015) because, unlike these papers, my results demonstrate that cities *do not* respond to a budget shock through austerity measures. However, this asymmetry is consistent with the "flypaper effect" summarized by Hines and Thaler (1995), which suggests governments respond to taxpayer wealth shocks differently than to proportional increases in expenditure obligations. This finding has important policy implications because it demonstrates that federal mandates are unlikely to displace funding of other local goods and services. In other words, the current focus within congressional mandate reform on local fiscal burdens may be overly narrow.

Second, I provide evidence that partially funded federal mandates can reduce inefficiencies to the provision of public goods that are valued by local residents. While popular opinion generally agrees that federal mandates serve to further interests specific to a national agenda, I argue that mandate compliance can provide important benefits locally if mandates correct for certain market failures that otherwise restrict the local provision of public goods.

Lastly, my paper contributes to a rich literature concerned with estimating the value of policy-induced changes to environmental amenities (Keiser and Shapiro 2018; Banzhaf and Walsh 2008; Kahn 2000; Chay and Greenstone 2005; Greenstone and Gallagher 2008; Gamper-Rabindran et al. 2011; Sieg et al. 2004). To the best of my knowledge, this paper is the first to study the direct regulatory impacts of the CWA technology standard. More broadly, my paper demonstrates that environmental regulation enforced through a federal mandate can impact local public finances and, in particular, can increase the cost of living for local residents. This finding is important because several national environmental policies, including the Clean Air Act, the Safe Drinking Water Act, and the Superfund Amendments—like the CWA—require local government expenditures for their implementation (Conlan 1994). Consideration of the local public finance element of these regulations is of first order concern because, if environmental improvements from federal regulation spur increases to local taxes or fees, hedonic estimates may systematically underestimate true willingness to pay for environmental amenities.

The remainder of my paper proceeds as follows. Section 1.2 provides a brief overview of the theoretical framework motivating my empirical estimation, Section 1.3 provides institutional background on the CWA regulations and determinants of wastewater treatment technology adoption, Section 1.4 describes the data, Section 1.5 presents the two empirical approaches I use to identify the effects of the CWA infrastructure mandate, and Section 1.6 discusses my results and provides several robustness checks. In Section 1.7, I explore heterogeneity in municipal responses to the CWA. Finally, Section 1.8 concludes.

1.2 Theoretical Framework

In this section, I discuss the basic theoretic framework motivating my empirical research design. The experience of Walton, NY provides context for the theoretic motivation.

Walton is a village of 3,700 people situated along the west branch of the Delaware River. Following the CWA, Walton was mandated to build a $9 million wastewater treatment plant (Newman 1976). Their annual operating budget in that same year was $556,000 ($3.07 million in 2012 dollars). While Walton received large federal subsidies to comply with the CWA infrastructure requirement, the challenge of meeting the local cost share as well as securing federal aid had several implications for the city's finances. For the first time in its 300 year history, Walton issued debt, charged sewerage user fees on its residents, and applied for intergovernmental grants. Did residents of Walton value this new wastewater treatment more than its costs?

Figure 1.3 illustrates that the CWA federal mandate on wastewater treatment infrastructure can have two opposing effects that depend upon the initial efficiency of local public goods provision. If pre-CWA Walton is best represented by the left panel, city A, water pollution abatement in equilibrium was under-provided due to some market failure such as credit constraints or coordination failures stemming from upstream pollution spillovers. Note, that while existence of pollution externalities can induce under-provision of pollution abatement relative to an aggregate, social efficiency perspective, credit constraints and coordination failures can induce underprovision from a *local* efficiency perspective. These market failures generate a wedge, indicated by X, between the locally efficient level of P^* and the equilibrium level, P^{Low}. For city A, the mandate serves to, both, reduce credit constraints by providing federal matching grants and reduce coordination failures because it requires that *all* cities on the same river equally abate their pollution. The former reduces marginal costs while the later increases the local marginal benefit of abatement. Consequently, the mandate increases local surplus of living in city A.

On the other hand, if pre-CWA Walton is best represented by city B, residents experience the mandate as strictly reducing local surplus because the city was already providing pollution abatement at its locally efficient level. The mandate forces this city to abate at a level beyond which additional benefits outweigh additional costs.

Following Oates (1969), I consider a utility maximizing consumer who weighs the benefits stemming from the menu of local public services against the cost of their tax liability and chooses as a residence the location that provides them with the greatest surplus of benefits over costs. The individual's tax liability is their effective "price" of consuming the local output of public services. After Walton increases its tax liability in order to expand its output of public services (i.e., pollution abatement through wastewater treatment), property values need not decline and may actually increase if residents value the additional public good provision more than its cost. Mandated expansion of the public good will increase property values in city A of Figure 1.3 because the mandate will serve to increase local social surplus. In addition to property values, changes to population provide evidence of a positive revealed preference for changes in public goods, as individuals "vote with their feet" (Tiebout 1956; Banzhaf and Walsh 2008).

At least 10,000 other villages, towns, and cities like Walton were also in violation of national restrictions on dumping untreated sewage and waste into surface waters following the CWA (EPA 1973). Whether city A or city B is more representative of the *average* US city prior to the CWA is an empirical question. Consequently, I estimate the effect of the CWA mandate on local taxes and fees, population, and housing prices to test whether, on average, the mandate altered local tax liabilities and whether it increased or a decreased efficiency of public goods provision.

My empirical approach exploits the 1972 CWA policy change in order to isolate the mandate's effects on these outcomes net of any remaining confounding factors. Let i index cities, and t index years, I estimate:

$$y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\Gamma_t + \varepsilon_{it} \ , \tag{1.1}$$

where y_{it} is one of several outcomes of interest, including local fees, population, and housing prices; P_i is an indicator equal to 1 if a city is under the burden of CWA mandate compliance *ex ante*; and $POST_t$ is an indicator equal to 1 for years after the CWA came into effect. β provides the differential change in y for a noncompliant relative to a compliant city as a result of the CWA mandate. If the mandate increased local social surplus and increased local efficiency of wastewater treatment provision then β will be positive. In the following sections, I discuss determinants of wastewater treatment technology adoption and potential sources of endogeneity in *ex ante* noncompliance.

1.3 Regulation of Surface Water Pollution

1.3.1 Benefits & Costs of Wastewater Treatment

Wastewater treatment facilities protect environmental and public health by treating sewage, urban debris, and pathogens from piped waters before they return to rivers and lakes. By removing oxygen-consuming organic matter that damages aquatic ecosystems, wastewater treatment helps the environment as well as the aesthetic and recreational use value of surface waters.[4] Specifically, wastewater treatment can improve surface water clarity be reducing instances of algal bloom eutrophication and aids in improving aquatic biodiversity (Brown 2018). Prior work on water quality valuation demonstrates that consumers value improvements to recreational fishing, swimming, boating, and surface water clarity (Olmstead and Kuwayama 2015; Bockstael et al. 1987, Lipton 2004; Boyle et al. 1999). Further, waterfront (Leggett and Bockstael 2000) as well as non-waterfront property values (Walsh et al. 2011, Poor et al. 2007) increase following local surface water pollution control, suggesting that the benefits to local water quality extend beyond properties immediately adjacent to the affected water body. Consequently, improvements to surface water following wastewater treatment has the potential to increase surrounding property values or population levels if individuals value the water quality amenity more than its

[4] Wastewater treatment also increases the supply of potable water and can help to prevent disease by removing harmful bacteria and chemicals. In industrialized economies, however, health benefits from surface water pollution control are likely to be dwarfed by recreational and ecosystem benefits because basic drinking water treatment methods are ubiquitous and have a long history, predating most federal environmental regulations (Olmstead 2010).

marginal cost.

However, constructing and maintaining wastewater treatment facilities is costly and requires significant public financial investment. US municipalities allocate 10% of their total annual expenditures toward sewerage and wastewater treatment (US Census Bureau 2015). Costs vary considerably with the type and rigor of wastewater treatment technology. Primary treatment is a basic treatment process that utilizes physical methods (gravity, settling tanks, or centrifuges) to separate waste from water. Secondary treatment is a more advanced technology that uses biological processes to decompose the organic matter in waste that can both spread disease and absorb oxygen in water. Secondary treatment removes more than twice as much oxygen demand from wastewater as primary treatment and is thus more effective at protecting aquatic life and reducing bacterial counts in surface water (Stoddard et al. 2003). However, secondary treatment is considerably more expensive to install and operate, ranging between 2 and 10 times the cost of primary treatment (EPA 1976). Figure 1.4 plots the engineering costs required for secondary treatment technology based on a plant's service population. For a city of 30,000 people, upgrade costs were roughly $6 million in 2012 dollars, equal to the annual public safety operating budget of a similar-sized city.[5]

[5]Per Guo et al. (2014), the average wastewater flow per capita is 100 gallons per day. Public safety cost estimate calculated by averaging Census of Governments data on annual police and fire expenditures for cities with populations between 20,000 and 40,000.

1.3.2 The 1972 Clean Water Act

The large investment costs and potentially diffuse benefits of secondary wastewater treatment contributed to the need for federal regulation of surface water pollution. Prior to the CWA, over two-thirds of municipal wastewater systems used the primitive, less expensive treatment technology of primary treatment (see Figure 1.1). The CWA addressed this low take-up of rigorous wastewater treatment by establishing secondary treatment as the minimum technology standard for all wastewater re-entering surface waters. For this reason, the Act constitutes a "technology-forcing statute" because of its focus on pollution abatement treatment (Copeland 1999).[6] By directly removing bacteria and thus lowering oxygen demand in wastewater effluent, secondary treatment provided the means to achieve the CWA's ultimate goal of making all US surface waters "fishable and swimmable."

Congress enforced the secondary technology standard through a new monitoring system. Local governments, firms, or individuals dumping untreated wastewater into surface waters through any discrete conveyance could be fined up to $25,000 per day, sanctioned, sued, or imprisoned by the federal government (Copeland 1999).[7] The CWA also recognized the authority of citizens to bring civil suits against their local governments for violating CWA standards (Andreen 2013).[8]

[6] The original 1972 Act included requirements under Section 303(d) for states to monitor surface water pollution levels and abate when pollution levels exceed state limits. The EPA did not begin enforcing the "control" approach of the CWA until 1992 (Copeland 2012).

[7] For examples of recent CWA enforcement, see Fields and Emshwiller (2011) or Westerling (2011).

[8] Earnhart (2004) analyzes the impact of different regulatory instruments (permits, inspections, and enforcement actions) on water pollution abatement among municipal plants in Kansas in the

To assist in paying for these large infrastructure costs, the federal government distributed construction grants to state and municipal governments. Approximately two-thirds of existing wastewater treatment plants received at least some funding (Keiser and Shapiro 2018). At their peak, these grants were intended to support up to 75% of total capital costs. However, by 1981, Congress changed the matching rate to 55%, and by 1987, the grant program was phased out.

While a substantial portion of these construction costs for CWA compliance were supported by federal grants, there are several reasons to expect that the unfunded costs placed a significant burden on local budgets. First, municipalities were obligated to fund at least 25% of their capital needs for secondary treatment, equivalent to approximately 4% of the average noncompliant city's budget prior to the CWA. Second, operating costs, which on average are 60% of total annual wastewater treatment costs (US Census Bureau 2015), were not eligible for grant assistance and operations costs are likely to increase with secondary treatment.[9] Lastly, while no study has comprehensively tested how compliance with the CWA technology standard impacted local finances, early case studies found that some communities were unable to provide the required finances for infrastructure adoption without burdening taxpayers or displacing other services (EPA 1973; GAO 1980; PriceWaterhouse 1994; Hanford and Sokolow 1987; Lake et al. 1979).

mid 1990s.

[9]Secondary treatment requires more energy input relative to primary treatment to operate aeration pumps and added personnel to monitor electrical and mechanical processes. Also, monitoring of the secondary treatment biological digestion process often requires skilled labor from environmental and civil engineers, unlike primary treatment operations (Brown 2018).

Figures 1.1 and 1.2 show aggregate, national-level effects of the CWA technology standard. Figure 1.1 shows that between 1972 and 1977, the number of treatment plants with noncompliant "primary" technology fell by 50%, and steadily declined thereafter. The contemporaneous changes to municipal budgets are apparent in Figure 1.2: while the share of total spending in wastewater trended with overall spending in years 1967 and 1972, there is a divergence after the CWA in which the share of wastewater spending increased by 3 percentage points on average, despite a downward trend in total spending per capita.

The 1972 CWA marked a major shift in the nation's approach to surface water regulation. Prior to the CWA, state governments had *de facto* autonomy over their surface water regulations. In contrast, the 1972 CWA gave the federal government substantial power to respond directly to violations of the Act through administrative actions, civil actions, and criminal sanctions (Andreen 2013).[10] The prior federal approach to surface water regulation was more passive, relying on voluntary subsidies to local governments in order to promote national programs (Dilger 2013). Additionally, disputes between the executive and legislative branch nearly handicapped the enforcement of the Act.

The unique providence of the CWA legislation is important for this paper's empirical design. I leverage the unanticipated nature of the CWA to compare changes

[10]Of the six major federal Acts between 1948 and 1970 related to surface waters that preceded the 1972 CWA, all maintained that regulatory authority was mainly in the hands of the states (Fairfax and Hamilton 2000). Two Acts in 1966 and 1970 had attempted to impose wastewater treatment standards, but these actions were legally challenged by the States, and were never enforced (Stoddard et al. 2003).

in outcomes after the Act across *ex ante* compliant and noncompliant cities. Identification of causal estimates under this approach would be invalidated if cities had been able to anticipate the CWA regulations or if broader changes to wastewater treatment were taking hold prior to the CWA. However, for cities to have anticipated the CWA regulations, they would need to have foreseen a substantial deviation from historical precedent on state rights to self-regulate and strong collective action on the part of the legislative and judicial branches to counteract a presidential veto.

1.3.3 History of Wastewater Treatment

Prior to the 1972 CWA, litigious downstream neighbors suffering from pollution externalities were one potential mechanism inducing cities to adopt secondary treatment technologies. In the early development of modern water infrastructure, methods of wastewater treatment were mainly for aesthetic as opposed to direct health benefits. Urban water systems were designed such that drinking water intakes were upstream of wastewater outfall locations and protected from wastewater contaminants (Okun 1996). Primary treatment could provide a localized improvement in surface water quality by reducing accumulation of solid materials, debris, and pervasive odors, but wastewater treatment was not considered necessary for cleaning a city's drinking water supply (Tarr 2016).[11] Any health benefits from sewage treat-

[11] In reference to primary treatment development, Stoddard et al. (2003) states: "In many cases, this construction was promoted by city officials and entrepreneurs, who were rapidly learning that unsightly urban debris and a delightful growing phenomenon, tourists with leisure dollars to spend, did not mix."

ment were perceived by early city planners to accrue mainly to downstream neighbors (Metcalf and Eddy 1922).

Urban population growth during the twentieth century combined with improved understanding of riparian biochemical cycles in the scientific community brought about water pollution disputes between upper and lower riparian cities. Development of the metric "biochemical oxygen demand" provided a method to directly link the existence of organic wastes in wastewater to bacterial and oxygen levels in natural waters (Fairfax and Hamilton 2000; Melosi 2000), enabling downstream cities to pinpoint sources of pollution in their own drinking and surface waters. Consequently, pressure from downstream cities could induce upstream polluters to adopt secondary treatment technology (Melosi 2000). Various historic anecdotes illustrate this pattern. For example, the city of Chicago—which shares the Mississippi basin with St. Louis—invested in secondary treatment technology beginning in 1916, only after the state of Missouri enacted a lawsuit against Chicago in 1901 for polluting the drinking water of St. Louis (Cain 2005; Stoddard et al. 2003; Missouri 1901). Similarly, Melosi (2000) discusses how city adoption of wastewater treatment methods were "born amid the unhealthy background of injunctions and court orders" between cities.[12] Several state and federal court cases at the turn of the century set the precedent for individual protections against water pollution. These cases provided the legal framework for individuals and local governments to prevent both

[12]Melosi (2000) also states: "Conflict between upstream and downstream cities over the dumping of sewage and industrial waste had been fought in the courts and addressed through interstate sanitation compacts."

public and private entities from polluting surface waters.[13]

In summary, a city's riparian exposure to downstream populations may have a positive impact on the likelihood of *ex ante* compliance with the CWA technology standard. Prior to the 1972 CWA, cities demanding aesthetic improvements to nearby surface waters generally adopted the low-cost primary technology. More costly secondary treatment adoption, on the other hand, could follow litigation disputes from downstream neighbors suffering from pollution externalities. Such early-adopter cities were unaffected by the CWA infrastructure standard when the law passed in 1972 because they already were compliant. My identification strategy exploits these determinants of infrastructure adoption to predict *ex ante* CWA compliance. By leveraging forces external to the city dictated by river networks, this instrument provides variation in *ex ante* compliance status that is plausibly exogenous to local spending decisions or growth.

[13]Example cases include: Storley v. Armour & Co., 107 F.2d 499 (1939), Sammons v. City of Gloversville, 67 N.E. 622 (NY 1903); Butler v. White Plains, 69 N.Y.S. 193 (1901); Gould v. City of Rochester, 105 N.Y. 46 (1887). McQuillin (1912) provides a thorough law review of municipal responsibilities for water pollution control. These cases set the precedent that "a city has no right to gather its sewage and cast it into a stream so as to injure the lower proprietor" and further, the "power of a municipal corporation to construct sewers or to use a natural stream as a sewer does not authorize it to so construct the sewers or to use the stream as to create a nuisance to the damage of a lower riparian owner. (p. 3051)"

1.4 Data

Ninety percent of publicly owned wastewater treatment plants in the US are financed, operated, and managed by local governments.[14] Consequently, I consider the local government, including cities, villages, boroughs, towns, and townships as my primary unit of observation.[15] Throughout this paper, I refer to "local governments," "municipalities," and "cities" interchangeably. Local governments in the US are small. Eighty percent of all municipalities have populations less than 20,000 and a mean population of about 5,000 people (see Appendix Figure A.1). The analysis in this paper, therefore, is distinct from much prior literature concerned with urban sorting responses in its focus on the representative US municipality as opposed to metropolitan urban centers.

I construct a new dataset of municipal CWA compliance using records obtained from the Clean Watershed Needs Survey (CWNS) Team of the EPA. The CWNS is a census of approximately 22,500 publicly-owned wastewater treatment facilities. These data provide detailed facility-level information including unique facility identifier codes, treatment technology characteristics, operating status, and identifying information on the facility's managing authority including name, county, state, and

[14]Roughly 11% of the 22,500 publicly-owned plants in the US are managed by counties, states, or other non-municipal authorities such as universities, national parks, or correction facilities (EPA Clean Watershed Needs Survey, 1973-2004).

[15]Local governments are defined by the US Census Bureau as political entities authorized by state constitution to provide government for a specific population in a defined area. The geographic boundaries of local governments are endogenously determined and can vary over time. I standardize the political jurisdiction of local governments over time through FIPS place codes rather than geographic boundaries because jurisdictional borders may respond to fiscal shocks.

government type. The CWNS surveys began in 1972 and have since been administered approximately biannually.[16]

The treatment technology variables of the 1972 survey provide the crucial information I use to observe compliance status of a city's wastewater treatment plant before the CWA regulations came into effect. While the EPA has not maintained data dictionary records associated with the 1972 computer-readable survey file, I was able to identify the survey's compliance status information using a copy of the original 1972 survey questionnaire, found in the appendix of EPA (1973). I define a plant as *ex ante* noncompliant if its effluent discharge is recorded as not meeting secondary treatment levels at the time of the survey. A plant is *ex ante* compliant if its effluent discharge is recorded as meeting secondary or more stringent treatment levels. This is the first study to codify these early CWNS surveys in order to categorize and document impacts of municipal CWA compliance.

While I use only the 1972 survey to designate municipal treatment compliance status, I utilize the full survey panel to construct a sample suitable for analyzing the CWA minimum technology standard. Appendix A.2 provides additional detail on my sample restrictions.

Figure 1.1 shows the national growth in plants with secondary treatment technology following the CWA. As of 1972, over three-quarters of publicly owned wastewater treatment plants lacked the secondary treatment technology mandated by the CWA.

[16]Data reported in each CWNS report are representative of the prior calendar year. For example, the first available CWNS survey titled the "1973 Clean Watershed Needs Survey" describes plant technology as of 1972.

Between 1975 and 1977, the number of plants with only primary treatment fell by over 55%, from roughly 10,000 plants to under 4,500, and declined steadily thereafter. Compliance with the CWA's technology standard had a substantial time lag. Part of this lag is mechanical: primary to secondary treatment upgrades require several years for engineers to execute planning and construction, as well as for the municipality to secure financing and apply for federal aid. Another part of this lag is political: the CWA passed amid substantial backlash from the executive branch. After a Congressional override on Nixon's initial veto of the Act, Nixon impounded half of the funding Congress had appropriated for plant construction costs. It was not until after 1975, when the Supreme Court ruled against presidential power to impound funds (Train 1975), that appropriations for the CWA ramped up and several cities received construction grants (Copeland 2015). These lags justify assigning 1972 as a pre-policy year even though this was the year that the Act became law.

To gauge the impact of the CWA technology standard on municipal finances, I use the US Census Bureau's "Historical Finances of Individual Governments" database. These data provide detailed information on annual revenues, expenditures, and debt for the census of local governments every five years, starting from 1967. I merge the municipal finance data with the CWNS plant technology data based on the name, state, county, and government type (i.e., "city", "village", "township", or "borough") of the plant's managing authority. To ensure the accuracy of this merge, I exclude plants managed by counties, districts, universities, or corrections facilities. Additionally, I exclude cities with non-unique name–government type combinations within their county. Under this criteria, I am able to match 3,593 municipalities to the Census finance data from approximately 4,000 municipalities in the CWNS data. Because the Census finance data are self-reported, they may be prone to measurement error, particularly if local government expense and revenue categories do not precisely match up with the Census of Governments categories. For this reason, I consider impacts on categories directly related to the CWA mandate (i.e. wastewater user fees), as well as aggregates of those categories (ie, total user fees). I also exclude municipalities reporting zero total expenditures or property taxes, which eliminates approximately 6% of the municipalities. Lastly, I restrict the sample to cities that appear in each decade of the "Historical Finances of Individual Governments" database, which eliminates approximately 5% of municipalities. These data restrictions yield a sample of 2,975 cities.

Data on growth outcomes, including population, education levels, and median housing prices, are sourced from the Decennial Census. IPUMS NHGIS provides these data at the relevant FIPS place and county subdivision levels from 1970. IPUMS also provides shapefiles, which I use to calculate the centroid of each FIPS place and county subdivision in GIS. Information on local labor markets and industrial composition are sourced from County Business Patterns, available from 1956, and annually from 1974. I gauge pre-CWA support for environmental issues among state senators using the League of Conservation Voters score card from 1971 and 1972, and I source distance-to-major- waterbodies using data from Rappaport and Sachs (2003). Lastly, I source water quality data back to 1962 from the EPA STORET Legacy database, as well as the National Water Information System *Water Quality Portal*. These data provide water quality readings from over 740,000 monitoring locations across the US. I calculate ambient water quality as the annual average dissolved oxygen level within 25 miles of a city centroid, where monitor readings are inversely weighted by their distance from the city centroid.[17] I focus on dissolved oxygen as my preferred measure of water quality because it is directly impacted by secondary treatment, and because it provides a holistic measure of aquatic ecosystem health. Appendix A.3 provides further discussion on dissolved oxygen and its relevance as a measure of water quality.

Table 1.1 compares descriptive statistics across *ex ante* compliant and noncompliant cities (i.e., compliant or not with the CWA technology standard of secondary treatment prior to 1972). Voluntary secondary treatment is correlated with both ability to pay for, and propensity to benefit from, water pollution abatement. Secondary treatment adoption is positively correlated with wealth (e.g., share of population with a college degree, median housing prices, and revenues per capita), preference for environmental protection (e.g., conservation score), higher levels of manufacturing industry, receipt of intergovernmental funding, and proximity to waterbodies. Expenditures appear overall balanced, with the noticeable exception of wastewater expenditures: cites that already adopted secondary treatment spent nearly double per capita on wastewater treatment prior to the CWA. Figure 1.5 shows the spatial distribution of *ex ante* compliance aggregated to county-level averages for exposition purposes. *Ex ante* compliant cities are more likely to be near large lakes, population centers, or manufacturing-intensive areas, such as Tennessee, Michigan, Pennsylvania, and New York. These substantial differences underscore the importance of using

[17] I focus on a distance of 25 miles following Keiser and Shapiro (2018).

an empirical approach that eliminates potential confounding factors correlated with outcome differences across compliant and noncompliant cities.

1.5 Empirical Strategy

The goal of my analysis is to estimate the impact of the CWA infrastructure mandate on local government budgets and growth. The relationship of interest is:

$$y_{irt} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + (\gamma_r \times t) + \tau_t + \nu_i + \varepsilon_{irt} \quad (1.2)$$

where y_{irt} is one of several outcomes related to municipal expenditures (e.g., wastewater expenditures) or growth (e.g., population) for city i in geographic region r in year t.[18] P_i is an indicator variable equal to 1 if a city is *ex ante* noncompliant, meaning it had only primary treatment technology at the start of the CWA in 1972, and 0 if a city is *ex ante* compliant, meaning it had at least secondary treatment as of 1972. $POST_t$ is an indicator equal to 1 for all post-CWA years (e.g., years 1977 and later because my panel structure is quinquennial).

As *ex ante* compliant and noncompliant cities exhibit substantial differences in observable characteristics (see Table 1.1), I include a vector of pre–CWA city characteristics, $\mathbf{X_i}$, whose effects are allowed to vary by year (Lechner et al. 2011). The pre–CWA city characteristics in $\mathbf{X_i}$ include income per capita, share of employment

[18] I follow the Bureau of Economic Analysis (BEA) definition to categorize states into one of eight US regions: New England, Mideast, Great Lakes, Plains, Southeast, Southwest, Rocky Mountain, and Far West.

in all manufacturing as well as water-polluting manufacturing, river population size to control for historic differences in riparian connections to urban markets (Bleakley and Lin 2012, Rappaport and Sachs 2003), distance to coast, and receipt of inter-governmental grants to account for differences across cities in their fiscal managerial capabilities to obtain federal or state funding. The year fixed effects τ_t control for macroeconomic time-varying determinants of y_{irt} common to all cities, such as federal budget cycles, while ν_i captures all unobserved, time-invariant differences across cities that affect y_{irt} such as distance from a water body or soil and landscape attributes that affect infrastructure construction costs. $(\gamma_r \times t)$ is a vector of region-specific linear time trends, which ensures estimated differences in y_{irt} across cities are not driven by divergent patterns of growth across regions. Lastly, ε_{irt} is an error term. I cluster standard errors at the city level to account for city-specific correlations in unobserved components of spending and growth over time. The coefficient of interest is β, which measures the differential change, conditional on controls, in outcome y between noncompliant and compliant cities before versus after the CWA.

For the differences-in-differences estimate of the CWA mandate effect to be unbiased, $(P_i \times POST_t)$ must be uncorrelated with the error term ε_{irt}. That is, potential outcomes y would have trended similarly for *ex ante* compliant and noncompliant cities in absence of the CWA technology standard. There are several reasons why this assumption is problematic in this context. First, the architecture of the CWA legislation included not only enforcement of a technology standard, but also substantial federal financial assistance in the form of capital construction grants. Congress distributed over \$153 million in outlays from 1973 through 1986 for the construction and modification of municipal wastewater treatment plants (Copeland 2015).[19] These construction grants provided funding not only for *ex ante* noncompliant cities to upgrade from primary to secondary, but also for *ex ante* compliant cities to invest beyond the minimum requirements in advanced treatment technologies. Table 1.1 suggests that compliant cities are historically more capable of obtaining intergovernmental funding, therefore the availability of CWA federal grants is likely to have impacted *ex ante* compliant cities differently than *ex ante* noncompliant cities. This means that simply controlling for baseline differences in grant receipt will not remove bias induced by the contemporaneous federal grants shock. In particular, difference-in-differences estimates of β are likely to be attenuated toward zero for fiscal and water quality outcomes because control cities may respond to the federal grants program by further investing in pollution abatement technology.

Second, the early 1970's were a time of several major environmental regulations, including the Safe Drinking Water Act in 1974 and the Clean Air Act in 1970, the latter of which affected the economic development of regulated counties through its impacts on industry, labor markets, and local amenities (Greenstone 2002; Kahn 2001; Lin 2016). If violations across these Acts were correlated within cities, deviations from trend that appear after the CWA may be spurious. Lastly, differences in baseline wealth, surface water quality, and environmental preferences shown in Table 1.1 suggest that taxpayers may have sorted differentially into *ex ante* compliant cities relative to noncompliant cities even absent the CWA technology standard. If

[19] See Keiser and Shapiro (2018) for an empirical analysis on the cost-effectiveness of the CWA grants program.

ex ante compliant cities were more competitive in attracting taxpayers, differences-in-differences estimates of β will be biased away from zero toward a negative growth effect.

To address these various sources of bias, I employ an instrumental variable approach that uses variation in downstream population across cities combined with variation in pre-CWA water pollution abatement across states to predict CWA compliance status. The key to my identification strategy is that I exploit variation that is external to the city; I predict CWA compliance from factors that are unlikely to be correlated with local taxpayer preferences for public goods or local abilities to obtain federal grants.

1.5.1 Instrumental Variable Approach

The positive relationship between downstream population size as well as state environmental regulation and pre-CWA secondary treatment adoption (as described in Section 1.3.3) forms the basis of my identification strategy. Specifically, cities situated upstream of population centers were more likely to adopt secondary treatment technology prior to the CWA regulations relative to low-downstream-population cities. Further, cities located in states with stronger regulation of surface waters prior to the CWA were more likely to adopt. The CWA secondary technology standard was, therefore, more likely to bind for cities with smaller populations downstream and weaker state regulation of water pollution.

I construct the downstream population component of the instrument using digital spatial maps sourced from the National Hydrography Dataset Plus of the US Geological Survey (USGS). These maps contain hydrologic information for over 2.6 million stream segments averaging 1 kilometer in length. Every river segment possesses three identifying attributes that allow me to trace out all possible linkages in the US river system: a segment identification code, the code of the immediate upstream river segment, and the code of the immediate downstream river segment. In addition, all segments include an identifier for the terminal point of its river network (i.e., the river "mouth"). The combination of network linkages across segments and terminal point identifiers allows me to identify upstream versus downstream relationships across cities located on the same major river (e.g., the Mississippi) as well as across cities on differing tributaries sharing the same major river basin (e.g., the Illinois and Ohio rivers, which both feed into the Mississippi).

I assign each city centroid to its closest stream segment using GIS software. My criteria for matching cities to a stream segment is to select the six closest stream segments to a city centroid and assign the city to the stream segment with the lowest branching level. This approach accounts for the tendency of cities to divert wastewater effluent into the main river segment closest to their city as opposed to a small tributary. I then calculate each city's cumulative downstream population through a recursive "search tree" algorithm as follows: I first find the terminal point, or the mouth, of each river network and assign this segment a downstream population of $x_i = 0$ and a current population of x_j equal to the population of a city at that mouth, if one exists. Moving upstream along stream segments, indexed by j for current and

i for the relative downstream segment, I sum the population x_i of any cities located along those segments until a branching occurs. The branch point is again treated as a temporary "river mouth" with a downstream population of $\sum_0^i x_i$, and the process repeats itself until the source ($j = N$) of the river is reached, with a total downstream population of $x_j + \sum_0^N x_i$. The final result is a downstream population value for every river segment. The "search tree" algorithm provides important precision in my measure of downstream population by explicitly accounting for tributary branching within river networks. A naive reliance on distance to river mouth across cities without accounting for branching would induce substantial measurement error into the downstream population calculation.

Figure 1.6 shows downstream populations, aggregated as county means for expositional purposes. Cities with a higher downstream population are generally located near the headwaters of populous river networks, such as the upper Missouri, the upper Mississippi, and the upper Ohio rivers. Importantly, downstream population is not only a function of river length. For example, cities along the Columbia river in Washington state have high downstream populations despite a shorter river length. Cities along the similarly-sized Colorado River in Arizona, in contrast, have relatively low downstream populations owing to the relatively low population density in the American southwest.

The second component of the instrument exploits pre-existing differences in water pollution regulation across states to predict municipal *ex ante* compliance. My instrument includes the share of wastewater treatment plants within a state that had secondary treatment technology prior to the CWA. Figure 1.7 shows variation across states in pre-CWA municipal secondary treatment adoption. States with historically more environmental legislation, such as Pennsylvania and New York, had higher levels of secondary treatment prior to the CWA.[20,21]

My identification strategy thus exploits two sources of variation to predict *ex ante* compliance: pre-CWA downstream population size and pre-CWA state composition of compliant wastewater treatment plant technology. In my empirical approach, these fixed characteristics predict differences in outcomes across cities as a function of their differential effects prior to versus after the CWA. The first and second stage estimation equations are shown below in Eqs. 1.3 and 1.4. Let i index cities and t index years. P_i is an indicator equal to 1 if a city is *ex ante* noncompliant and y_{ist} is one of several outcomes of interest (e.g., wastewater expenditures, population, etc.).

$$y_{ist} = \beta_{IV}(\widehat{P_i \times POST_t}) + \mathbf{X_i}\theta_t + (\gamma_r \times t) + \tau_t + \nu_i + \epsilon_{ist} \tag{1.3}$$

$$\begin{aligned} P_i \times POST_t = {} & \alpha_1(D_i \times S_s \times POST_t) + \alpha_2(S_s \times POST_t) \\ & + \alpha_3(D_i \times POST_t) + \mathbf{X_i}\lambda_t + (\gamma_r \times t) + \tau_t + \nu_i + \mu_{ist} \end{aligned} \tag{1.4}$$

In Eq. 1.4, the first three terms are the excluded instruments, where D_i is a city's downstream population as of 1970 and S_s is pre-CWA state technology composition, measured as the share of all wastewater treatment plants in state s with secondary

[20]As of 1970, twenty-nine state constitutions included standards on water quality (Andreen 2013), though enforcement of these regulations was generally infrequent (Rechtschaffen 2003). Pennsylvania began regulating water quality in 1937 with the Clean Streams Law, and several requirements of the CWA were enacted in Pennsylvania prior to 1968 (Walters 2017). New York state laws on watershed regulations date back to the early 1900s (Hudson River Watershed Alliance 2015.)

[21]The pre-CWA state composition of treatment plants can be equivalently interpreted as the predicted CWA technology incidence in the state. In their identification strategy, Duflo and Pande (2007) employ a similar approach by predicting district-level dam adoption in India using a state's baseline share of all national dams.

treatment as of 1972. As before, $\mathbf{X_i}$ is a vector of pre-CWA determinants of y whose effects are allowed to vary with time, including intergovernmental grants, industry mix, distance to coast, and river population size to account for time varying effects of historic exposure to urban markets via river-based trade routes (Bleakley and Lin 2012, Rappaport and Sachs 2003). This river population control is important because this ensures that *ex ante* compliance is predicted from a city's relative positioning on a river system, and not by the overall population size of their river. Standard errors are clustered at the city level.

The remaining variation that powers my instrument is the population size downstream, state secondary treatment share and its interaction with population downstream. Figure 1.8 shows the variation I exploit. For a given downstream population, there is a range of state technology composition. Similarly, states with similar technology composition harbor cities with varying degrees of downstream population sizes. The interaction of these two terms is important because the marginal impact of downstream population varies with pre-existing, state-level regulations of water pollution. Specifically, cities located in states with more regulation of water pollution were more likely to adopt wastewater treatment as a result of downstream population pressure relative to cities located in states with less regulation of water pollution. The state compliant plant share term ($S_s \times POST_t$) enters into the instrument to take advantage of variation in pre-existing state environmental regulations.

In Table 1.2, I formally examine the first stage relationship between a city's downstream population, state plant composition, and CWA compliance status. These results confirm the importance of water pollution externalities: a city with higher likelihood of inflicting pollution on its neighbors was more likely to have adopted secondary treatment prior to 1972. Columns (1) and (2) use cross-sectional variation. A one standard deviation increase in downstream population reduces the likelihood that a city had only primary treatment as of 1972 by approximately 7.5 percentage points from a mean noncompliance rate of 75%. These results are robust to including controls for distance to river mouth and distance to river edge, which supports that downstream population, as opposed to the size or proximity to a river, provides the relevant variation in compliance.

Columns (3) and (4) show estimates of the first stage Eq. 1.4, where downstream population is interacted with state pre-CWA share of secondary treatment plants. Cities in states with a higher share of secondary treatment plants are less likely to be *ex ante* noncompliant (more likely to be *ex ante* compliant). Conditional on state share, cities with a larger downstream population are additionally less likely to be noncompliant. These results demonstrate that the marginal impact of downstream population for *ex ante* compliance was stronger in states with more water pollution abatement. Accounting for baseline differences across cities in income, intergovernmental funding, or geographic region does little to impact the likelihood of secondary treatment adoption. Column (4) employs the full set of controls and is my preferred specification of Eq. 1.4. At the mean state share level, a standard deviation increase in downstream population reduces the likelihood of *ex ante* noncompliance by 2.7 percentage points, from a mean of 75%. The Kleibergen-Paap first stage F-statistic (which is robust to non-i.i.d. errors) is above the Stock and Yogo (2005) critical value of 9.08 (for 10% maximum relative bias). I also fail to reject that the instruments are uncorrelated with the error ε_{ist} with a p-value of 0.44. Table 1.2 shows that my instruments are sufficiently strong.

1.5.2 Validity of Instrumental Variable Approach

A causal interpretation of the instrumented differences-in-differences parameter β_{IV} requires two assumptions (Hudson et al. 2015). First, the evolution of outcomes across cities with different pre-existing state compliance shares should have trended similarly; and —conditional on state compliance—the evolution of outcomes across cities with high- versus low- downstream populations should have trended similarly absent the CWA technology mandate. Second, the exclusion restriction requires that shocks which co-vary with the CWA do not differentially impact cities with high- versus low-downstream population sizes, or cities in high- versus low- compliance states. In other words, the exclusion restriction requires that downstream population size and state compliance share explain post-CWA differences in outcomes *only* through their influence on pre-CWA wastewater treatment plant technology adoption.

I assess the plausibility of these assumptions by testing for the presence of pre-trends of city characteristics across municipalities exposed to above median versus below median exposure to the instrument. Table 1.3 column (2) shows that cities predicted to be in the treatment group have similar pre-CWA growth trends compared to cities predicted to be in the control group on the basis of the instrument.

In contrast, column (1) shows apparent differences in growth trends across observed treated versus control cities prior to the CWA for population, wastewater expenditures, user fees, and receipt of federal grants. These baseline characteristics are more comparable across cities on the basis of the instrument. The fact that the instrumental variable approach invites comparable pre-trends in federal grant receipt is particularly important; this reduces concern that the CWA grant program will have differentially impacted cities predicted to be control relative to those predicted to be treated. Taken together, Table 1.3 suggests that potential outcomes in absence of the CWA are more likely to trend in parallel under the instrumental variable approach compared to a basic difference-in-differences approach.

I further test the plausibility of these assumptions with respect to my instruments by comparing pre-CWA trends in wastewater spending across cities with high relative to low downstream populations as well as high relative to low state compliance share cities. Panel A in Figure 1.9 plots the linear combination of estimates $\delta_{t\times\underline{50}} + \delta_t$ in black and δ_t in gray from the following equation:

$$y_{it} = \sum_t \delta_{t\times\underline{50}}(I_{\underline{50}} \times S_s \times D_t) + \sum_t \delta_t(S_s \times D_t) + (N_R \times D_t) + \nu_i + \varepsilon_{it} \qquad (1.5)$$

where $I_{\underline{50}}$ is an indicator for a city having downstream population size in the bottom 50th percentile and N_R is river population. An estimate of $\delta_{t\times\underline{50}} > 0$ indicates higher expenditures in year t relative to 1972 for cities in the bottom 50th percentile of the downstream population distribution relative to cities in the top 50th percentile. Prior to the CWA, cities with high and low downstream populations had similar wastewater expenditures per capita, suggesting that cities with differing downstream population sizes had similar potential outcomes in absence of the CWA. However, their expenditures per capita diverge beginning after 1972: cities with *lower* downstream populations incur larger wastewater expenditures per capita following the CWA relative to cities with downstream populations in the top 50th percentile.[22] Panel B similarly shows how cities with 1972 state compliance share in the bottom 50th percentile of the distribution incur larger wastewater expenditures per capita after the CWA relative to cities in states with higher compliance shares. The absence of significant pre-trends in Figure 1.9 supports the identifying assumption that differences across low and high downstream population cities, and low and high state compliance share cities impact expenditures only through their impacts on CWA compliance status.

The local average treatment effect (LATE) identified by my instrument is the effect of the CWA infrastructure mandate among cities with low riparian exposure to downstream populations and low levels of pre-CWA state regulation. The CWA was binding for these cities because they faced little pressure from downstream to treat their wastewater. The comparison group of cities are those that adopted compliant infrastructure prior to the CWA as a consequence of downstream pressure. Recall that *ex ante* compliant cities could apply for and use federal infrastructure grants under the CWA to further upgrade their secondary treatment plants. The correlated shock of the CWA grants program that may have impacted both the treatment and

[22] Appendix Figure A.2 shows the same figure after including year fixed effects in Eq. 1.5. Consequently, $\delta_{t\times\underline{50}} + \delta_t$ and δ_t capture the difference across low relative to high downstream population cities conditional on being located in a high compliance state. The same pattern holds; however, the coefficients capture the effect of low versus high downstream population relative to all cities in low compliant states.

control cities under the difference-in-differences approach is less likely to manifest with the IV approach. My instrument exploits variation across cities in the degree of external pressure they faced to abate surface water pollution. Consequently, the cities induced by the instrument to be *ex ante* compliant are less likely, relative to the self-selected *ex ante* compliant cities, to undertake federal infrastructure grants and inframarginally consume additional wastewater treatment. My IV estimates do *not* capture the effect of the CWA infrastructure mandate due to alternative reasons that would cause the policy to bind, such as insufficient finances or local taxpayer indifference to environmental protection. Consequently, the LATE identified by my instrument may differ from the impact experienced by the average city bound by the infrastructure mandate.

1.6 Results

My empirical tasks are threefold: first, I identify the magnitude of direct compliance costs and benefits by examining changes in wastewater expenditures and water quality. Second, I test how cities financed those direct costs by estimating differences in expenditures of non-wastewater public goods and municipal revenue sources. Finally, I estimate the indirect, non-pecuniary impacts of compliance by testing how water quality, population, housing prices, and demographic composition of city residents change differently among noncompliant cities relative to compliant cities. I present results from both the differences-in-differences approach and the instrumental variable approach. The differences-in-differences approach provides useful information on the policy-relevant average impacts of CWA noncompliance and is a likely lower bound on the fiscal impacts of compliance. In contrast the instrumental variable estimates are more likely to capture a "mandated expenditure" effect because compliance status is identified from fixed characteristics that are external to the city's own decision-making process.

1.6.1 Local Government Budgets & Water Quality

I start by examining the effect of noncompliance on ambient water quality. I test the assumption of common potential outcomes by estimating a dynamic effect specification that allows for visual examination of pre-trends in the data. For the differences-in-differences estimator, the dynamic effects estimating equation is a flexible version of Eq. 1.2, where the impact of noncompliance is allowed to vary in each year:

$$y_{irt} = \sum_{t='67(\ni'72)}^{'02} \delta_t(P_i \times D_t) + \mathbf{X_i}\theta_t + (\gamma_r \times t) + \tau_t + \nu_i + \varepsilon_{irt}. \tag{1.6}$$

The coefficient δ_t measures the difference, conditional on controls, in outcome y between noncompliant and compliant cities in year t relative to 1972.

Figure 1.10 Panel A presents evidence of the initial amenity effect experienced by noncompliant cities. The figure plots estimates of δ_t from Eq. 1.6, where the dependent variable is the five-year annual average dissolved oxygen in milligrams per liter for city i in years t to $t-5$. The bars show 95% confidence intervals, and the dashed line denotes the start of the CWA. An estimate of $\delta_t > 0$ indicates higher water quality in noncompliant cities relative to compliant cities in year t relative to 1972. While limited availability of data prior to the CWA at the level of local governments restricts me from observing pre-trends beyond one period, the lack of significant differences in spending leading up to the CWA is suggestive that compliant and non-compliant cities had similar potential trajectories. There is a lagged increase in water quality following about 15 years after the increases to wastewater expenditures. On average, water quality increased by 0.5 milligrams per liter from pre-CWA levels, or less than 1%. This is significantly lower than the 7% improvement in water quality implied by Keiser and Shapiro's (2018) estimate of reductions in dissolved oxygen deficit as a result of the CWA construction grants. However, the effect is imprecisely estimated.

Figure 1.10 Panel B shows the dynamic effects after using downstream population and state plant composition as instruments for *ex ante* compliance status. Changes to water quality are substantially larger in magnitude and show a persistent upward trend through the end of the study period relative to the differences-in-differences estimates. Because the IV approach removes potential bias from control cities responding to the CWA treatment, the estimated change in water quality better characterizes the causal impact of unanticipated wastewater treatment adoption. Water quality improved nearly 2mg/l on average after the CWA, or an 18% increase from pre-CWA levels. My results are consistent with those implied by Keiser and Shapiro (2018). The authors find that the average infrastructure grant improved dissolved oxygen deficit by 7% in the year following grant receipt, and improvements grew in magnitude over time. They also find positive dose-response effects to additional grants. Given that many cities received at least three infrastructure grants, and that my study period averages effects over nearly 30 years, an 18% improvement to water quality following CWA compliance is consistent with their findings.[23]

Panel A of Table 1.4 provides differences-in-differences estimates of Eq. 1.2; and Part B, instrumental variable estimates of Eq 1.3.[24] Wastewater expenditures increased substantially more from pre-CWA levels relative to all other expenditure categories. The differences-in-differences results in Panel A show minor increases in public safety and general and administrative spending, suggesting potential crowd-*in* effects of the CWA mandate. Increases in wastewater expenditures induced increases in intergovernmental grant receipt, as well as user fees. Overall tax revenues did not significantly change, as decreases in property taxes offset increases in sales & license tax revenues.

The IV estimates in Panel B show similar results, however with substantially larger magnitudes. Attenuation of the differences-in-differences estimates is consistent with uptake of CWA federal grants and increased expenditures on wastewater treatment among the control group. *Ex ante* compliant cities may be differentially impacted by the grants program because of their pre-existing advantage in obtaining

[23] Appendix A.4 provides further discussion on the comparability of my water quality estimates to those of Keiser and Shapiro (2018).

[24] Appendix Tables A.1 and A.2 show the sensitivity of wastewater expenditures and growth outcomes to various specifications. Inclusion of city-level fixed effects in Appendix Tables A.1 increases the CWA effect on wastewater spending substantially, suggesting that unobserved fixed differences across cities are highly correlated with expenditure decisions. The additional controls mainly serve to increase precision of the point estimate, but do not substantially change the magnitude. The point estimates in Appendix Table A.2 show greater sensitivity to controls, which is likely due to the more limited panel. Controls serve to attenuate the growth effects.

federal funding. Thus, the IV estimates provide a more accurate estimate of the fiscal costs required to comply with CWA infrastructure mandate because the LATE group of cities are more comparable with respect to pre-existing wealth and intergovernmental funding (see Table 1.3). Wastewater expenditures increased by $151 per capita, or over 200 percentage points after the CWA. Both wastewater capital and operating costs increased. The estimated increase to total wastewater expenditures closely aligns with engineering cost estimates from the EPA CWNS. These surveys show that additional expenditures required for secondary treatment adoption for a city of 30,000—the mean population size in my sample—is approximately $6 million per year.

Importantly, Panel B shows no evidence of crowd out in the funding of other goods. In other words, cities did not respond to the federal mandate through austerity measures. Rather, total city expenditures increased by approximately 47%, driven partially by increased public safety expenditures contemporaneously with wastewater expenditures.

To meet these increased expenditures, cities increased user fees and receipt of intergovernmental grants. While wastewater user fees do not precisely match with changes in wastewater expenditures, the disparity could be due in part to measurement error from the Census of Governments. Data from the Census of Governments are self-reported by local government administrators, and local expense and revenue categories do not always match up with the Census of Governments structure. For this reason, estimates for on specific line items, like wastewater user fees, are likely to be attenuated from measurement error.

Since user fees are indexed directly to the consumption of the public good, user fees are non-distortionary. However, the $61 per capita (or, over $180 per household) annual increase could place a nontrivial burden on local taxpayers, particularly those with lower incomes that cannot easily substitute away from consuming water.[25] The fact that intergovernmental grant receipts per capita increased more than wastewater expenditures per capita is suggestive of strategic complementarities across applications for grant funding. That is, it is possible that cities responded to the CWA federal grants program by not only applying for treatment plant construction grants but for other federal grant programs as well. Case studies of 16 communities, collectively, in Hanford and Sokolow (1987) and Weiland (1998), found that municipal CWA compliance served to improve their financial positions and organizational skills for acquiring intergovernmental grants.

In summary, Table 1.4 demonstrates that wastewater expenditures per capita tripled after the CWA, with little change in expenditures on other goods and services. Lack of displaced funding contrasts with prior work that largely finds local governments reduce spending on goods and services in response to fiscal shocks. However, the asymmetric response to expenditure liabilities relative to tax revenue loss may be additional evidence of the "flypaper effect" (Hines and Thaler 1995), whereby governments act as though money is not fungible and respond to taxpayer wealth shocks differently than proportional shocks to local fiscal obligations.

[25] Per Foster and Beattie (1979), the price elasticity of demand for water among US consumers is inelastic, at -0.1 (see Table 2).

Do these budgetary and amenity changes, in turn, impact municipal growth? To answer this question, I employ dependent variables drawn from the Census and measure city growth through pre- versus post- CWA changes in population, housing prices, and composition of residents with high educational attainment.

1.6.2 Local Government Growth

I explore how changes to local government budgets impacted municipal demographics and property values in Table 1.5. Unlike the municipal finance data which I observe once every five years from 1967, Census outcomes are observed only once per decade. Consequently, results in Table 1.5 are estimated from a more limited panel where observations from 1972 provide pre–CWA outcomes, and observations from 1982 and 1992 provide post–CWA outcomes. Difference-in-differences estimates in Panel A show that the noncompliant cities experienced declines relative to compliant cities in all three measures of growth. Compared to *ex ante* compliant cities, *ex ante* noncompliant cities grew 3% slower with respect to population, 0.6% slower in median housing prices, and 5% slower with respect to share of the working-age population with a college degree. The population and housing price results suggest that residents of noncompliant cities did not value benefits from additional wastewater infrastructure more than their costs. Households adjusted locations in response to the federal mandate and capitalization effects are negative.

The negative growth response is intuitive if cities consume public goods only up to the point where aggregate social benefits outweigh the marginal cost of providing the goods (Samuelson 1954). Compliant cities likely adopted prior to the CWA because they garnered higher benefits from wastewater treatment compared to noncompliant cities, due to either easier access to surface waters or greater treatment needs from manufacturing industry. This selection is problematic for causal identification of federally mandated expenditures, particularly if differences across cites in their ability to benefit from, or ability to pay for, wastewater treatment infrastructure generates more positive growth outcomes for *ex ante* compliant cities. For example, Carlino and Saiz (2008) and Guerrieri et al. (2013) show that cities with greater levels of natural amenities and initial wealth, respectively, exhibited significantly faster population and housing price growth during the late 20th century. Presence of such selection would bias the differences-in-differences CWA effect away from zero.

The instrumental variable results in Table 1.5, Panel B, indeed, show effects of the CWA mandate that suggest a more positive effect of the CWA mandate for city growth. Cities predicted to be noncompliant as a function of their downstream population and pre-CWA state adoption share experience 12.6% greater changes to population and 1.1% higher housing price growth relative to control cities, although the housing price result is not statistically distinguishable from zero. While the population effect is large, the implied 12.6% increase in treated city populations following an 18% improvement to water quality is consistent with prior literature that tests for migration effects from environmental regulation. Banzhaf and Walsh (2008) find a 7% increase in population over a decade among communities that lost exposure to TRI chemical emissions, while Kahn (2000) estimates a population increase of 8% over 15 years among California counties that experienced an average improvement to ozone exposure of 21%.

Median housing values grew by 1.1% more on average among noncompliant cities, although the estimate is imprecisely estimated. This represents the combined capitalization effect of both increased user fees, and improvements to local ambient water quality, as well as any subsequent general equilibrium changes resulting from either effect. Appendix Table A.3 provides details on a back-of-the-envelop calculation for resident fee capitalization. Under the more conservative rate of return assumption, these estimates indicate that residents value water quality improvements at twice their marginal cost. The effect of the CWA mandate on user fees is of first order importance. Although the general equilibrium effects complicate interpreting the hedonic estimates as a direct willingness-to-pay for water quality improvements, ignoring the impact of the CWA on local public finance would substantially underestimate the implied value of water quality improvements.

The last column of Panel B shows that the share of residents with a college degree increased by over 3 percentage points as a result of the CWA infrastructure mandate. This result is consistent with prior work by Sieg et al. (2004) and Banzhaf and Walsh (2008) that find evidence of gentrification effects following exogenous improvements to environmental amenities. Consistent with their findings, the positive parameter estimate on housing prices suggests that increased cost of living may explain relative out-migration of individuals with lower educational attainment. The growth results collectively suggest that taxpayers of mandated cities valued the treatment infrastructure above its local cost, which taxpayers realize as increased user fees.

1.6.3 Robustness Checks

My identification strategy relies on comparisons across cities with differing downstream population levels within states of similar pre-CWA water pollution abatement activities. The size of a city's downstream population serves as a proxy for the degree of litigation pressure it could have received from exposing downstream cities to its surface water pollution. To assess whether power from downstream is the driving mechanism, I re-estimate the two-stage least squares results using the count of large downstream cities as an instrument for *ex ante* compliance. I define a city as "large" if its 1970 population was at least 30,000, the approximate mean city size in my sample. Table 1.6 shows that downstream litigation pressure explains the relevance of the downstream population instrument. Having one additional city of at least 30,000 people downstream increased the likelihood that a city was *ex ante* compliant by 2%. In contrast, having one additional small city downstream has nearly no impact on the probability of compliant technology adoption. Appendix Table A.4 and A.5 shows that estimated effects on budgetary and growth outcomes using the count of large cities downstream as an instrument closely mimics my main results.

The validity of the instrumented differences-in-differences design requires that outcomes across cities with high versus low exposures to the instrument would trend in parallel absent the CWA technology standard. One concern suggested by Figures 1.6 and 1.7 is that, both, downstream population size and state compliance is regionally determined. Northern cities, for example, have higher downstream populations on average relative to southern cities because several major US rivers begin in the north and terminate along coastlines in the south. Additionally, coastal cities generally have lower downstream populations relative to interior cities. The regional correlations generated by my instrument may be problematic because coastal cities are unlikely to be comparable controls to interior cities,[26] and likewise northern cities are unlikely to be comparable controls to southern cities. Appendix Tables A.6 through A.9 show that my instrumental variable results are robust to excluding coastal cities, and are robust to excluding hydrologic regions with the largest downstream populations (the upper Mississippi and the Ohio river watersheds). This suggests that my findings are not driven strictly by divergent growth trends across coastal relative to interior cities, or by northern relative to southern cities.

I lastly show that results are robust to using the full, unbalanced panel of municipal wastewater treatment plants (including plants that were built after the CWA) in Appendix Tables A.10 and A.11. This final robustness check provides qualitatively similar results to that of my main balanced sample. However, significant increases in other goods expenditures and population levels due to CWA noncompliance are suggestive that the unbalanced panel of cities likely includes those that adopt wastewater treatment due to unobserved local demand shocks, and not purely due to the binding

[26] A small number of coastal cities face different treatment technology standards from interior cities. Known as 301h Wavier Recipient facilities, some treatment plants that discharge into coastal waters are exempt from secondary treatment requirements. As of 1994, only nineteen facilities in the lower 48 states had this exemption (EPA 1994). Most are located in Maine. I remove such facilities as a robustness check in Appendix Table A.6 and find consistent results.

constraints of the mandate. This result supports my selection criterion used for my main results, whereby only cities that have had an operating wastewater treatment plant since 1972 are included in the sample.

1.6.4 Limitations & Interpretation

I consider some of the potential limitations associated with these data and my experimental design. These potential limitations will generally tend to bias my results toward a null effect. First, I am comparing treated cities to control cities, whereby the treated cities may receive an exogenous improvement to surface water quality after complying with mandated investment in wastewater treatment infrastructure. However, municipal boundaries may not provide the correct spatial extent of pollution abatement effects from wastewater treatment. If the infrastructure's actual impact is more narrow than municipal borders, any perceived benefits of wastewater pollution measured at the this level will be diluted. Conversely, if the actual impacts of surface water pollution extend beyond municipal borders, to downstream cities, for example, the mandated infrastructure will have some impact on control cities, again diluting the identified differential across these communities.

Second, multiple municipalities may share a single wastewater treatment plant, particularly if those municipalities are located close together. I can identify the municipality that manages a publicly owned plant from the CWNS data, but I cannot distinguish whether other municipalities are serviced by that plant. This will not compromise the diagnosis of treatment versus control cities in my design, as I consider only municipalities that are, themselves, the managing authority of a plant. However, to the extent that there is cost sharing of mandate compliance across unobserved communities, the estimated differential on expenditure changes will be diluted. In Appendix Figure A.3, I compare the plant service population of each plant reported in CWNS to the Census population estimate for the plant's managing municipality and find a correlation very close to unity. This suggests that mis-measurement of the per capita compliance costs borne by municipal residents is likely to be minimal.

Finally, the interpretation of my hedonic estimates on property values are best interpreted as a capitalization effect of mandated wastewater treatment infrastructure rather than the exact willingness-to-pay. This is because, first, the impacts on water quality following the CWA are likely non-marginal and, second, my identification strategy exploits a long panel. Over this time period, preferences among treated city residents likely changed as the residential composition of treated cities shifted toward higher educational attainment (see Table 1.5). When hedonic analysis is used to estimate "large" changes in public goods, Kuminoff and Pope (2014), Banzhaf (2018), and Sieg et al. (2004) suggest that the resulting partial equilibrium estimates - as in my empirical design - will likely understate residents' willingness to pay and are better interpreted as a lower bound on the Hicksian equivalent surplus.

1.7 Who Benefits from the CWA mandate?

The aggregate effects of CWA infrastructure mandate estimated in section 1.6 potentially mask important sources of heterogeneity. Figure 1.4 suggests that per capita costs of secondary treatment adoption diminish with population size, at least with respect to capital costs. In the following section, I test how the effects of CWA compliance differ according to city population and explore plausible mechanisms including differences in housing supply, exposure to upstream polluters, and receipt of federal grants. I also test for differences across cities in the implied value of mandate compliance and find that water quality improvements are more likely to be valued by residents in areas where water recreation is more feasible and among cities whose water quality is more reliant on the actions of upstream polluters.

1.7.1 Heterogeneous Responses by City Size

Table 1.7 shows IV estimates of:

$$y_{ist} = \delta_{POST}\widehat{(P_i \times POST_t)} + \delta_{POST\times\omega}\widehat{(P_i \times POST_t \times \omega_i)} \tag{1.7}$$

$$+ (\overline{N}_i\psi_t) + \mathbf{X_i}\theta_\tau + (\gamma_r \times t) + \tau_t + \nu_i + \varepsilon_{ist}$$

where ω is an indicator variable equal to 1 if a city has above median population size. I add a control for baseline differences in population size interacted with year fixed effects, $\overline{N}_i\psi_t$, to account for differential growth trends across large relative to small cities over the study period. All other variables are as before in Eq. 1.3. For the local budgetary outcomes, I expect $\delta_{POST \times \omega} < 0$ if the effects of CWA compliance are less costly for larger cities.

Table 1.7 confirms that compliance costs are lower in larger cities. Larger city wastewater expenditures are, on average, \$33 less per person per year relative to smaller cities, or about 80% the costs of smaller cities. These per capita expenditures savings for large cities translate into user fee savings, as well. Notably, intergovernmental grant receipts per capita are approximately 17% larger in larger cities. Larger cities also appear better able to finance the mandated costs with long-term debt, which may speak to their relative advantage in credit markets. Results of Table 1.7 are consistent with the existence of scale economies in mandated wastewater treatment infrastructure.

In Appendix A.5, I test how these cost differences translate into differences in growth outcomes and explore mechanisms for divergent responses of small relative to larger cities.

1.7.2 Federal Grants and City Size Heterogeneity

The CWA infrastructure mandate spurred asymmetric responses across large relative to small noncompliant cities. While both large and small cities increased municipal expenditures substantially following the 1972 CWA, smaller cities experienced significantly larger cost increases, ranging from 14-20% above the per capita costs of larger cities. A natural question is whether changes to the structure of the CWA mandate could have mitigated some of the fee increases; and, more importantly, whether such changes could have generated gains to economic efficiency.

The CWA construction grants program distributed approximately \$153 billion in federal aid to local governments for the purpose of municipal wastewater treatment and sewerage improvements between 1973 and 1986 (Copeland 2015). Table 1.7 shows that relative to larger cities, smaller cities received disproportionately less federal aid following the 1972 CWA. The change in wastewater expenditures for small cities was approximately \$33 higher per capita, however, smaller cities received approximately \$61 less per capita in federal aid compared to large cities.[27]

Whether redistribution of federal funds from large cities to small cities would have improved overall social welfare depends upon the relative effectiveness of grants in large versus small cities. Only if federal dollars induced larger benefits on the margin to smaller cities relative to larger cities would redistribution of funds have improved economic efficiency.

In Table 1.8, I test whether grant funding was more effective at mitigating negative growth outcomes in small versus large cities. Grant receipt is likely to be endogenous to growth outcomes, however the comparison across large and small cities is suggestive of grant mis-allocation. This table shows results from estimating Eq. 1.7 for two separate samples: above median and below median population cities. Here, ω is a city's annual receipt of federal grants. The coefficient $\delta_{POST \times \omega}$

[27]While the "Inter Govt" category includes federal grants, not just those for wastewater infrastructure, I show in Appendix Table A.22 that the same allocative pattern holds when considering the grant category most likely to include CWA federal construction grants.

is the difference in one of several outcomes - user fees, population, property values, or high skill population share - among treated cities that received additional grant funding. Although the point estimates are not directly comparable across samples, these results are suggestive that federal grants are effective at reducing user fees and mitigating housing price decline among smaller cities. In contrast, user fees and growth outcomes appear to be insensitive to grant receipt among large cities. The marginal effect of federal grant dollars in large cities are precisely estimated zeros. Grants have little additional impact on mitigating the CWA budget shock among large cities.

Results of Table 1.8 indicate that redistribution of federal grants that imposes equity of compliance costs could have resulted in a more efficient use of those grant funds. The insensitive response of larger cities to federal grants suggests that such a redistribution could be revenue-neutral. Future research should address the endogeneity of grant receipt to better understand its impacts at alleviating local costs of mandated infrastructure adoption.

1.7.3 Heterogeneous Responses by a priori Potential Benefits

In recent work by Albouy et al. (2018), the authors show that estimates of resident value for public goods can critically depend on their complementarities with other public goods. They demonstrate that public park access, for example, is valuable to residents only if accompanied by improvements to public safety. The central finding of their paper is that ignoring the existence of such complementarities across public goods can significantly attenuate valuation estimates of those public goods. In the context of this paper, valuation of mandated wastewater treatment infrastructure may be biased toward zero if cleaner ambient surface water is valuable to residents only in conjunction with other public goods, such as warm swimming weather or proximity to a large river or lakes.

I estimate a version of Eq. 1.7 to test how my valuation outcomes of population, housing prices, and high skill composition respond differently among cities with warmer summer temperatures, proximity to large water bodies, and for cities more likely to experience water quality improvements from upstream abatement. Figure 1.11 shows the heterogeneity prevalent in each of these regressions.[28] The value residents place on mandated changes to wastewater infrastructure depends upon the municipality's *a priori* potential to benefit from a federal mandate. Cities with warmer July temperatures, that are closer to large water bodies, and that are located closer to the mouth of a river have larger increases in population and housing prices. The July temperature and Distance to Water results are consistent with prior literature that generally finds the salience of water quality improvements is an important determinant of its valuation. The "Distance to River Mouth" point estimates indicated by diamonds provide suggestive evidence that the largest benefits accrued to cities more likely to have experienced water quality improvements from upstream abatement because they are situated near the bottom of a river. This

[28]Tabular results from these regression estimates are shown in Appendix Table A.17.

result should not be interpreted as the spillover benefit because distance to river mouth is not a precise measure of exposure to upstream pollutants. However, it points to a potentially important function of federal mandates in enabling strategic complementarities from environmental regulation. Specifically, the greatest benefits accrue to cities that are not only impacted by their own investment in pollution abatement, but investments of upstream neighboring cities as well. This suggests federal mandates have the capacity to improve efficiency locally when the benefits of public goods provision are intertwined with the decisions of other neighboring cities.

1.8 Conclusion

The American Society of Civil Engineers estimates that infrastructure in the United States requires an investment of at least $3.2 trillion to prevent deterioration of the country's aging roadways, electrical grids, transit, and waterworks (Economic Development Research Group, Inc. 2016). To meet some of these needs, the Trump administration proposed in January of 2018 to use local user fees and state and local financing to fund a $1.5 trillion infrastructure renewal plan. While there is bipartisan consensus that infrastructure renewal is necessary, we know very little about how localized financing for federally-mandated projects affects local economies. The striking gap in our knowledge of these effects matters because understanding who ultimately bears the burden of federal spending mandates and whether mandates are valued locally can alter conclusions of the cost effectiveness of federal policies.

This paper is the first empirical effort to assess the impact of federal mandates on local government budgets and to determine whether mandated provision of goods and services are valued by local residents. Further, this is the first paper to test how the regulatory burden of the CWA technology standard on wastewater treatment impacted municipal governments. My results indicate that local governments relied mainly on user fees to finance the unsubsidized portions of CWA compliance, with a collective increase totaling $300 per year per household or a three-fold increase from pre-CWA levels. Importantly, I do not find evidence that local governments displaced funding from unmandated public goods and services in order to fund mandated infrastructure, even several decades following the CWA legislation. Several of the largest federal mandates - including regulations on solid waste management and drinking water quality - are funded with a fee for service, suggesting that mandate compliance is unlikely to generate distortions in the menu of goods and services offered by local governments. Yet, reliance on user fees for essential, demand-inelastic goods like piped water presents important distributional concerns. In Baltimore, for example, federal demands on renewal of the city's water infrastructure have increased resident water bills so much that the city has repossessed several homes for unpaid water bills (Baltimore Sun Editorial Board 2019).

The local impacts of the CWA infrastructure mandate were associated with statistically significant positive increases to total population. Further, compliance with the uniform mandate induced highly heterogeneous responses. First, I find robust and economically significant evidence of taste-based sorting, with regulated cities experiencing increases in educational attainment. Second, cities more likely to experience pollution abatement from upstream neighbors, closer to large water bodies, and in more temperate climates incurred larger gains in both housing prices and population relative to other cities less likely to benefit from proximate water quality improvements. My results suggest that the externality corrections induced by the CWA mandate are at least valued above their cost by taxpayers at the local level. Future work should assess whether the CWA mandate is overall welfare enhancing after accounting for spillover effects and the benefits from upstream pollution abatement on downstream cities.

Figure 1.1: Inventory of Publicly-Owned Wastewater Treatment Plant Technology

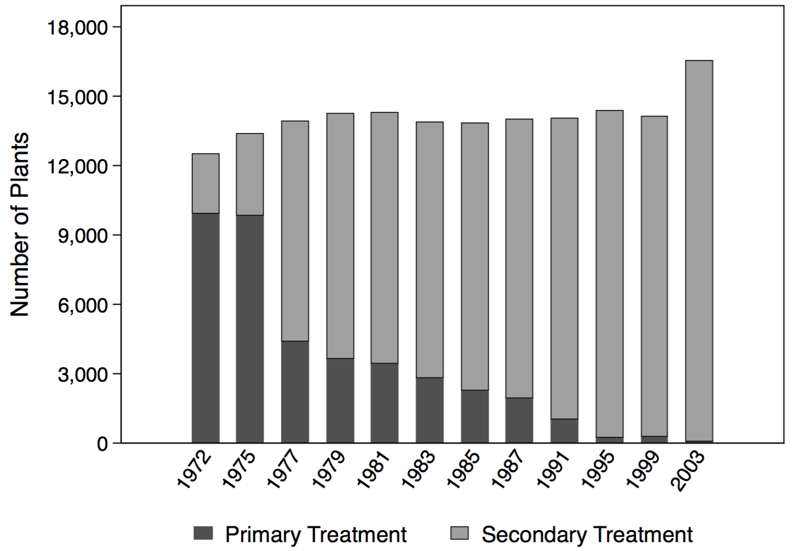

Source: EPA Clean Watershed Needs Surveys (1973-2004). Figure plots the total number of publicly-owned wastewater treatment plants by technology type. Secondary treatment technology was mandated under the 1972 CWA for all surface water-bound effluent. Treatment plants with only primary treatment were required to either cease operations or upgrade to secondary treatment. See text for further details on treatment technology characteristics.

Figure 1.2: Municipal Expenditure Trends

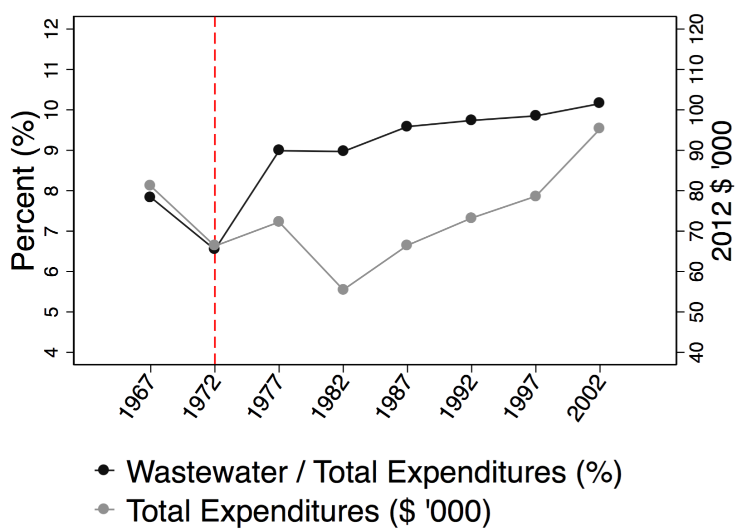

Source: Census of Governments. Figure plots raw means for annual direct expenditures (total expenditures net of intergovernmental payments) in thousands of 2012 USD in gray; and annual share of expenditures in wastewater in black. Dashed line indicates the start of the Clean Water Act in 1972.

Figure 1.3: Federal Mandates and Local Efficiency

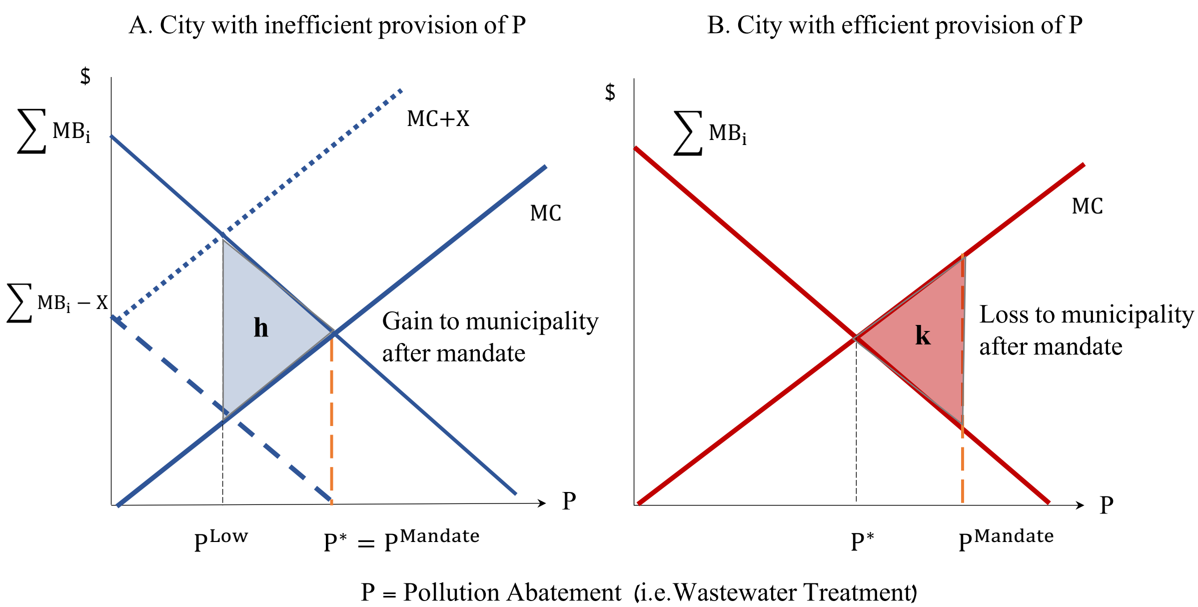

Note: Figure shows the effect of the CWA mandate given locally inefficient versus efficient provision of pollution abatement. City A under-provides pollution abatement in equilibrium due to market failures X which either increase the marginal cost from MC to $MC + X$, or reduce aggregate social marginal benefits from $\sum MB_i$ to $\sum MB_i - X$. For city A, the mandate reduces either source of inefficiency and moves the equilibrium provision of pollution abatement from G^{Low} to G^* and local surplus increases by **h**. City B efficiently supplies pollution abatement in equilibrium. The mandate reduces local surplus by **k**.

Figure 1.4: Capital Cost of Secondary Treatment by Plant Service Population

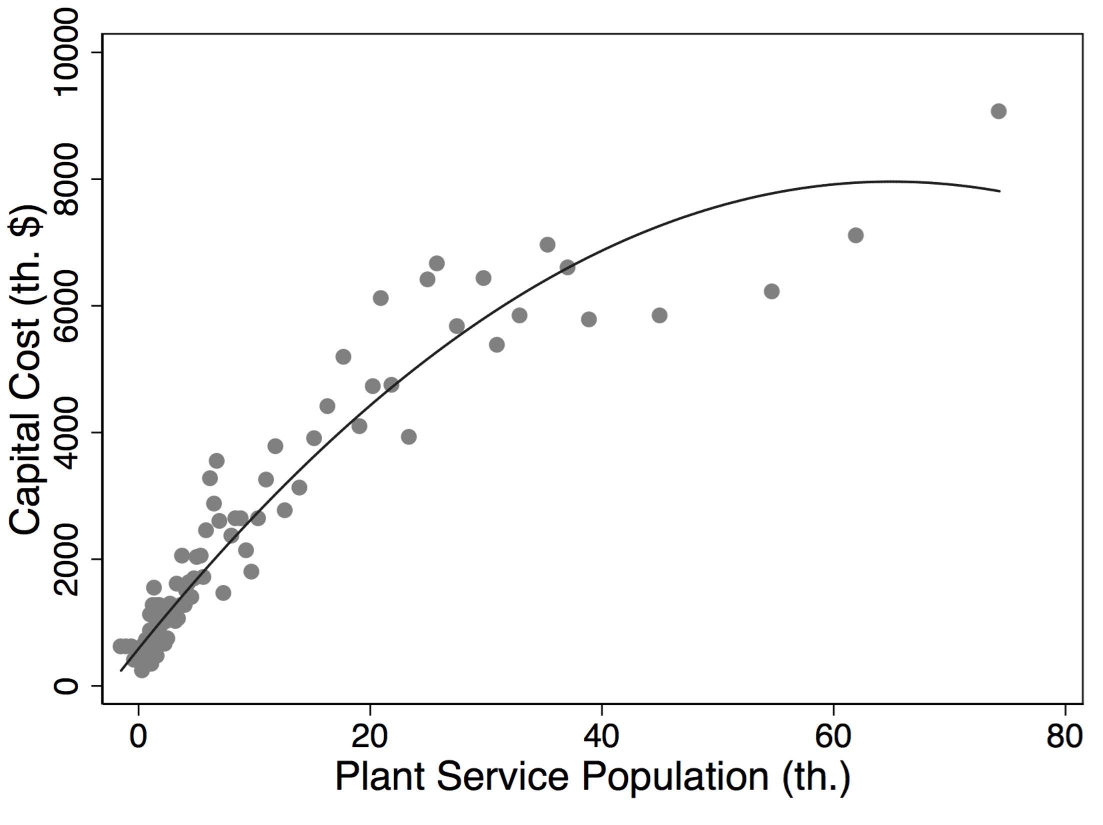

Source: EPA (1973). Figure plots the binned scatterplot and quadratic fit of cost needs for compliance with the secondary treatment standard relative to a plant's service population. Plot divided into 100 equal-sized bins. Residualized by year fixed effects. Sample is based on 48,115 observations, which includes 5,884 treatment facilities from 1975-2003 reporting non-zero secondary treatment cost needs and non-zero service population. Sample excludes the top and bottom 5% of treatment capacity, and plants that appear less than 7 years over the 28-year panel. Plant service population is calculated as plant capacity in gallons per day divided by 100 (Guo et al. 2014). Cost values are in thousands of 2012 dollars.

Figure 1.5: Distribution of Pre-CWA Secondary Treatment Adoption

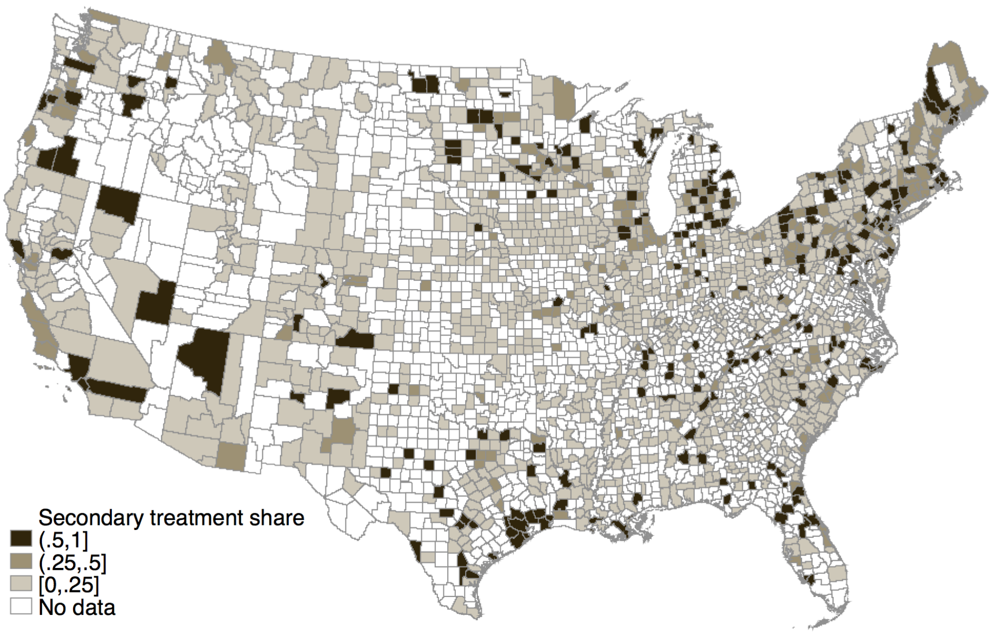

Source: EPA (1973), Census. Figure shows the county share of municipal wastewater treatment plants with secondary treatment as of 1972.

Figure 1.6: Mean Municipal Downstream Population Size by County, 1970

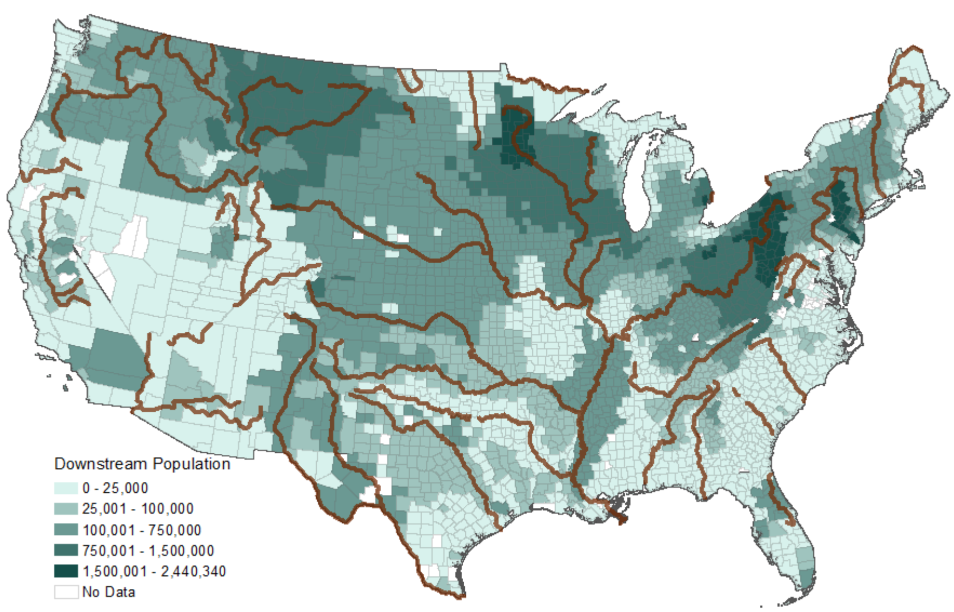

Source: USGS, Census, author's own calculation. Figure shows county-level averages of city downstream population as of 1970.

Figure 1.7: State Composition of Wastewater Treatment Technology, 1972.

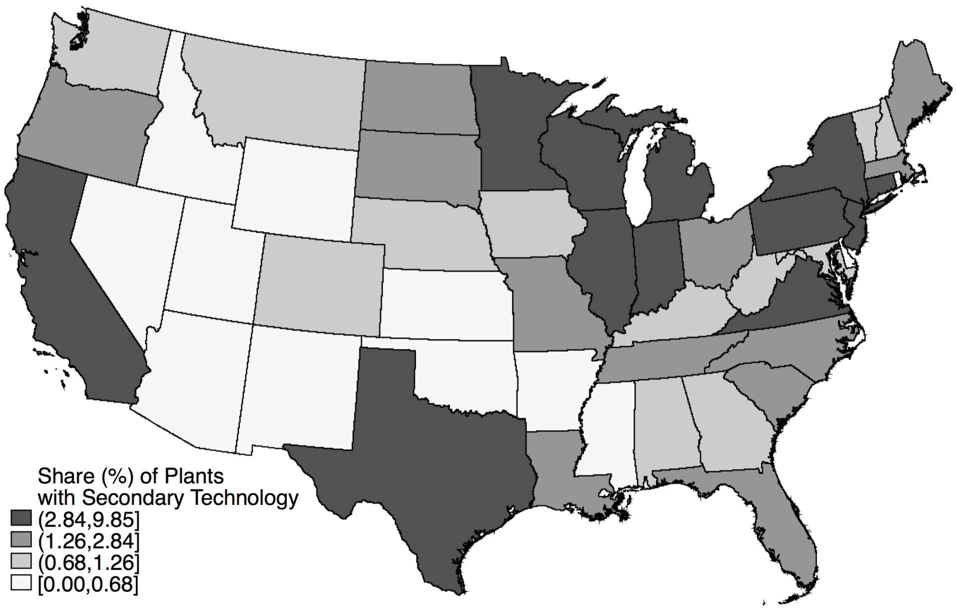

Source: EPA 1973 Clean Watershed Needs Survey. Figure shows the distribution of wastewater treatment plant technology composition by state. Each shade of gray corresponds to a quartile.

Figure 1.8: State Compliant Plant Share (1972) vs Downstream Population

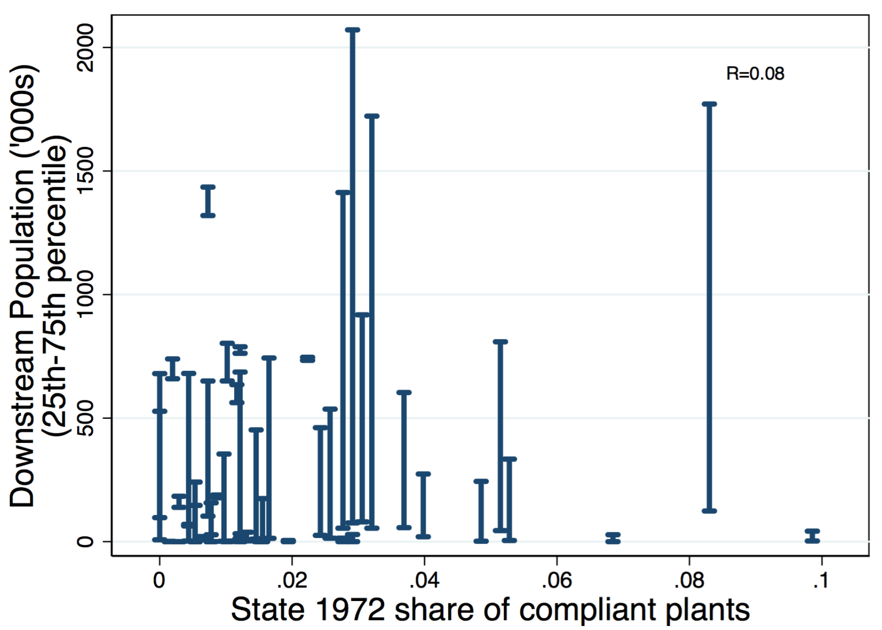

Source: EPA 1973 clean watershed needs survey, USGS, author's calculations. Figure shows the distribution of downstream population (25th-75th percentile) by state share of compliant plants as of 1972.

Figure 1.9: Wastewater Expenditures per capita and Instrumental Variable Variation

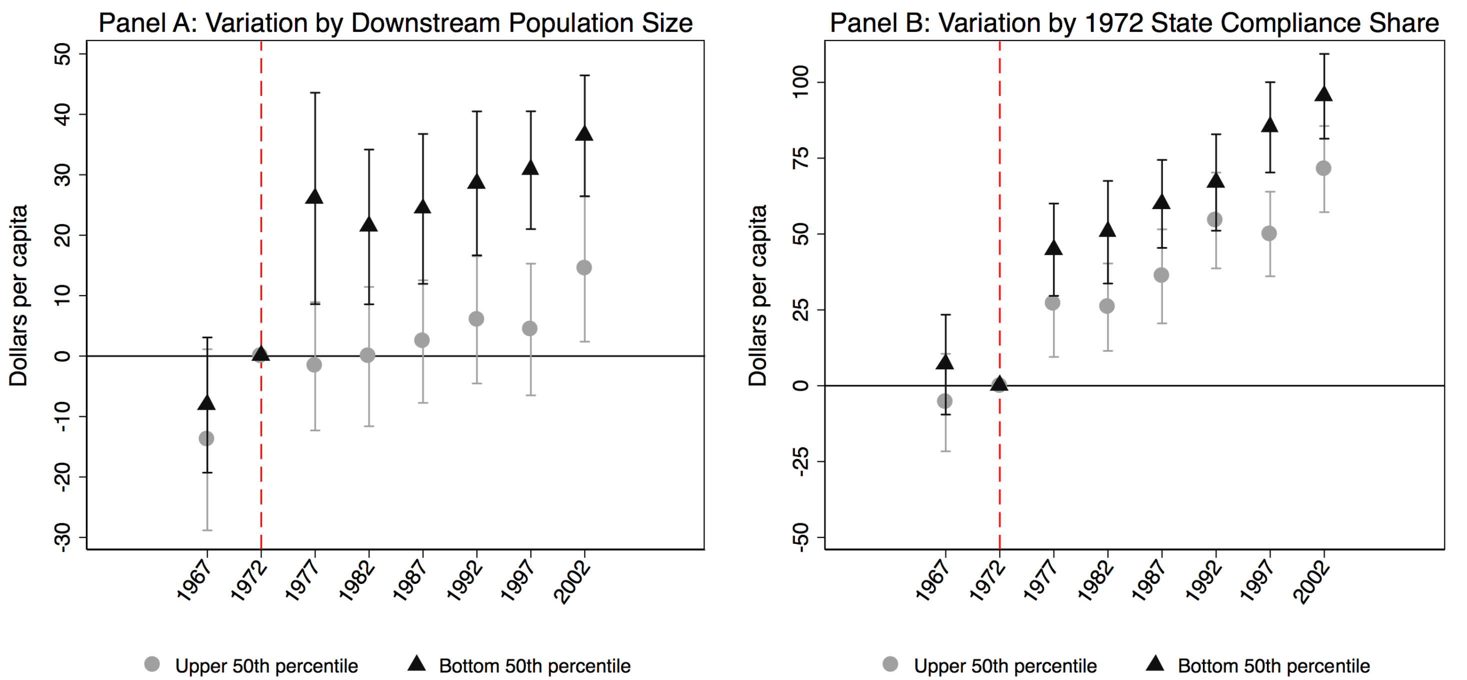

Source: USGS, CWNS, Census of Governments, author's calculations. **Panel A** plots $\delta_{t\times \underline{50}} + \delta_t$ and δ_t from equation: $y_{it} = \sum_t \delta_{t\times \underline{50}}(I_{\underline{50}} \times S_s \times D_t) + \sum_t \delta_t(S_s \times D_t) + (D_i\sigma_t) + \nu_i + \varepsilon_{it}$ where the dependent variable is wastewater expenditures per capita for city i in year t; $I_{\underline{50}}$ is an indicator for a city having downstream population in the bottom 50th percentile, S_s is state share of compliant plants as of 1972, D_t is an indicator for year t, and D_i is downstream population. Bands show 95% confidence intervals. All coefficients are evaluated at the mean state share of secondary treatment plants as of 1972. The reference year is t=1972. Robust standard errors are clustered at the city level. **Panel B** plots $\sigma_{t\times \underline{50}} + \sigma_t$ and σ_t from equation: $y_{it} = \sum_t \sigma_{t\times \underline{50}}(I_{\underline{50}} \times D_t) + \sum_t \sigma_t(D_t) + \nu_i + \varepsilon_{it}$ where the dependent variable is wastewater expenditures per capita for city i in year t; $I_{\underline{50}}$ is an indicator equal to 1 if a city's state has share of compliant treatment plants as of 1972 in the bottom 50th percentile, and D_t is an indicator for year t. Bands show 95% confidence intervals. The reference year is t=1972. Robust standard errors are clustered at the city level.

Figure 1.10: CWA Noncompliance and Dynamics of Water Quality

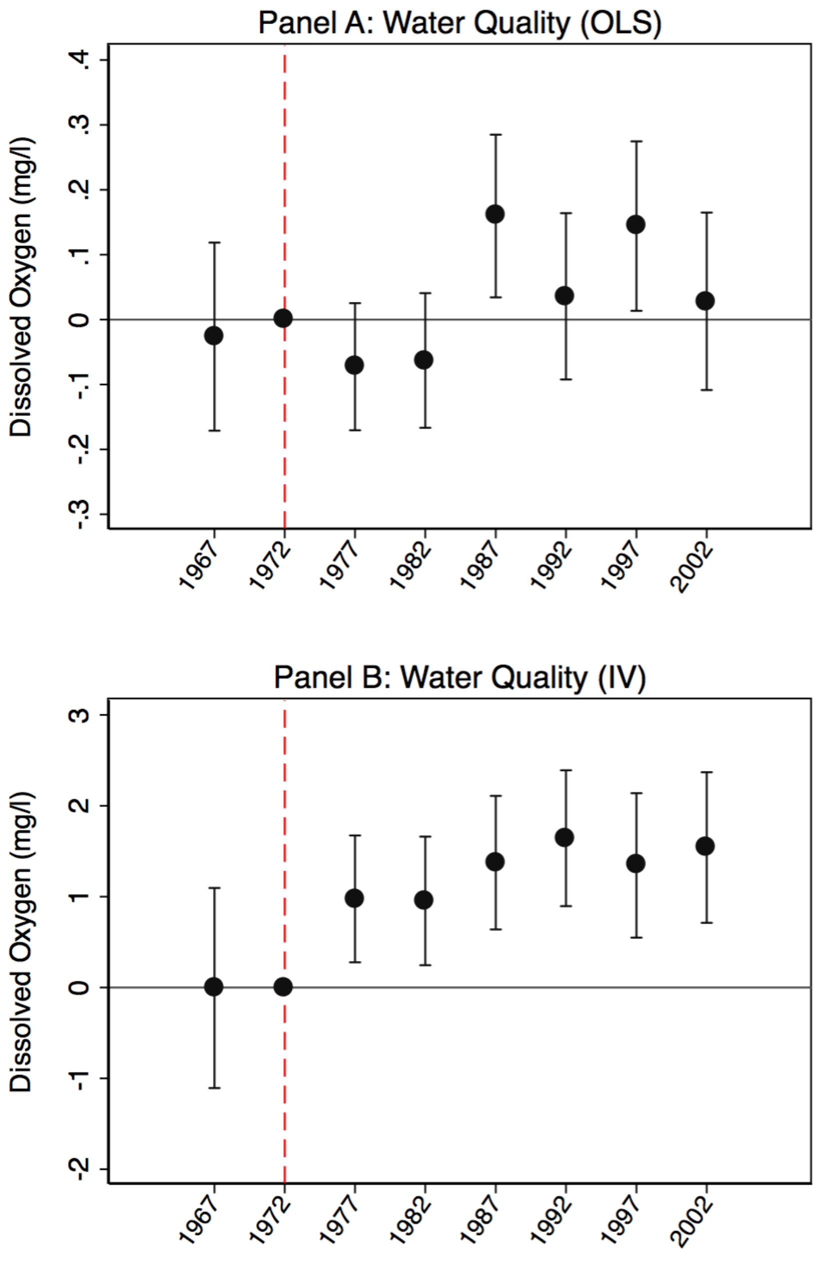

Note: Figure plots δ_t from estimating $y_{ist} = \sum_t \delta_t(P_i \times D_t) + \mathbf{X_i}\theta_\tau + (R \times t) + \gamma_t + \nu_i + \varepsilon_{ist}$. y_{ist} is dissolved oxygen. Panel A is estimated from differences-in-differences; Panel B is estimated from two-stage-least squares. Bands show 95% confidence intervals. All coefficients are normalized relative to t=1972. Robust standard errors are clustered at the city level.

Figure 1.11: Benefits of Mandate Compliance and Complementary Amenities

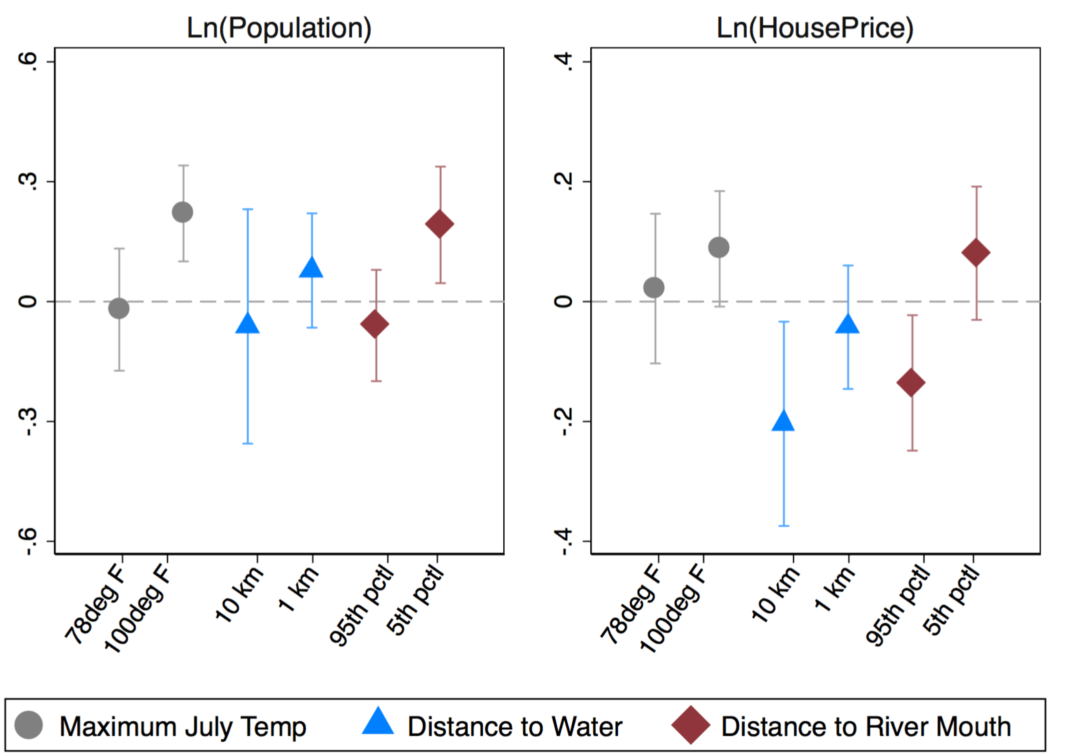

Note: Figure plots estimates of $\delta_{POST} + \delta_{POST \times \omega}$ from Eq. 1.7, where ω is one of three variables: (i) average maximum July temperature from 1970-2000; (ii) kilometers to nearest large water body, where "large" is a water feature with a stream order of 6 or larger as defined on a scale of 1-13 by USGS; and (iii) distance to river mouth in percentiles. ω is evaluated at high and low values of each variable. Standard errors clustered by city.

Table 1.1: Pre-Clean Water Act Descriptive Statistics

	Primary Treatment	Secondary Treatment	P-value for difference in means
City Characteristics			
Population	36,752	30,055	0.439
Median House Price ($)	95,434	103,559	0.000
Share of population with a college degree	0.110	0.114	0.166
Dissolved oxygen (mg/l)	8.11	8.14	0.591
LCV conservation score	43	50	0.000
County-level labor market			
County income per capita ($)	23,071	24,168	0.000
County employment share in manufacturing	0.362	0.386	0.000
County employment share in water polluting manufacturing	0.146	0.152	0.293
Expenditures per capita			
Total expenditures	1,027	1,195	0.000
Wastewater	66	115	0.000
Total other	649	667	0.354
Public works	132	134	0.429
Public safety	384	382	0.886
General & admin.	61	75	0.000
Health & welfare	28	30	0.552
Recreation	43	44	0.639
Revenues per capita			
Total revenues pc ($)	1,014	1,141	0.000
Intergovernment revenues	168	217	0.000
Revenues from own sources	846	923	0.007
Total taxes	387	508	0.000
Property taxes	299	431	0.000
Sales & License taxes	87	77	0.020
Total user fees	111	100	0.095
Wastewater user fees	31	29	0.336
Long-term debt outstanding	1,394	1,368	0.799
Short-term debt outstanding	87	116	0.011
Geography			
River Population as of 1970 (th.)	8,577	5,558	0.000
Distance to waterbody (km)	0.057	0.021	0.002
Distance to river mouth (km)	1,276	828	0.000
Distance to navigable river (km)	207	193	0.084
Distance to Great Lake (km)	735	663	0.004
Distance to Ocean (km)	525	388	0.000
Number of Cities	2,290	685	
Panel Frequency	5.2	5.4	
Observations	11,320	3,546	

Note: All variables measured as means in 1967 and 1972. P-value denotes significance of difference in means. Dollars in USD 2012 values. See Section 1.4 for details on data sources.

Table 1.2: Determinants of *Ex Ante* CWA Compliance Status

	Cross Section (1972)		Panel	
	(1)	(2)	(3)	(4)
Downstream Population	-0.049***	-0.057***		
	(0.012)	(0.014)		
Dowstream Population x StateShare'72 x Post			-0.704**	-0.593*
			(0.322)	(0.330)
Dowstream Population x Post			-0.003	-0.011
			(0.019)	(0.019)
StateShare'72 x Post			-2.066***	-2.345***
			(0.347)	(0.372)
River FE	Y	Y		
Geography Controls		Y		
City Controls		Y		
Year & City FE			Y	Y
RiverPopulation x YearFE			Y	Y
CityControls x YearFE				Y
Income Trend				Y
Region Trend				Y
Baseline mean	0.75	0.75	0.75	0.75
Pct δ in y for 1 SD increase DSpop	-6.58%	-7.56%	-3.20%	-2.70%
F-statistic	15.64	5.84	19.22	20.11
Observations	2151	2151	14866	14866

Note: The dependent variable in (1) and (2) is an indicator for primary treatment as of 1972. In (3) and (4), this is interacted with a post CWA indicator. Standard errors clustered by city. Geography controls include distance to river edge and distance to river mouth. City Controls include pre-CWA averages from 1967-1972 of: share of county-level employment in manufacturing and water-polluting industries; annual federal, state, and local intergovernmental grants and distance from coastline. Specifications (3) and (4) interact these baseline controls with year fixed effects. Downstream population is normalized by its standard deviation. Region consists of 8 indicators based on BEA US regions: New England, Mideast, Great Lakes, Plains, Southeast, Southwest, Rocky Mountain, and Far West. Income is average county-level income per capita in 1970. Marginal effects evaluated at the mean value of 1972 state share of secondary treatment plants (3.4%). * (p<0.10), ** (p<0.05), *** (p<0.01).

Table 1.3: Pretrends of City Characteristics by Treatment and Instrument

	Secondary Treatment as of 1972		Above Median Exposure to Instrument		Mean of Secondary Treatment
	(1)	(2)	(3)	(4)	(5)
Dissolved Oxygen (mg/l)	-0.069	-0.094	-0.055	-0.071	8.11
	(0.074)	(0.074)	(0.065)	(0.093)	
Ln(Population)‡	0.046***	0.035**	-0.011	-0.028	24,549
	(0.016)	(0.015)	(0.013)	(0.019)	
Total Expenditures pc ($)	71.458	48.132	9.800	-90.762	1032.09
	(46.889)	(46.677)	(41.159)	(57.829)	
Sewerage Expenditures pc ($)	68.133***	60.453***	33.992**	18.221	66.67
	(15.216)	(15.328)	(13.415)	(19.109)	
Other Expenditures pc ($)	-19.102	-18.141	-49.897*	-89.160**	650.54
	(30.941)	(30.975)	(27.105)	(38.321)	
Total Revenues pc ($)	-27.091	-30.600	-67.511	-95.004	1019.02
	(58.585)	(59.460)	(51.351)	(73.667)	
Total User Fees pc ($)	-10.137	-7.767	-5.075	5.793	111.91
	(9.832)	(10.073)	(8.625)	(12.488)	
Wastewater User Fees pc ($)	6.024**	6.496**	-0.414	-4.158	31.41
	(2.521)	(2.560)	(2.216)	(3.179)	
Long Term Debt pc ($)	110.565	141.048	-29.219	71.008	1394.71
	(156.300)	(160.351)	(137.105)	(198.821)	
Short Term Debt pc ($)	29.074	20.770	23.630	-28.144	87.83
	(20.672)	(20.838)	(18.132)	(25.829)	
Total federal grants pc ($)	14.003	17.254	6.682	4.176	29.23
	(10.559)	(10.756)	(9.265)	(13.346)	
Federal Infrastructure grants pc ($)	20.983***	24.446***	0.342	4.583	14.63
	(6.394)	(6.337)	(5.631)	(7.900)	
Controls		Y		Y	
Observations	2590	2590	2590	2590	

Note: Table provides city-level summary statistics obtained by estimating $f_{irt} = \beta \text{Control}_{ir} + \gamma_r + \tau_t + \epsilon_{ic}$ where f_{irt} is a pre-CWA characteristic for city i in region r. Column (1) reports estimates of β when "Control" equals 1 if a city had secondary treatment as of 1972 (i.e., compliant treatment technology). Column (2) reports estimates of β when "Control" equals 1 if a city has a higher than median probability of having a compliant treatment plant prior to the CWA's adoption on the basis of the instrument, Downstream Population × StateShare'72. Includes pre-CWA years, 1967 and 1972. Federal Construction Grants pc include federal grants for wastewater treatment as well as disaster relief, homeland security, and miscellaneous goods. Controls include all variables listed in Table 1.2. ‡ Population regression includes census years 1950, 1960, and 1970. Controls include only city and year fixed effects, region linear time trends, and time-varying effects of distance to ocean and river network population. * (p<0.10), ** (p<0.05), *** (p<0.01).

Table 1.4: Effect of CWA compliance on Local Government Budgets

PANEL A: DIFFERENCE-IN-DIFFERENCES

Expenditures Per Capita	Total	Wastewater			Other					
		Total	Capital	Operating	Total	Public Safety	Public Works	Gen Admin	Welfare	Rec
Primary'72xPost	59.787	58.715***	54.737***	3.469	-27.342	5.232	-44.154	11.195*	0.982	-0.597
	(42.141)	(9.415)	(9.025)	(2.193)	(35.306)	(3.182)	(33.322)	(5.731)	(4.354)	(2.282)
Marginal effect (%)	5.79%	88.07%	141.62%	12.53%	-4.20%	3.95%	-11.49%	18.24%	3.34%	-1.39%
Baseline mean	1032.09	66.67	38.65	27.69	650.54	132.39	384.35	61.36	29.43	43.02

Revenues Per Capita	Total Revenues			User Fees		Taxes			Debt	
	Total	Inter Govt	Own	Total	Wastewater	Total	Property	Sales & License	Long Term	Short Term
Primary'72xPost	41.352	59.698***	-18.362	17.454*	4.885**	5.626	-4.638	10.195**	34.324	12.547
	(32.089)	(11.040)	(29.178)	(10.345)	(2.461)	(10.311)	(9.254)	(4.425)	(112.620)	(13.652)
Marginal effect (%)	4.06%	34.99%	-2.16%	15.60%	15.55%	1.44%	-1.54%	11.46%	2.46%	14.29%
Baseline mean	1019.02	170.60	848.45	111.91	31.41	389.85	300.87	88.99	1394.71	87.83

PANEL B: INSTRUMENTED DIFFERENCE-IN-DIFFERENCES

Expenditures Per Capita	Total	Wastewater			Other					
		Total	Capital	Operating	Total	Public Safety	Public Works	Gen Admin	Welfare	Rec
Primary'72 x Post	480.927*	151.768***	101.021**	46.471***	241.263	64.237***	136.003	2.286	20.586	18.151
	(255.934)	(50.576)	(46.097)	(15.528)	(213.601)	(21.855)	(203.596)	(21.338)	(24.974)	(15.033)
Marginal effect	47%	228%	261%	168%	37%	49%	35%	4%	70%	42%
Baseline mean	1032.09	66.67	38.65	27.69	650.54	132.39	384.35	61.36	29.43	43.02

Revenues Per Capita	Total Revenues			User Fees		Taxes			Debt	
	Total	Inter Govt	Own	Total	Wastewater	Total	Property	Sales & License	Long Term	Short Term
Primary'72 x Post	691.545***	403.769***	287.513*	-1.178	61.335***	26.508	7.610	18.850	-309.102	76.313
	(182.481)	(79.160)	(153.476)	(63.395)	(19.742)	(51.740)	(42.648)	(27.005)	(834.045)	(71.667)
Marginal effect	68%	237%	34%	-1%	195%	7%	3%	21%	-22%	87%
Baseline mean	1019.02	170.60	848.45	111.91	31.41	389.85	300.87	88.99	1394.71	87.83
First Stage F-statistic	20.11	20.11	20.11	20.11	20.11	20.11	20.11	20.11	20.11	20.11
Clusters	2975	2975	2975	2975	2975	2975	2975	2976	2975	2975
Observations	14866	14866	14866	14866	14866	14866	14866	14866	14866	14866

Note: Dependent variables are in 2012 dollars per capita. Table reports estimates of β from Eq. 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$.

The instrument for $(P_i \times POST_t)$ in Panel B is Eq. 1.4. Standard errors clustered by city. $\mathbf{X}$ includes all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5: Effect of CWA compliance on Water Quality & Municipal Growth

	Panel A: Difference-in-Differences			
	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary'72xPost	0.039	-0.030***	-0.006	-0.005**
	(0.047)	(0.011)	(0.008)	(0.002)
Marginal effect (%)	0.49%	-2.97%	-0.64%	-4.18%
	Panel B: Instrumented Difference-in-Differences			
	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary'72xPost	1.543***	0.121*	0.013	0.027**
	(0.327)	(0.071)	(0.054)	(0.012)
Marginal effect (%)	19%	12.59%	1.12%	25%
First Stage F-statistic	19.58	20.43	20.43	16.16
Baseline mean	8.11 mg/l	36,710	$95,461	11%
Clusters	2933	2965	2965	2329
Observations	14177	8264	8264	6366

Note: Population, housing price, and high skill regressions include decade interval years only (1972, 1982, 1992). Dollars in USD 2012 values. "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Table reports estimates of β from Eq. 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + R_\tau + \gamma_t + \nu_i + \varepsilon_{it}$. The instrument for $(P_i \times POST_t)$ in Panel B is Equation 1.4. Standard errors clustered by city. Includes all controls listed in Table 1.2, column(4). Marginal effect for ln(population) and ln(median house price) calculated as $(\exp(\beta - var(\beta)/2) - 1) \times 100$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.6: Downstream Population Mechanisms

	All Cities Downstream			"Big" Cities Downstream		
	(1)	(2)	(3)	(4)	(5)	(6)
Count of Cities Downstream	-0.001***					
	(0.000)					
Count of Cities Downstream x StateShare'72 x Post		-0.029***	-0.022***			
		(0.008)	(0.008)			
Count of Cities Downstream x Post		0.001***	0.001*			
		(0.000)	(0.000)			
Count of Big Cities Downstream				-0.017***		
				(0.006)		
Count of Big Cities Downstream x StateShare'72 x Post					-0.430***	-0.445***
					(0.156)	(0.162)
Count of Big Cities Downstream x Post					0.004	0.005
					(0.008)	(0.008)
StateShare'72 x Post		-1.516***	-1.870***		-1.466***	-1.760***
		(0.396)	(0.412)		(0.401)	(0.415)
Geography Controls	Y			Y		
Baseline Controls			Y			Y
Year & CityFE		Y	Y		Y	Y
RiverPopulation x YearFE		Y	Y		Y	Y
Baseline mean	0.75	0.75	0.75	0.75	0.75	0.75
Pct δ in y for 1 SD increase Downstream Count	-0.14%	-0.13%	-0.10%	-2.26%	-1.95%	-2.02%
F Statistic	10.627	17.876	17.453	9.408	19.331	19.486
Observations	2151	14866	14866	2151	14866	14866

Note: Dependent variable in columns (1) is an indicator for primary treatment as of 1972. In the remaining columns, this is interacted with a post CWA indicator. "CityCountDownstream" is the count of cities downstream. "BigCityCountDownstream" is the number of cities with 1970 populations over 30,000 downstream.Standard errors clustered by city. Baseline controls include all controls listed in Table1.2 column 4. Marginal effects evaluated at the mean value of 1972 state share of secondary treatment plants (3.4%). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.7: Effect of CWA compliance on Local Government Budgets by City Size

PANEL A: EXPENDITURES PER CAPITA

	Total	Wastewater			Other					
		Total	Capital	Operating	Total	Public Safety	Public Works	Gen Admin	Welfare	Rec
Primary'72xPost	413.561*	177.624***	124.537***	49.013***	237.406	53.771**	149.349	12.871	5.216	16.200
	(230.584)	(50.148)	(45.907)	(14.531)	(185.368)	(22.256)	(174.569)	(20.462)	(22.286)	(14.656)
Primary'72xPostxAboveMedian	66.808	-33.468**	-16.796	-17.117***	-70.991	15.531***	-90.021	-14.517*	16.033***	1.983
	(71.275)	(13.922)	(12.250)	(4.323)	(63.089)	(5.765)	(59.246)	(8.121)	(6.211)	(3.492)
Baseline mean	1032.09	66.67	38.65	27.69	650.54	132.39	384.35	61.36	29.43	43.02

PANEL B: REVENUES PER CAPITA

	Total Revenues			User Fees		Taxes			Debt	
	Total	Inter Govt	Own	Total	Wastewater	Total	Property	Sales & Licenses	Long Term	Short Term
Primary'72xPost	591.233***	359.173***	231.715	6.131	57.561***	32.172	16.208	15.883	-477.680	72.000
	(174.907)	(76.445)	(147.383)	(61.219)	(19.382)	(50.445)	(41.354)	(25.345)	(754.298)	(78.240)
Primary'72xPostxAboveMedian	103.603**	61.055***	42.636	-27.117*	-8.020*	-4.212	-8.914	4.728	454.371***	-44.050**
	(47.914)	(18.502)	(39.972)	(16.353)	(4.759)	(12.496)	(11.500)	(8.551)	(173.885)	(18.277)
Baseline mean	1019.02	170.60	848.45	111.91	31.41	389.85	300.87	88.99	1394.71	87.83
First Stage F-statistic	11.02	11.02	11.02	11.02	11.02	11.02	11.02	11.02	11.02	11.02
Observations	14866	14866	14866	14866	14866	14866	14866	14866	14866	14866

Note: Dependent variables are in 2012 dollars per capita. Table reports estimates of δ_{POST} and $\delta_{POST\times\omega}$ from Eq. 1.7: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST\times\omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + (\overline{N}_i \times \gamma_t) + D_i\sigma_t + (R \times t) + \gamma_t + \nu_i + varepsilon_{it}$, where Equation 1.4 instruments for $(P_i \times POST_t)$. Standard errors clustered by city. $\mathbf{X}$ includes all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.8: CWA Compliance and Heterogeneous Effects by Grant Receipt and City Size

	Wastewater Userfee	Ln(Population)	Ln(Median house price)	High Skill Share
		PANEL A: BELOW MEDIAN POPULATION CITIES		
Primary'72xPost	78.271***	0.055	-0.215**	-0.002
	(28.407)	(0.126)	(0.099)	(0.016)
Primary'72xPostxGrant	-6.823	0.012	0.065**	0.009
	(10.804)	(0.035)	(0.021)	(0.006)
First stage F-statistic	5.55	3.31	3.25	4.07
Baseline mean	$13.63	2,648	$83188.99	9%
Observations	5980	2270	2270	1010
		PANEL B: ABOVE MEDIAN POPULATION CITIES		
Primary'72xPost	18.555	-0.036	-0.115*	0.011
	(18.766)	(0.069)	(0.064)	(0.011)
Primary'72xPostxGrant	0.007	-0.001	-0.001**	0.000
	(0.038)	(0.000)	(0.000)	(0.000)
First stage F-statistic	4.06	4.15	4.15	4.15
Baseline mean	$44.52	59,594	$103,570.1	12%
Observations	5944	2482	2482	2482

Note: Table reports estimates o δ_{POST} and $\delta_{POST\times\omega}$ from Eq. 1.7: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST\times\omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + (\overline{N}_i \times \gamma_t) + D_i\sigma_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$, where ω is annual federal grant receipt, normalized by its standard deviation, and includes federal grants for wastewater treatment, as well as disaster relief, homeland security, and miscellaneous goods. Eq. 1.4 instruments for $(P_i \times POST_t)$. Primary'72xPostxGrant are evaluated at the average annual grant receipt, in standard deviation units. Sample includes years 1967-1987. Population, house price, and high skill share regressions include only years 1972 and 1982, which correspond to decennial census years 1970 and 1980. Standard errors clustered by city. Includes all controls and fixed effects listed in Table 1.2 column 4. $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$

CHAPTER 2

THE EFFICIENCY OF LOCAL GOVERNMENT: THE ROLE OF PRIVATIZATION AND PUBLIC SECTOR UNIONS

2.1 Introduction

In 2012, 19.3 million people (13% of the total workforce) worked in state and local government. This sector encompasses labor-intensive tasks such as K-12 education, garbage collection, fire protection and providing public transit services. Dating back at least to Baumol (1967) economists have argued that productivity growth is slow in labor-intensive industries. In the absence of public sector productivity gains, rising costs for public services mean that local taxpayers face a higher tax burden.

A challenge in studying the efficiency of the public sector arises because key indicators such as street safety or local school quality are both a function of the local population's composition and the local government's ability to deliver quality services. Evaluating public sector efficiency is further complicated by the absence of a universal metric that can provide meaningful comparisons of local governments, either across localities or over time for the same locality. Experimental or quasi-experimental studies permit evaluation of government service efficiency, however only a few examples exist in the literature.[1]

[1] Chong et al. (2014) perform a cross-country evaluation of government service efficiency using a natural experiment on postal services. Cellini et al. (2010) study the efficiency of local tax policy across school districts in California using a regression discontinuity design.

In this paper, we argue that the operating cost of moving a public transit bus one mile represents a standardized metric for ranking local government service efficiency. A distinctive feature of bus travel is its simple production function featuring three key inputs: a driver, a bus, and fuel. The bus and the fuel needed to move the bus are traded on a national market. Labor represents over 80% of the cost. Bus transit provides an informative lens through which to investigate public sector productivity for three additional reasons. First, because multiple agencies may operate within the same city, bus transit provides a unique opportunity to compare public and private service provision within the same locality. Second, bus transit is a ubiquitous public service found in both urban and rural locations alike and across all U.S. regions.[2] Third, fixed-route bus transit forms the core of public transportation in the U.S. Out of 10.6 billion passenger-trips made on mass transit as of 2013, bus transit was the most frequently-used mode of transit accounting for 50%, while heavy rail accounted for 36%. Low-income households are less likely to own vehicles and many do not live near viable heavy or light rail transit. These individuals often use the bus to commute within America's cities. Thus, higher bus transit operating costs could lead to inadequate provision of an important public service for low-income households.

A transit agency can have a high operating cost per mile for two main reasons. It can pay bus drivers and mechanics more than their local opportunity cost or it can hire too many drivers and mechanics relative to the efficient level of employment. For instance, the median hourly wage paid by the Chicago Transit Agency in 2014

[2] While just 13 metropolitan areas offer heavy rail systems (i.e., commuter trains or subways), 342 metropolitan areas across the U.S. have at least one bus transit system (APTA 2015; FTA 2016).

to bus drivers was \$32 and the 25th percentile was \$24 per hour. Private sector hourly wages in Chicago for comparable work were \$12 and \$8, respectively.[3] While a for-profit firm in a competitive industry would have sharp incentives to engage in cost minimization, non-profit public transit agencies have weaker incentives to do so as they face pressure from unions and local political leaders.[4] Prior research has shown that unions have negligible impacts on economic outcomes of employers in the private sector (DiNardo and Lee 2004). In contrast, our study demonstrates that public sector unions are a key determinant of public transit operating cost differentials across space and time.[5]

Privatization represents one strategy for checking local union power and improving transit agency productivity. A transit agency that privatizes a fraction of its bus miles is outsourcing a contract to private sector firms. This action (and even the threat of this action) could discipline public sector unions and is likely to enhance public sector efficiency. Using variation in privatization levels from nationally representative data on public transit, we provide new estimates of the effect of pri-

[3] Based on 2013 IPUMS CPS data for full time, private-sector workers, ages 25-64, living in the Chicago metropolitan area, with at least a high school diploma, but no bachelor's degree, with the occupation "Motor Vehicle Operators".

[4] Winston (2010) provides a review of several institutionalized protections that inhibit labor production efficiencies in the transit sector. These include powerful transit worker unions, excessive administrative staffing in order to fulfill federal bureaucratic obligations, and large transit subsidies which discourage efficiency improvements. Lave (1991) further documents excessive administrative staffing mandated by subsidizing-government institutions. Winston (2000) reviews how policymakers in charge of public entities tend to be responsive to political influences rather than market influences at the expense of efficiency.

[5] Brueckner and Neumark (2014) document that public employees extract rents from local taxpayers in areas with attractive natural amenities. Public sector wages rise in absolute terms relative to private-sector wages in the presence of such amenities. This relationship is stronger for unionized public-sector workers. Diamond (2017) finds similar results in areas with inelastic housing supply.

vatization on the cost of government service provision. We address the fact that privatization is an endogenous choice by employing a regression discontinuity (RD) design based on political elections following the strategy introduced by Lee (2008).[6] Since we compare government service provision at the local level, our approach follows Ferreira and Gyourko (2009) where we use closely-won mayoral elections as a source of quasi-random treatment assignment in the level of privatization.

Our results from the RD design show that bus transit agencies experience substantial cost savings from privatization: transit agencies can reduce per-mile operating costs by nearly 70% from fully privatizing the service provision, holding other factors constant. In addition to testing for the average effect of privatization, we also document that privatization has larger cost-saving effects for a city's dominant transit agency and in areas with strong union power. We test whether these cost savings are a result of service reductions, both in terms of the level of service or the quality of service, and find no evidence of either channel after privatizing. The supply of service and the quality of service are not indistinguishable across publicly relative to privately operated transit agencies. These results highlight that the public sector cost premium cannot be easily explained by efficiency wages because publicly and privately operated agencies in the same city can have significantly different operating costs for the same unit of service. Our findings are consistent with private enterprises operating at a lower cost and higher efficiency relative to their public counterparts. Private entities hire fewer workers and engage in less "featherbedding" relative to

[6]Lee (2008) establishes that as long as there is some inherent uncertainty in the final vote count of a political election with candidates from multiple parties, treatment status (i.e., political partisanship) is as good as random in a small neighborhood of the winning margin threshold.

public transit agencies.

We conduct several robustness checks to examine threats to the RD design including reverse causality and omitted variable bias. The results suggest that privatization is the main channel through which transit agencies with Republican mayors reduce operating costs. Further, Republican and Democrat mayors in closely-won elections do not engage in significantly different cost-cutting behaviors, other than their propensity to privatize transit services. Alternate identification strategies, including an instrumental variable approach, and a difference-in-difference estimator exploiting within-city cost variation support the RD result that privatization directly reduces operating costs by over 70%.

After estimating the effect of privatization on the cost of bus service provision, we quantify the deadweight loss generated by public provision of high-cost services. Larger cities with strong union power experience substantial losses in consumer surplus from costly public transit. Our estimates suggest that fully privatizing all bus transit would generate cost savings of approximately $5.7 billion, or 30% of total U.S. bus transit operating expenses. The corresponding increased use of public transit from this cost reduction would lead to a gain in social welfare of $524 million, at minimum, and at least 26,000 additional transit jobs.

Our study finds that the political affiliation of city mayors strongly predicts the local level of public service privatization. Specifically, the propensity to privatize public services at the local level increases when a Republican mayor wins office. In contrast, Ferreira and Gyourko (2009) conclude that mayoral partisanship does not affect the extent or allocation of government spending at the local level. Our conclusions differ from that ofFerreira and Gyourko (2009) for two important reasons: first, our analysis considers the "intensive margin" of public spending (i.e., costs per unit of service) as opposed to their "extensive margin" (the i.e., the size of government spending). Unit cost reductions for public services do not necessarily imply cost reductions at the aggregate level. For example, is it quite plausible that a newly-elected city mayor outsources certain public services in order to lower costs, and then allocates the realized cost savings to other public services, or to improving technologies of the privatized service, such that the net change in aggregate spending is minimal. Second, while Ferreira and Gyourko (2009) consider mayoral elections from 1950 to 2005, our study examines a more recent time period, 1998 through 2011, a period of growing political polarization.[7]

Recent work by Bel and Rosell (2016) compares operating costs across privately and publicly-operated transit agencies in Barcelona, Spain. Unlike their work, this paper identifies cost differentials without imposing that the privatization decision of agencies is orthogonal to their cost-generating behaviors. Our significant cross-sectional, and time-wise variation allows us to control for selection endogeneity. We find selection induces a substantial negative bias on the cost-saving effect of privatization. Further, our study provides nationally representative estimates for all U.S. cities with public bus transit.

[7]Several studies have documented increases in political polarization in the U.S. over the past two decades, including: Abramowitz (2010); Ash et al. (2015); McCarty et al. (2016); and Gentzkow et al. (2016).

The remainder of the paper is organized as follows. Section 2.2 describes the relationship between labor unions and transit operating costs across the U.S. Section 2.3 presents our theoretical model of transit service provision and our empirical identification strategy, respectively. Section 2.4 describes the data, and Section 2.5 establishes the relevance of our RD design. Section 2.6 discusses our empirical results. We provide estimates of social welfare losses in Section 2.7, and Section 2.8 concludes.

2.2 Spatial & Temporal Variation in Operating Costs of Bus Transit

We use vehicle miles traveled (VMT) as the relevant unit with which to standardize the variables of interest, operating cost and privatization levels. The cost per mile is calculated as total annual operating costs divided by total annual VMT.[8] Other possible operating statistics available through NTD include annual vehicle revenue miles, annual number of passenger trips, and annual passenger miles traveled (PMT). We focus on VMT for several reasons. First, unlike passenger trips or PMT, VMT is strictly a mechanical measure of bus transit and does not conflate unobservable demand-side factors such as service quality or route desirability with the cost of providing bus transit. Second, vehicle revenue miles exclude miles in operator training, maintenance testing, or deadhead, thus do not fully capture the operational costs. Finally, normalizing by VMT ensures that transit agencies serving differently sized

[8]See Appendix B.1 for details on the line items of operating costs.

populations are comparable. If we were to standardize costs by PMT, for example, transit agencies in larger cities would appear to have lower operating costs by construction of serving a larger number of patrons for every bus mile.

The cost of operating a public bus varies greatly across U.S. cities. Table 2.1 compares the operating cost per mile of public bus transit across the twenty largest urbanized areas in the U.S. in 2012. The unit operating cost ranges from $5.91 per mile in San Antonio to $18.67 in New York City. A large part of the variation is driven by differences in wages paid to drivers and mechanics.

Figure 2.1 compares the average operating costs per mile over time for transit agencies in weak bargaining states versus strong bargaining states.[9] Costs per mile are consistently higher in strong bargaining states. Figure 2.2 illustrates that the distribution of labor costs per mile in weak bargaining states is consistently lower than that of strong bargaining states. While the dispersion of the labor cost per mile in weak bargaining states has decreased over time, the dispersion in strong bargaining states has increased over time.

In contrast, privatization shares are consistently lower in strong bargaining states.

[9]State bargaining rights data are sourced from Freeman and Valletta (1988). The level of a state's collective bargaining provision is coded as values ranging from 0 (no state laws relevant to bargaining rights) to 6 (state law dictates employer is obligated to negotiate and come to written agreement with unionized public employee). We classify states with strong bargaining rights laws are those where legislative mandate either implicitly or explicitly dictates public employers and union employees must come to an agreement on contract negotiations (values 5 and 6). States with weak bargaining rights laws either prohibit collective bargaining all together, or do not mandate that the public employer bargain with unions (values 4 or less). Courts have typically interpreted an absence of provision for collective bargaining (value of 0) as prohibiting collective bargaining (Freeman and Valletta 1988), thus we include states classified as level 0 as weak bargaining rights states.

We define a transit agency's privatization share as the ratio of privatized VMT to total VMT. Table 1 illustrates how the average privatization share for the twenty largest urbanized areas is negatively correlated with the cost per mile. San Diego, for example, was 63% privatized as of 2012, and had a cost per mile of $7.09. San Jose, on the other hand, which has a similar cost-of-living index, had a cost per mile nearly 56% higher than that of San Diego at $12.63, and this agency's bus transit is operated almost completely in-house (the privatization share is 1.3%).

Figure 2.3 displays the growth in privatization shares, comparing strong and weak bargaining states. Strong bargaining rights states have consistently lower privatization share. Taken together, Figures 2.1 through 2.3 illustrate a paradox whereby transit agencies with higher costs are less likely to privatize. Union power appears to simultaneously increase transit agency costs and limit the degree to which the transit agency can outsource their operations.

For some major cities including Boston, Chicago, Denver, and Houston, we were able to access their administrative salary databases. These data provide the count of drivers and their salary or hourly wage. In Table 2.2 we report the empirical distributions of current hourly wages for these major transit agencies, as well as descriptive summary statistics on employee utilization, union presence and mean home prices in each city. Boston and Chicago have stronger union presence relative to Houston or Denver because the share of unionized workers in these urbanized areas are higher and because Massachusetts and Illinois are non-right-to-work states. Cities with a stronger union presence have more employees per VMT and higher wages paid to transit agency employees. The employees earning the 25th percentile wage in Boston and Chicago earn more than the 75th percentile wage in Houston or Denver. It is unlikely that housing costs entirely explain this divergence in pay, as Chicago has lower average home prices than Denver.

2.3 A Model of Transit Service Provision and Privatization

In this section, we present a model of input choice in the production of public transit services and discuss mechanisms that can generate cost differences across cities for similar public services. Each transit agency is required to forecast aggregate demand for its services and then to prepare to supply these miles by purchasing buses, and fuel, and hiring drivers and mechanics. We model the manager's decision as choosing operating inputs to minimize the total operating cost conditional on capital and fleet inputs.[10] Managerial decisions on capital and fleet procurement are less frequent relative to operating input decisions of labor, maintenance, and fuel. While Li et al. (2015) study the durable bus investment decision and its implications on energy efficiency for transit agencies, we focus on the labor input decisions and productivity under the influence of labor union rules.

[10]Transit agencies do receive federal subsidies for capital costs but their operating costs are generally not subsidized (Li et al. 2015).

2.3.1 Optimization & Input Decisions

A transit agency provides service (i.e., vehicle miles) using three essential inputs: the physical input (bus), fuel, and labor. We follow Berechman (2013) and posit a constant returns to scale production technology. The production function is:

$$Q = \min\{g_k(K, \eta^k), g_l(L, \eta^l), g_f(F, \eta^f)\} \tag{2.1}$$

where Q denotes the amount of bus service demanded, K the number of buses, L labor input and F fuel input. η^k, η^f and η^l are productivity shocks, including both those observed and unobserved to researchers that affect how efficiently these inputs produce the service. The function $g_k(.)$ translates buses into effective capital (i.e., buses in full operation). The function $g_l(.)$ translates labor inputs (i.e., number of full-time equivalent employees) into effective (i.e., non-idle) labor. The function $g_f(.)$ translates the fuel input into energy used to propel the bus and is affected by bus fuel economy, engine type, and unobserved shocks such as driving conditions.

A transit manager first chooses the physical input K, procured through capital investment, to meet the expected local demand for bus services. Next, they decide the fuel and labor inputs to operate the buses while observing shocks, η. We focus on the operation stage where the transit agency manager chooses inputs to minimize the operating costs while meeting the demand of their service area.

The static optimization problem can be defined as:

$$\begin{gathered} \min\nolimits_{L,F} W * L + P * F \\ \text{s.t. } \min\{g_k(K, \eta^k), g_l(L, \eta^l), g_f(F, \eta^f)\} = \overline{Q}, \end{gathered} \tag{2.2}$$

where W is wage and P is fuel price. We assume the following functional forms for $g_l(L, \eta^l)$ and $g_f(F, \eta^f)$:

$$\begin{aligned} g_l(L, \eta^l) = L * \exp(\eta^L) \\ g_f(F, \eta^f) = F * \exp(\eta^f) \end{aligned} \tag{2.3}$$

The optimal fuel and labor inputs can be written as:

$$\begin{aligned} L^* = \overline{Q}/\exp(\eta^L), \\ F^* = \overline{Q}/\exp(\eta^f), \end{aligned} \tag{2.4}$$

The total cost function can be expressed as:

$$TC^* = \overline{Q}[W/\exp(\eta^L) + P/\exp(\eta^f)] \tag{2.5}$$

Conditional on the capital investment decision (i.e., the bus type), fuel productivity shocks, η^f, are mainly dependent upon congestion and other local driving conditions. These forces are arguably out of the transit agency manager's control and are treated as exogenous to their cost minimization decision.[11] Labor productivity shocks η^L, on the other hand, are dependent upon local labor and political conditions which the transit agency manager may be able to counteract through the use of private contractors. Labor input decisions are the major component of a transit agency manager's cost minimizing problem.

Wages for bus transit workers depend upon local market wages, as well as public sector union strength. There are two main avenues through which unions increase the costs of transit service provisions. First, the majority of transit union labor contracts place substantial limits on use of part-time workers. Contracts either stipulate minimum eight-hour shifts, or require that a minimum ratio of all runs be "straight

[11] The transit agency manager may decide to alter bus routes in order to reduce fuel costs, but must maintain service at the desired quantity, $\overline{Q}$.

runs" as opposed to "split runs."[12] These rules work in direct opposition to the heavily peaked demand of transit service. During midday lulls, workers may be paid even when they are not driving. On the other hand, if a driver works more than an eight-hour shift - extending between morning and evening peak demand - the additional hours are compensated as over-time pay.[13] If managers are able to negotiate for part-time employees, they must pay concessions in the form of wages or benefits that often outweigh gains from lower labor utilization (Giuliano and Lave 1989). By limiting the use of part-time labor, public sector union strength enters into the η^L term in Eq. 2.2, and inhibits the translation of employees into fully-utilized, effective labor. The second avenue through which unions increase the cost of service provision is through pensions and fringe benefits. The costs of union workers' pensions can amount to 50% of their direct wage bill (Black 1991; DiSalvo 2010). Higher unit costs serve to increase W in Eq. 2.5.

A unionized transit agency is likely to pay a higher wage per hour and to have more drivers and mechanics on payroll than would be predicted by the cost-minimizing decision under competitive labor markets. The net effect is a higher average cost per mile of bus service.[14] Figure 2.1 illustrates the positive correlation between average unit cost of bus service and state union strength.

[12]"Split runs" partition driver duties into multiple pieces for a given route, as opposed to a "straight run" where a driver works a continuous eight-hour day. (Deprez 2013).

[13]In 2010, over time costs paid to New York City's MTA employees amounted to 13% of payroll, or the equivalent of employing 7,000 additional full-time workers (Deprez 2013).

[14]Shleifer (1998) review of related empirical work finds similar differences across government vs. privately-run firms: government-run firms hire more workers, have higher operating costs, and lower productivity and profitability than their privately-run counterparts.

Now that we have modeled the cost of producing bus miles under alternative local rules, we introduce the privatization decision, which is one strategy transit agency managers can use to counteract union influence and cut costs.

2.3.2 The Privatization Decision

The ultimate goal of privatization by transit agencies is to reduce costs, especially in the face of a tight budget. Many cities use privatization as a way to reduce costs for various types of services.[15] Local and state level economic conditions affect the revenue base that provides the majority of the operating funding for bus services. Therefore, sluggish economic growth could provide the impetus for transit agencies to explore the cost-cutting options such as privatization, especially in areas more vulnerable to macroeconomic fluctuations such as those in a weak local housing market. A strong public union presence can increase the unit cost of public service provision, but can also make it harder to privatize.[16] Opponents of privatization often argue that the public sector provides steady, well-paying jobs for middle class minorities and that the service from private contractors can be less safe, less reliable

[15] Levin and Tadelis (2010) examine the privatization decision of city services by examining the probability of privatizing as a function of service- and city- specific characteristics. They conclude that services for which it is harder to measure and monitor performance are less likely to be privatized. In their empirical work, they find that cities which are larger, newer, or in the western part of the U.S. are more likely to contract out public services to private providers. They also find a negative correlation between city expenditure per capita and privatization of public service.

[16] In a study on the impacts of state legal environment on county-level privatization of public services, Lopez-de Silanes and Vishnyn (1997) find prohibition on public employee political activity and low unionization encourage privatization of public services.

and, ultimately, more expensive.[17]

The decision to privatize transit operations originates either with agency management or with a governing body of the agency, such as a county board of supervisors. The agency seeks bids from multiple private competitors. Contract lengths generally span one to three years, after which the transit agency management again seeks bids from competitors (Iseki et al. 2005). Most often, the transit agency pays the private firm a negotiated fixed rate per unit of service delivered.

The privatization of transit service routes can include a wide range of contractual arrangements between the transit agency and private contractors. Contract specifications can range from managerial or maintenance assistance alone, to full "turn-key" relationships, where the contractor performs all essential roles including financial management, procurement, marketing and scheduling on behalf of the public transit agency. In most instances, private firms are contracted to manage personnel who operate and maintain the buses for a subset of service routes offered by the public transit agency. In some instances, private firms are contracted to operate all routes offered by the public agency. The private firm is in charge of hiring, compensating, and scheduling employees, as well as negotiating labor contracts with union representatives. Transit agencies retain control over key policy decisions, including service levels, fares, annual operating plans, and contractual compliance. The transit agency is also in charge of budgeting and financing operations, and maintains ownership of all equipment, vehicles, and facilities.

[17] See, for example, "Public-Sector Jobs Vanish, Hitting Blacks Hard" by Larry Hanley, President, Amalgamated Transit Union International, Washington.

Most private contractors of bus transit are national or multinational firms.[18] While market forces incentivize private firms to keep costs at competitive levels, private firms further benefit from economies of scale in ways that public providers cannot. Accumulated experience garnered from managing several transit agencies at once and negotiating with labor unions enable private firms to employ labor and other operating inputs in a more cost-efficient way than their public counterparts. For example, the demand for bus service is highly peaked in the morning and evening commute hours. Private contractors can hire part-time workers to meet the demand during peak hours while public transit agencies have more limited ability to negotiate for part-time labor use in the face of public sector unions.

2.3.3 Empirical Specification

Based on the model of service provision in Section 2.3.1, we specify the productivity shocks to labor η^l in Eq. 2.4 as a function of privatization and other factors:

$$\eta^l = -\beta d - \mathbf{x}\theta - \varepsilon \tag{2.6}$$

where d is a variable characterizing the level of privatization from 0 (no privatization) to 1 (full privatization of all bus service). $\mathbf{x}$ are other observed covariates that affect productivity and ε are productivity shocks unobserved to researchers.

[18] The three largest private contracting companies, Veolia, First Transit, and MV Transit collectively accounted for 65% of all US public bus transit contracts in 2013 (NTD.gov). Veolia is a French company that operates worldwide. MV Transit is a US-based company with international operations. First Transit is a US-based company that operates in the US and Canada.

With the optimal labor inputs and productivity shocks from Eqs. 2.4 and 2.6, we can write the unit labor cost, l, as the following:

$$l = L^* * W/\overline{Q} = W/\exp(\eta^l) = W/\exp(-\beta d - \mathbf{x}\theta - \varepsilon) \tag{2.7}$$

Performing a log transformation on the input cost equation yields:

$$\ln(l) = \ ln(W) + \beta d + \mathbf{x}\theta + \varepsilon \tag{2.8}$$

Denote a transit agency by i and year by t, we rewrite Eq. 2.8 as:

$$\ln(l_{it}) = \beta_0 \ln(W_{it}) + \beta_1 d_{it} + \mathbf{x_{it}}\theta + \tau_t + \eta_s + \varepsilon_{it} \tag{2.9}$$

Eq. 2.9 provides the basis for our empirical analysis. In addition to wages, privatization share, and UZA and agency controls, our empirical specification includes a set of year fixed effects τ_t to control for common macroeconomic trends across transit agencies and state fixed effects η_s to control for time-invariant shocks within a state.

2.3.4 Identification

The decision for a transit agency to privatize some (or all) of their operations is likely to be influenced by local characteristics as well as expectations about future operations and service demand, each of which may not be fully captured by $\mathbf{x_{it}}$ and may affect agency unit costs, l_{it}. The inability to control for such factors will render ε_{it} to be correlated with the treatment, d_{it}, which will bias estimates of β_1. To address this identification challenge, we employ an RD design by using data on mayoral election results and political partisanship.[19] Local partisanship impacts public spending and the propensity to privatize public services (Gerber and Hopkins 2011; Richmond 2001). All else equal, a Republican mayor is more likely to engage in privatizing compared to a Democratic mayor.[20] We exploit the discontinuity in treatment levels generated by narrowly won mayoral elections as a source of pseudo-random assignment of privatization.

In our RD framework, whether the winning margin of the Democratic candidate (i.e. the running variable) falls on one side of a fixed cutoff or the other partly determines treatment intensity. We define the Democrat winning margin as the share of votes received by the Democratic candidate less the largest share of votes received by any other candidate in the election (the "runner-up"). A Democrat winning margin of 0% is the fixed cutoff determining the political party of the winner. Defined in this way, the Democratic candidate wins the election if and only if the Democrat winning margin is positive.[21] The discontinuity we use is fuzzy because the probability of privatization jumps by less than one at the Democratic winning margin threshold of 0%; mayoral partisanship is only part of the factors that affect the probability that a city will privatize it's public transit. As we will demonstrate,

[19] In Appendix B.3, we also provide results based on two alternative strategies as robustness checks: the first method uses labor contract cycles as an instrument for privatization, and the second compares transit costs between buses and subways (which, by their nature cannot be privatized) for cities that have both modes of transportation. Although these strategies are based on different identification assumptions, the empirical findings from these robustness checks are all consistent with those from the RD design.

[20] Between 1998 and 2014, the Amalgamated Transit Union, the largest transit worker's union in the nation, allocated an average of 93% of its political campaign contributions to Democratic candidates (see OpenSecrets.org Center for Responsive Politics).

[21] The Democrat winning margin ranges from -1 to 1 and it takes value -1 if none of the candidates are Democrats and 1 if all the candidates are Democrats.

transit agencies in cities where the Republican candidate won are significantly more likely to privatize a portion of their operations, even after controlling for several potentially confounding factors including market wage rates, unionization levels, and population density.

Our empirical model of the privatization decision is specified as follows:

$$d_{it} = D(S)_{it} + P(M_{it}; \kappa) + \mathbf{z_{it}}\delta + \nu_{it} \tag{2.10}$$

where $D(S_it$ is a function describing the mayoral political party after election year t for the city in which transit agency i operates. We allow the mayoral party effect to vary across states with strong relative to weak collective bargaining rights, such that $D(S)_{it} = \lambda_1 D_{it} + \lambda_2 D_{it} S$. The dummy variable D_{it} is equal to 1 if the winning mayoral candidate in election year t is a Democrat and zero otherwise. The dummy variable S is equal to 1 for states with legislation strongly in favor of collective bargaining rights. $P(M_{it}; \kappa)$ is a flexible polynomial function of the Democrat winning margin (M_{it}). $\mathbf{z_{it}}$ is a vector of UZA- and agency-level control variables including fixed effects. Eq. 2.10 provides the first-stage regression in the analysis of the cost Eq. 2.9. Our preferred specification for the control function P is a second order polynomial in M, however our results vary little after using cubic or quartic polynomials. For all specifications, we allow the polynomial coefficients to differ above and below the cutoff. Consequently, our main specification for the first-stage regression is as follows:

$$d_{it} = \lambda D(S)_{it} + \alpha_1 M_{it} + \alpha_2 M_{it}^2 + \alpha_3 D_{it} M_{it} + \alpha_4 D_{it} M_{it}^2 + \mathbf{z_{it}}\delta + \nu_{it} \tag{2.11}$$

We model the privatization decision under the assumption of uniform impacts of Democratic mayor across states ($D(S_{it} = \lambda_1 D_{it}$) as well as heterogeneous impacts by collective bargaining strength ($D(S)_{it} = \lambda_1 D_{it} + \lambda_2 D_{it} S$). We estimate the privatization treatment effectβ_1 in Eq. 2.9 via two stage least squares (Hahn et al. 2001; Lee and Lemieux 2010).[22]

Comparing outcomes across cities within a sufficiently narrow window around the winning margin cutoff permits one to draw causal inference regarding the effects of privatization on public transit operating costs. Our fuzzy RD (FRD) design provides a weighted average of the effects of privatization for transit agencies in cities with mayoral elections near the winning margin threshold, where weights reflect the ex ante likelihood that the agency's mayoral winning margin is close to the threshold.[23]

2.4 Data Sources

Our primary data source is the government-sponsored National Transit Database (NTD). Any transit agency that receives grants or financing from the Federal Transit Administration must report data to the NTD. Our analysis focuses on transit agencies providing fixed-route, public bus transit. We obtained information on annual VMT, passenger miles traveled (PMT), operating costs, fuel use, fleet characteristics, and

[22] Our fuzzy RD method has two IV's instead of the canonical example with one IV in Hahn et al. (2001). The second IV is an interaction term between the Democrat mayor dummy and the strong collective bargaining rights dummy. The interaction term allows the discontinuity in privatization to differ based on the strength of a state's collective bargaining rights.

[23] Our RD design reflects imperfect take-up with a continuous treatment variable. Lee and Lemieux (2010) refer to this type of RD design as a fuzzy RD with a "continuous endogenous regressor." Despite the semantic differentiation, the treatment effect in both cases is analogous to the treatment effect in an instrumental variables setting, and in both cases can be estimated using two-stage least squares. Thus, throughout this paper we refer to our RD design as a "fuzzy RD."

number of accidents and fatalities per year for 328 transit agencies from 1998 to 2011, for a total of 3,706 transit agency-year observations. Our sample of transit agencies includes all 50 of the largest public agencies that operate buses, as ranked by passenger trips as of 2011.[24] These data cover 236 distinct urbanized areas (UZA's) across the U.S.

Agencies are required to partition their reported operating cost and service information between services that were directly operated and those that were outsourced to a private contractor. We construct the privatization share by dividing the reported privatized VMT by total annual VMT. As of 2011, approximately 43% of transit agencies outsourced their operations to some degree between 1998 and 2011. Of those transit agencies, roughly half privatized 100% of their VMT for at least one year between 1998 and 2011. Fleet characteristics sourced from NTD include the size and average age of the bus fleet, the share of the fleet that is hybrid or operating on compressed natural gas (CNG), and the organizational type. We codify transit agencies as: an independent agency, a city agency, or "other" which includes agencies operated by the state DOT or planning agency subsidiaries.

NTD reports operating costs which includes labor costs (salaries, wages, benefits, and pensions), costs of fuel, materials and supplies, utility costs, taxes, and liability costs. They do not report these individual components separately. However, they do report annual fuel usage. Using fuel prices from the U.S. Energy Information Administration, we construct the annual fuel costs for each transit agency. Based

[24] 2013 Public Transportation Fact Book. American Public Transportation Association. 64th edition.

on reported operating costs and fuel costs estimates, we calculate the labor costs as total operating costs net of fuel costs. The labor cost is the major component of transit agency operating costs, averaging over 80% of total annual operating costs. Appendix 1 provides a discussion of the calculation of these fuel and labor cost measures.

The UZA-level data used in our empirical analysis consists of wages, unionization shares, average housing prices, a road congestion index, and population density. Because employee wages are not observed through the NTD data, we use earnings data from the Current Population Survey (CPS) to approximate wages paid to transit agency operations workers. Our measure of W is the average weekly wage earned by a full-time worker without any college education in an urbanized area. We use the urbanized-area average wage for two reasons.[25] First, we do not observe wages or benefits that transit agencies pay their employees. Second, the actual wages paid to transit employees are an endogenous choice made by the agency. The CPS wage data provide a measure of the competitive market wage for a low skill employee in an urbanized area and thus are unlikely to be correlated with transit agency-specific cost shocks. Labor market controls, including share of unionized workers, state collective bargaining rights and right-to-work status, account for the difference between the compensation of transit workers (reflected in the labor cost data) and the wage variable used in our regression (the average wage of all low-skill workers in a UZA). Unionization share is the ratio of union-represented full time-equivalent (FTE) workers to all FTE workers, also sourced from CPS. State-level data on collective

[25] We removed observations of earners below the 5th and above the 95th percentile.

bargaining rights as well as right-to-work status are obtained from Freeman and Valletta (1988).[26]

Population density is sourced from NTD, and covers years 2000 and 2010. Annual average housing price data are sourced from ACCRA Cost of Living Index. The road congestion index measures density of traffic across an urban area. The index is sourced from the Texas A&M Transportation Institute.

Mayoral election data prior to 2008 are sourced from Ferreira and Gyourko (2009). From the universe of Ferreira and Gyourko's original data, we use only elections after 1998 in urbanized areas that have public bus transit systems reporting to NTD. We extended election results for cities in Ferreira and Gyourko's dataset from 2008-2011 through contacting city clerks and performing online searches. These data include the names, political affiliation, number of votes received by the winner and runner-up, total election votes, date, and location for elections. Election data from both Ferreira and Gyourko (2009) and the authors' extension are limited to cities with populations over 25,000 with direct mayoral elections.[27]

[26]We use the 1996 values updated by Kim Rueben in our analysis. While the timing of the collective bargaining and right-to-work data does not overlap with our analysis, year-to-year variation of collective bargaining right strength varies minimally within state. See footnote 10 for further explanation on the construction of this variable.

[27]NTD data are aggregated to the UZA level while mayoral election data are at the city level. In instances where a UZA had multiple cities with distinct mayoral elections in a given year, we assigned the UZA a winning margin equal to the average winning margin across elections with the same winning political party. Cases where the winning political party differed across cities within the same UZA-year were excluded. We excluded elections where the winning mayor ran as an Independent. In instances where the candidate ran unopposed, the winning margin is 1 for Democrat winners (-1 for Republican winners). Our results are robust to excluding elections run unopposed.

The RD analysis uses only the dominant transit agencies in a UZA. We define saydominant transit agencies as the largest transit agency in a UZA, based on average annual VMT. Using only the dominant transit agencies effectively creates a one-to-one match within our sample between a city and a transit agency.[28] The majority of mayoral election cycles occur every three or four years, consequently our RD estimation is identified off of 548 unique elections. The final RD sample includes 1,444 observations covering 128 urbanized areas in 36 states from 1998 through 2011.

Table 2.3 characterizes the difference between transit agencies that are fully public and those that are not using sample means for UZA and transit agency characteristics. For both the RD sample and the full sample, labor costs and union participation rates are lower, on average, among agencies that engage in privatization. In general, the RD sample has similar mean characteristics as the full sample.

2.5 The RD Framework

Whether and to what extent privatization reduces the cost of public service provision are contentious issues. Using panel data on hundreds of transit agencies' annual

[28]We focus on the dominant transit agencies for the RD approach for the following reasons. First, the privatization decision of these agencies has more significant implications on the local budget than non-dominant agencies. Appendix Figure B.2 plots the privatization share in non-dominant agencies against the winning margin and there does not appear to be a discontinuity in the privatization share among close elections. Second, the non-dominant transit agencies tend to be those that serve suburban areas and their service area may overlap with multiple municipalities. Thus, the connection between a city mayor's political decisions and the city's public services is less clearly defined for such non-centralized transit services. For these two reasons, we focus on dominant transit agencies in our RD analysis.

operating costs, we seek to understand why agencies differ with respect to their cost of service provision and what roles privatization and public unions play in the cost of service provision. Since privatization is a choice, we must explicitly model the determinants of this endogenous variable. The following sections establish our RD as a valid experimental design that provides quasi-random assignment of privatization. We then show the relevance of our first stage relationship between mayor partisanship and privatization levels.

2.5.1 Testing the Validity of the RD Design

There are two important assumptions underlying the RD strategy and we discuss each of them in turn. The first assumption is that transit agencies in cities with mayoral elections in a narrow window around the zero winning margin threshold are similar on observable and unobservable dimensions.[29]

We provide evidence that treatment assignment is locally randomized by demonstrating that observable baseline covariates are balanced, and evolve smoothly over the winning margin threshold. To test covariate balance we conduct both graphical RD analysis and formal estimation. Figure 2.4 provides graphical evidence that selected covariates evolve smoothly through the winning margin cutoff. A comprehensive test of covariate balance is shown in Appendix Table B.1, where we compare sample means for each transit agency- and UZA-level baseline covariate across 30-

[29] See Lee and Lemieux (2010) for a comprehensive discussion of RD designs and their applications.

percentage-point bins. We test for covariate balance by regressing each variable in Appendix Table B.1 on the Democrat indicator variable, the Democrat winning margin, and their interactions.[30] The last column of Appendix Table B.1 shows the p-value associated with each discontinuity estimate (λ) from these regressions. We fail to reject continuity across the winning margin threshold for each covariate, except two (Share City agency; and the Share of CNG buses).

In cases where there are many covariates, it is possible that some discontinuities will be statistically significant by random chance (Lee and Lemieux 2010). To increase the power of the covariate balance analysis, we follow a "covariate index" procedure suggested by Card et al. (2015): we predict transit agency labor costs from a simple linear regression model relating the log of labor costs per VMT to each of the 11 predetermined covariates listed in Appendix Table B.1 as follows: $y_{it} = \mathbf{x}\kappa + \varepsilon_{it}$. The vector of coefficients, κ, is estimated from a random subsample of 150 observations. Next, we generate out-of-sample predictions of log(labor costs) per VMT: $\hat{y}_{it} = \mathbf{x}\kappa$. The resulting covariate index function can be interpreted as the best linear prediction of mean log labor costs given the vector of predetermined variables. In Appendix Figure B.3, we plot the mean values of these covariate-predicted labor costs around the Democrat winning margin threshold. Predicted labor costs evolve smoothly through the cutoff. Further, regression results show no evidence of a discontinuity in the predicted labor cost across the winning margin threshold.[31]

[30] Specifically, we run regressions of the form: $y_{it} = \sigma D(S)_{it} + \lambda_1 M_{it} + \lambda_2 M_{it}^2 + \lambda_3 D_{it} M_{it} + \lambda_4 D_{it} M_{it}^2 + \nu_{it}$ or ($|M| \leq 0.6$) where y is a covariate, M is the winning margin, and $D = 1$ if the winning mayor is a Democrat.

[31] We estimated the following regression: $I_{it} = \zeta D_{it} + \eta_1 M_{it} + \eta_2 M_{it}^2 + \eta_3 D_{it} M_{it} + \eta_4 D_{it} M_{it}^2 + \rho_{it}$ where I_{it} is the covariate index. The p-value associated with ζ, the discontinuity at the winning

The second assumption for a valid RD design is that transit agencies cannot precisely manipulate mayoral election voting outcomes, and subsequently select into the treatment. Voter turnout for our sample of election outcomes range from 3,000 to 257,000, in cities with populations ranging from 59,000 to 1.8 million. Given the large voter turnout, it is unlikely voters were able to precisely manipulate voting outcomes.[32] However, we provide descriptive evidence as well as a formal test that there is no discontinuity in the distribution of the running variable at the cutoff. Any evidence of a jump in the density of the running variable at the cutoff would suggest some degree of sorting into the treatment and would invalidate the quasi-experimental RD design.

In Figure 2.5, we plot the number of elections in each percentage point bin of the Democrat winning margin, as well as estimates of the Democrat winning margin density function. We estimate the density function following McCrary (2008).[33] There is no indication of discontinuity in the density of the running variable at the cutoff. A McCrary test confirms that there is no statistically significant jump in the running variable density function at zero.[34] This is consistent with absence of any

margin threshold is 0.750 without Table 2.3 controls, and 0.240 with controls.

[32]Near the winning margin threshold, voter turnout is comparable for UZA's on either side of the threshold. Average voter turnout is 48,500 for elections where the winner won by less than 20% of the vote share. Voter turnout is negatively correlated with winning margin. DiNardo and Lee (2004) observe this same trend in their analysis of union elections. The smallest voter turnout in their sample is 20, substantially smaller than the voter turnout observed in our mayoral elections.

[33]The McCrary test is based on an estimator for the discontinuity at the cutoff in the density function of the running variable. The test involves, first obtaining a finely-gridded histogram of the running variable; and second, smoothing the histogram using local linear regression on either side of the cutoff. The parameter of interest is the log difference in histogram height to the right and left of the threshold. See McCrary (2008).

[34]The point estimate of discontinuity in the density function is -0.40 with a standard error of 0.25, consequently we can reject the null hypothesis that the running variable density function is

voter manipulation that may undermine the RD design. Consequently, for transit agencies in UZA's with close margins of victory, the political affiliation of Republican or Democrat is assigned essentially at random.

2.5.2 Partisanship and Privatization

We present evidence on the effects of Democratic mayoral victories on a transit agency's level of privatized bus miles. Each dot in Figure 2.6 corresponds to the average transit agency privatization share following election t, given the margin of victory obtained by Democrats in election t. The solid and dashed lines in the figure represent the predicted values of privatization share from linear, quadratic, and cubic polynomial control functions without covariates. There are noticeable discontinuities at the cutoff, with conspicuous drops in privatization levels within each bin for positive Democrat winning margins near the 0% threshold. The partisan discontinuity is even more pronounced in Appendix Figure B.4, which "zooms" into a winning margin window around the threshold of -60% to 60%. Privatization levels, while significantly lower, are not zero for Democratic victories; however it is clear that there is a discrete increase in the probability of privatizing public bus transit operations in urbanized areas where a Republican won a mayoral election.[35]

The likelihood of privatization differs across states with either strong or weak

discontinuous at the threshold.

[35] Figures with varying bin sizes appear in the online appendix. The discontinuity in privatization share is robust under varying bin sizes.

collective bargaining rights. Panel B of Figure 2.6 shows that most of the jump generated in Panel A is driven by Republican victories in states with strong collective bargaining rights. Agencies in weak bargaining rights states show very little difference in privatization levels regardless of whether the Republican or Democratic mayoral candidate won an election (Panel C). Weak bargaining rights states may have less to gain from privatizing: these states tend to have lower wages, less union power, smaller city sizes, and lower urban density. In contrast, the potential cost savings is greater in strong bargaining rights states. Our analysis accounts for this heterogeneous treatment effect across strong versus weak bargaining rights states.

To improve estimation efficiency, we utilize the full sample of 1,444 transit agency-years in our main RD results. Approximately 10% of transit agencies in the RD sample fully engage in privatization (i.e., their privatization share is one), which is why some observations in Figure 2.6 appear as extreme outliers. While use of the full sample size improves estimation precision, our estimates are generally robust to restricting observations to elections with margins of victory closer to the winning threshold.

Regression estimates of the discontinuities in privatization share are shown in Table 2.4. Each cell in the first row reports an estimate of the effect of electing a Democratic mayor on privatization based on various specifications of Eq. 2.11. All specifications include quadratic control functions and their interactions with the mayor party dummy. Standard errors are clustered at the UZA level. A Democratic mayoral victory reduces the likelihood of privatizing bus miles. In column (3), we allow the discontinuity estimate to differ for states with strong vs. weak collective bargaining rights laws by including an interaction of the Democrat mayor indicator with strong bargaining rights indicator ("Democrat x Strong Barg"). This variable is equal to one for a transit agency operating in a strong bargaining rights state in a year where a Democrat mayor was elected, zero otherwise. The negative coefficient on "Democrat x Strong Barg" implies that Democrat mayors in cities with strong bargaining rights are even less likely to engage in privatization than their mayoral counterparts in areas with weak bargaining rights. The estimates in column (3) are jointly significant with 95% confidence level. Column (4) uses local polynomial regression estimation with a robust confidence interval developed by Calonico et al. (2014) and **?**, henceforth "CCT".[36] We report local polynomial regression estimates under a triangular kernel density and bandwidths following the optimal bandwidth selection procedure of CCT (2014).

The results reported in Table 2.4 suggest that a Democratic mayor reduces privatization levels by between 6 and 10 percentage points. Results are robust to the specification of the control function or the inclusion of covariates. All estimates are significant at the 10% level.

[36]The performance of available local polynomial confidence intervals are sensitive to the chosen bandwidth, and can produce biased intervals with low empirical coverage in finite samples. The confidence interval proposed by CCT corrects for this bias by rescaling the conventional bias-corrected t-statistic with a standard error formula that accounts for the additional variation introduced by the estimate bias. The resulting confidence interval allows for mean squared optimal bandwidth selectors, and improved coverage rates. See Calonico et al. (2014) for a detailed discussion on the construction of their bias-corrected confidence intervals.

2.5.3 Placebo Tests

We carry out a series of placebo tests to verify that the observed jump in privatization levels is, in fact, driven by the Democrat winning-margin threshold. In Figure 2.7, we plot the discontinuity estimates from a series of regressions, where we fix the winning margin cutoff to an alternative threshold. For expository purposes, we show estimation results over the range -0.4 to 0.4. Each dot corresponds to a discontinuity estimate in the privatization share around the "false" Democrat winning margin thresholds. We present results from local polynomial regressions including a quadratic polynomial in the winning margin. Estimates were noisier for thresholds below zero due to smaller sample sizes. However, only at the true cutoff of zero is the discontinuity estimate statistically significant.

Given the significant discontinuity in privatization levels across the winning margin threshold, and given that other covariates are distributed smoothly across the threshold, one can interpret any discontinuity in the conditional distribution of unit labor costs as a causal effect of privatization.

2.6 Empirical Results of the Privatization Effect

In this section, we first report our regression discontinuity estimates of the cost savings associated with privatization, and provide evidence that our main results are robust to various specifications of the RD design. We further demonstrate that privatization is the main mechanism driving cost savings for transit agencies. In Section 2.6.3, we test whether privatization affects public transit quality and reliability as a consequence of reducing costs. Our final set of empirical results provides evidence that privatizing transit agencies are able to reduce costs by limiting instances of "featherbedding."

2.6.1 Main Results: The Effect of Privatization on Labor Cost

We investigate whether privatization reduces transit agency labor costs. Because pictures are illuminating in a regression discontinuity context, we first provide graphical evidence of the relationship between labor costs and mayoral election outcomes. Each dot in Figure 2.8 corresponds to the average unit labor cost that follows a mayoral election t, given the margin of victory obtained by Democrats in election t among transit agencies in strong bargaining rights states. The lines in Figures 2.8 represent the predicted values from a linear, quadratic and cubic polynomial fit of the winning margin, without covariates. Republican victories over Democrat mayors have the greatest impact on privatization levels in strong, as opposed to weak bargaining rights states. Consequently, the discontinuity in labor costs near the winning margin threshold is driven by agencies in strong bargaining rights states (Panel B).[37] In such

[37] Appendix Figure B.5 shows residuals plots from a regression of log unit labor costs on all transit agency and UZA controls – excluding the privatization share –, as well as state and year fixed effect along the Democrat winning margin. In strong bargaining rights states, costs are lower after closely won Republican victories. Graphical evidence of the persistent discontinuity in labor costs, as well

states, costs are lower after closely won Republican victories.

Table 2.5 reports estimates of the privatization treatment effect, β_1 in Eq. 2.9, under three types of estimators: OLS, fuzzy RD, and the local polynomial estimator. We show results with and without UZA and agency control variables. All regressions include year and state dummies, so that the relationships between privatization levels and labor costs are identified using variation across agencies within the same state. Standard errors are clustered at the UZA level. OLS estimates in columns (1) and (2) suggest that a 1% increase in a transit agency's privatization share reduces their per-mile labor costs by between 0.16% and 0.3%. For a transit agency with no privatization, the labor cost-savings from fully privatizing transit operations would be between 15.0% and 30.9%.[38] However, the OLS estimates are likely to be biased by unobserved factors that are correlated with a transit agency's level of privatization that also affect labor costs. Columns (3) – (4) present analogous results under an FRD approach in which we instrument the privatization share with an indicator equal to one if a Democrat won a mayoral election. The FRD estimates in columns (3) and (4) are larger in magnitude relative to the naive specifications in columns (1) and (2). Thus, omitted variables produce a downward bias on the privatization effect. This is consistent with our prior finding that large agencies in cities with strong unions have more to gain from privatization but are less likely to privatize. The FRD estimates suggest a 1% increase in privatization reduces unit labor costs

as residual labor costs, under varying bin sizes appear in the online appendix.

[38]Because Eq. 2.9 takes the semi-logarithmic functional form and the privatization rate ranges from 0 to 1, a consistent estimator for the percentage impact from full privatization (e.g., privatization rate going from 0 to 1) on the labor cost is $100 * \left[\exp(\hat{\gamma} - \frac{var(\hat{\gamma})}{2}) - 1\right]$

between 0.4% and 1.9% for those transit agencies near the 50-50 margin of victory.

Columns (5) and (6) employ the flexible version of equation (9) where the Democrat mayor effect is allowed to differ by states with strong collective bargaining rights. These regressions show that outsourcing public transit to private entities by 1% will decrease labor costs per mile by between 1.0% and 2.1%. The full effect of privatization is a reduction ranging from 46.4% to 68.1% in unit labor cost after controlling for observables. Inclusion of UZA and transit agency control variables reduces the magnitude of the FRD estimates substantially. Finally, the local polynomial estimate in column (7) is of similar magnitude to the FRD estimates.[39,40]

The magnitude of cost savings from privatizing suggested by the FRD estimates is considerably larger than the OLS estimates. This difference in effect size indicates that treatment is endogenous, and the highest-cost agencies are not engaging in privatization. The difference in estimated effect sizes could also reflect the fact that FRD estimates are local; the FRD design predicts the average treatment effect for agencies with close mayoral elections, where the Republican candidate won by a small margin. These estimates may not generalize to a broader sample of transit agencies. Further, because the regression discontinuity analysis sample consists only of dom-

[39] The local polynomial estimate is consistent for various bandwidth sizes, as shown in Appendix Figure B.6.

[40] To investigate the potential bias into FRD results due to the relatively weak first stage, we employ a limited information maximum likelihood (LIML) estimator, which has the same asymptotic distribution as the FRD (i.e., 2SLS) estimator. The LIML estimator is less precise than 2SLS, but provides finite-sample bias reduction (Angrist and Pischke 2008; Greene 2003). Results of LIML estimation are shown side-by-side with that of the FRD estimation in Appendix Table 2.3. The effect of privatization on labor costs are qualitatively and quantitatively similar between LIML and FRD estimators, further the standard errors from LIML estimation are similar in magnitude to that of the FRD estimator, suggesting weak correlation is not substantially biasing our results.

inant transit agencies, the predicted cost savings are likely to be an upper bound estimate for non-dominant transit agency. We investigate the relationships between transit agency size and labor costs in Section 2.6.4 and find dominant transit agencies have approximately 20% higher labor costs per mile relative to non-dominant transit agencies, all else equal.

2.6.2 Specification and Robustness Tests

Contemporaneous Treatment Effects

An assumption underlying our RD strategy is that mayor partisanship affects transit operating costs only through its effect on the propensity to privatize. That is, transit agencies operating under Republican mayors experience lower operating costs relative to their Democrat-electing counterparts because Republican mayors are more likely to outsource public transit to private entities. We recognize that our approach would not recover the causal impact of privatization on costs if Republican mayors engaged in significantly different cost-cutting measures than Democrat mayors that are either unobserved, or unaccounted for in our specifications. Our tests for covariate balance (Figure 2.4, Appendix Table B.1, Appendix Figure B.3) established that pre-existing differences in city-wide characteristics do not explain our results. However, there remains a possibility that contemporaneous treatments, such as elimination of unprofitable routes, are driving our results.[41]

[41]In Appendix Figure B.7, we also show that incumbency or power transition effects within our mayoral election sample are not spuriously driving our regression discontinuity results. We

We test whether election of Republican mayors is systematically correlated with transit agencies eliminating unprofitable bus routes in an effort to reduce costs. Appendix Table B.2 shows results of regressing mayoral political affiliation in election year t on the level of transit service following the election. Among closely won elections (those with a Democrat winning margin between -60% and 60%), we find no significant difference whether we measure transit services in terms of annual VMT, annual PMT, or occupancy rate. Regressions that use the full support of the Democratic winning margin show that cities with Republican mayors have higher levels of VMT and PMT, on average, compared to Democrat mayors. This is the opposite result we would expect if Republican cities reduced costs by cutting transit services. Rather, the positive relationship between Republican mayors and transit service is the general equilibrium result we would expect after the consumer price of transit falls, and demand for transit increases. Additionally, it is possible that privatized transit agencies are able to offer service to remote, low-density areas that were too costly under public transit operations. Later in Section 2.6.4, we explore the mechanisms for cost savings under privatization.[42]

test that the closely won elections in our sample capture variation in political affiliation itself (ie Democrat or Republican party), as opposed to characteristics of mayoral power transitions that are correlated with political affiliation. For example, it is possible that political administration changes by newly elected mayors (e.g., non-incumbent mayors) induce privatization of bus transit, and newly elected mayors are more often Republican than Democrat. Our empirical results show that mayoral power transition characteristics evolve smoothly through the winning margin cutoff. Closely won Democrat victories are no more likely than Republican victories to have a non-incumbent winner, a new political party, or a mayor serving a longer tenure.

[42]It is unlikely that scale economies harnessed by Republican-won cities explain cost savings from privatization because our cost outcome variable of interest excludes fixed costs.

Alternative Identification Strategies

To further test the robustness of results to the identifying assumption, we estimate the effect of privatization based on two alternative strategies. The first strategy uses labor contract cycles as an instrument for privatization. The second strategy compares the operating costs per mile between buses and subways - which, by their nature cannot be privatized - for cities that have both modes of transportation. Although these strategies are based on different identification assumptions, the empirical findings from these robustness checks are all consistent with those from the RD estimates. We present the implementation and results of these alternative identification strategies in greater detail in Appendix B.3.

Varying the Margin of Victory Window

As a final robustness check, we verify that our chosen window around the winning margin threshold does not affect the main results. We re-estimated the regression specifications with full UZA and agency controls in Table 2.5 under several different ranges of the winning margin. Appendix Table B.3 shows the OLS, FRD, and local polynomial results under windows ranging from 10% to 70% around the threshold. These results are similar in sign and magnitude to those discussed in Section 2.6.1: the OLS estimates suggest a cost savings between 20% and 36% while FRD and local polynomial estimates imply cost savings between 70% and over 90%. The implied cost savings are larger for smaller windows; however, these estimates are less precisely estimated due to the small sample sizes. Because the estimates remain reasonably stable across different windows of the winning margin, we conclude that our main findings are robust to the choice of the winning margin range.

2.6.3 Effects of Privatization on Service Quality and Ridership

Critics of privatization often point to worsening service quality as a drawback of outsourcing. Hart et al. (1997) theorize that in-house service provision in the context of prisons is optimal if cost reductions from privatizing have substantial impacts on non-contractible quality. However, if quality reductions from cost reductions can be controlled through contracts or competition, then privatization is optimal. Therefore, the ability to measure and monitor service quality is an important factor that determines whether a public service can or should be privatized.

We estimate the effects of privatization on several indicators of service quality, including annual VMT, PMT, ridership, and the number of non-fatal and fatal incidents per year reported by the transit agencies. Non-fatal incidents include collisions, vehicles leaving the roadway, fires, electric shocks, or security incidents. Fatal incidents include fatal collisions with pedestrians, cyclists, suicides, or operator deaths. The incidents data were available from NTD from 2002 through 2011. Annual count of incidents provides a proxy for bus operator negligence. Outcomes of annual VMT, PMT, and rider occupancy measure whether transit agencies that outsource are more likely to reduce transit services or if unobservable factors cause rider demand to decline when operations are managed by a private entity. Evidence of reductions in service levels or demand would confound the cost-saving effects identified in our main analysis. Specifically, privatizing might reduce costs not through efficient management and labor utilization, but through reducing the services rendered to the public. Any evidence of reductions in rider occupancy would more importantly suggest that privatizing has welfare-diminishing effects by pushing people away from public transit.

We first provide graphical evidence that service quality does not differ across mayoral political affiliation. Appendix Figure B.8 plots the residuals generated from regressing each of the service variables VMT, PMT, rider occupancy, and number of incidents on transit agency observables. We would expect to see a discontinuity at the winning margin if political affiliation induced significant differences in service level or quality through privatization. The absence of a discontinuity suggests that reductions in overall bus service, or reductions in service quality cannot explain the labor cost savings observed in our main results.

Our regression analyses relating service quality attributes to agency privatization levels support the graphical results that privatization does not reduce costs through literal or figurative shortcuts to transit service. Appendix Tables B.3 and B.4 show estimation results following Eq. 2.9 where incidents, VMT, PMT, and rider occupancy are the outcome variables. Transit agencies with higher levels of privatization report a lower number of incidents per year under negative binomial, FRD, and local polynomial estimation, although the effect is imprecisely estimated (Appendix Table B.3).[43] Results of Appendix Table B.4 show that privatizing bus transit is positively correlated with VMT under both OLS and FRD estimation (columns 1-2), contrary to what we would expect if cost savings resulted from reducing transit services. Taking all regression results together as a whole, however, suggests that there is no significant relationship between privatization share and transit service offerings. Columns (4) and (5) show that privatization does not significantly affect rider occupancy.

These analyses suggest that private firms provide public bus transit at a lower cost than public operators, and these cost savings do not come at the expense of reduced miles of service or reduced quality.[44]

2.6.4 The Cost-Saving Channels from Privatization

Non-privatized, unionized transit agencies often face a higher labor cost due to restrictions on use of part-time employees, and through obligations to larger pensions and fringe benefits (Giuliano and Lave 1989; Black 1991; DiSalvo 2010; Winston 2010). We test whether publicly operated transit services exhibit under utilization of employees, or "featherbedding" of employee headcount. Due to the limitations of NTD data, we observe employee headcount only for the publicly operated tran-

[43]Because the incidents data (annual number of fatal and non-fatal incidents) are count data with several zeros, we use a negative binomial estimator, in addition to RD methods.

[44]In the online appendix, we show graphically that there is little evidence of discontinuities in number of incidents, VMT, PMT, or rider occupancy at the democrat winning margin cutoff.

sit services, thus we are not able to compare employee headcount for public and private operations directly. However, we are able to test the effect of state-level union strength on full-time employee headcount, holding all other factors constant, for publicly operated transit services. Table 2.6 shows results of regressing the log of the number of full time employees per publicly operated VMT on measures of state-level bargaining rights and a right-to-work state indicator. Stronger state bargaining rights are correlated with higher full-time employee headcount per VMT. Conversely, transit agencies operating in right-to-work states have a lower full-time employee headcount per VMT. The negative correlation between right-to-work laws and employee headcount per VMT persists even for transit agencies operating in states with strong collective bargaining rights.

In a similar spirit as Thomas (1998), we examine transit agencies in counties that share a state border, where right-to-work laws differ on either side of the state border. For instance, Kansas – a right-to-work state, and Missouri – a non-right-to-work state - share a border. We compare the transit labor costs between transit agencies in two counties that share the Kansas-Missouri border. These counties are arguably comparable on unobservables due to their close geographic proximity. After controlling for city and agency-specific observables like average wage, union strength, density, and cost of living, any variation in labor costs between these two agencies is due to their differing right-to-work status. The NTD data consists of 17 such border-county pairs comprising 21 transit agencies. Table 2.7 shows the results of fixed effect regressions. Each specification includes the 17 border-pair fixed effects. The sample size is small, however the right-to-work estimate is precisely estimated, and is robust to the inclusion of various fixed effects and UZA and agency controls. The evidence in Tables 2.6 and 2.7 suggest that right-to-work states have lower labor costs, all else equal, and leaner use of staff is at least partially responsible for these lower labor costs.

As a final piece of evidence demonstrating the role that unions play in lobbying large liberal city transit agencies, in Table 2.8 we report regressions in which the dependent variable is the log of operating cost per mile and the explanatory variables are dominant agency indicators and their interactions with right-to-work status and union strength. These regressions include UZA-by-year fixed effects, which serve to control for all local trends that can impact operating costs. The remaining variation is cross-transit agency variation within a UZA. Since major UZAs such as Chicago have multiple transit agencies, we compare cost per mile for dominant transit agencies and suburban agencies. In addition, we stratify these estimates by the state's right-to-work status. Column (2) shows that in pro-union, non-right-to-work states, operating costs are 36% higher in the dominant agency relative to the suburban agencies. However, in right-to-work states such as Texas, operating costs are only 3% higher in the dominant agency relative to the suburban agencies. If the efficiency wage theory explained why public transit employees receive higher compensation relative to private-sector employees, we would not expect to see an operating cost differential across transit agencies within the same city for the same unit of service. Rather, our preferred explanation is that unionized labor in liberal areas focus their negotiation efforts on the dominant transit agencies, generating an operating cost premium for these larger agencies.

An important dimension of the privatizing process is that the employees operating and maintaining buses often include the same people whether the bus is operated by the private or public entity. The majority of these employees are unionized, even if they are on payroll of the private entity. Thus, the impact of unions on transit operating costs is not absent when operations are outsourced to private firms. Rather, private firms are better equipped to negotiate labor contracts with union representatives.

Several private operators are large multinational firms that can draw upon best practices in contract negotiations and employee management. More importantly, private firms serve to create a disconnect between unionized transit employees and the transit authority responsible for budgeting and planning. Because of their affiliations with a political organization, public unions have the power to increase and sustain levels of public employment as well as impact local finance policy to increase taxes and salaries (DiSalvo 2010). Severing the tie between the policy-making institution and the employees is one of the most important mechanisms by which private firms reduce costs of public services.[45]

[45]DiSalvo (2015) discusses important advantages that public sector unions exercise over private sector unions at the negotiations table including: protection from business cycles due to the ease with which governments can borrow money, influence on employers through political contributions, and absence of direct mechanisms to control management because voters are generally uninformed of local government spending.

2.7 Welfare Analysis of Privatization

We estimate the welfare loss from overpaying for public bus transit by comparing the welfare under the current level of privatization as well as under full privatization with some simplifying assumption. The welfare loss stems from the high cost of service rendered by public sector provision of transit services relative to private sector provision identified in Section 2.6. In this section, we demonstrate how the cost premium from public sector provision results in insufficient transit service levels in equilibrium. While we found no significant effect of privatization on service levels in Section 2.6.3, our analysis describes the partial equilibrium. The long run, general equilibrium result of lower operating costs is an increase in demand and supply of transit services.

We first estimate the demand curve for bus services in terms of the number of passenger trips as a function of trip price. Total passenger trips on fixed-route bus transit as of 2011 were approximately 5.2 billion.[46] We obtain the bus ridership elasticity with respect to fare of -0.4 from prior studies on transportation demand.[47] Using the estimates for total operating costs per mile from the RD design described in Section 2.6.1, we calculate the average predicted bus fare for a transit agency as of 2011 under current levels of privatization as $\alpha_1 = \$2.98$. If all transit agencies privatized their operations completely, the average predicted bus fare as of 2011

[46] *APTA 2013 Public Transportation Factbook*, Table 5.

[47] Gagnepain and Ivaldi (2002) estimate an elasticity of bus ridership with respect to fare of approximately -0.441. Oum et al. (1992) estimate elasticities between -0.3 and -0.7. TCRP (2004) estimates an elasticity of -0.40.

would be $p_1 = \$1.87$. Using the current bus fare α_1, counterfactual bus fare p_1, the elasticity of bus ridership, and the aggregate number of passenger trips, we calculate the aggregate demand curve for bus passenger trips. Further details on these calculations are presented in Appendix B.2.

Figure 2.9 shows the predicted change in consumer welfare if all bus services were privatized and operated by private companies. The shaded area is the change in consumer surplus associated with privatization. Part of the change in consumer surplus is the result of a wealth transfer from consumers to public employees of transit agencies. Specifically, at least part of the differential between α_1 and p_1 is a wealth transfer from consumers to the employees of public transit in the form of higher wages. Area D in Figure 2.9 is pure deadweight loss associated with this transfer. While there are benefits to society from transit agencies employing individuals who might otherwise receive lower pay, the transfer is not costless. Since this transfer makes the public provision of transportation more expensive, demand for passenger trips is pushed to lower levels than it would be under full privatization, indicated by the lower dashed line. Under fully privatized operations, the long run, general equilibrium result of lower operating costs would be an increase in aggregate ridership would from 5.2 billion to approximately 6.2 billion passenger trips.

The national change in consumer surplus under the average impact of privatization on operating costs amounts to approximately \$6.3 billion. Pure deadweight loss is approximately \$524.3 million. This amount neither serves wealth redistributive purposes, nor is it compensated to transit riders for the inability to use transit at an efficient cost. That is, U.S. taxpayers lost \$524.3 million in aggregate from over-paying for bus transit in 2011. Losses of similar magnitude exist for subsequent years.

The social welfare gains from privatizing public transit are complicated by potential welfare losses to transit employees. Under privatized management, bus operators may suffer from a more efficient use of inputs, either through lower pay or longer work hours. However, the increase in ridership under lower-cost, privatized operations also serves to increase demand for bus transit labor. In 2011, public bus transit employed over 130,000 individuals as bus operators. These drivers served 5.2 billion passenger trips. Holding the ratio of drivers to trips constant, an increase in ridership demand to 6.2 billion trips would translate into over 26,000 additional bus operator jobs alone (not including new maintenance and managerial jobs that would also be necessary). To the extent that these workers value their jobs more than their second best alternative, the \$524.3 million deadweight loss in 2011 is a lower bound because it does not account for the welfare gain from employment and job creation that would result from increased demand for bus transit.

Table 2.9 shows the estimated deadweight loss in 2011 for three major transit agencies: Boston's MBTA, Chicago's CTA, and San Antonio's VIA. Boston's deadweight loss from inefficiently high transit costs amounts to approximately \$14.8 million, while Chicago's CTA has estimated deadweight losses of approximately \$26.3 million. In this same year, Boston and Chicago received \$846 million and \$705 million, respectively, in public funding for operations. The deadweight loss from inefficiently high operating costs accounted for 1.7 and 3.7%, respectively, of local, state, and federal funding sources for these two agencies in 2011.

Both Chicago and Boston operate in states with collective bargaining rights, whereas collective bargaining is outlawed in Texas. Notice that deadweight losses for San Antonio's VIA are estimated to be substantially lower at approximately $4.2 million. San Antonio's fleet is half the size of Boston, however their estimated deadweight loss is less than one third that of Boston.

2.8 Conclusion

Our analysis of public sector efficiency in the provision of bus services suggests three important findings. First, the private sector provides a unit of bus transit service at a substantially lower cost than the public sector. Second, the cost premiums are higher for larger transit agencies, when state-level collective bargaining rights are stronger, and when local political leaders are Democrats. Third, privately operated firms are able to curb labor costs per mile through efficient allocation of labor and through limiting featherbedding. Strong union bargaining power in non-right-to-work states appears to increase the number of full time employees on payroll, holding the service area constant. These effects are most pronounced for a city's largest transit agency, which suggests public sector unions shift their bargaining pressure toward the dominant transit agencies such as Chicago's CTA, New York's MTA, or Boston's MBTA. The operating cost per bus mile in Chicago's major transit agency in 2012 was $13.18, or 175% that of the operating cost per bus mile in Houston's major transit agency. While individual transit agencies absorb these costs, the workers who receive these higher wages benefit. During a time of great concern about income inequality, public employment offers stable middle-class jobs for less educated and disadvantaged workers (Alesina et al. 2000, Boustan and Margo 2009). Future research could examine the implications of public sector employment and pay for overall metro area income inequality trends.

Using RD estimation to identify the effect of privatization on total operating costs, we find that the difference in operating costs per mile between mean privatization levels as of 2011 and the counterfactual scenario when all bus miles are privatized is a decrease from $6.63 to $4.16 per mile. Holding total service mileage (2.3 billion VMT in 2011) constant, this per-unit cost savings translates into an aggregate cost savings of approximately $5.7 billion, or 30% of aggregate operating expenses.

The documented potential for cost savings from outsourcing bus operations is likely to be an underestimate of total government cost savings from improving the efficiency of public sector service provision. Transit unions are crucial political engines in major cities (Koch 2011). Thus, concessions from public employers to the transit unions are very likely to spread to other unions within a city and positively impact their negotiation power. Known as the "me-too" bargaining strategy, the practice of negotiating salaries and benefits based on that of other public employees has been attributed to state and local government budget deficits, for example in California, New York, and Vancouver, BC (Miller 2010). Future research should study the timing of public sector union contracts and the sequential game played by a big city mayor with the public sector unions. Republican mayors may enjoy cost savings across the board after creating an early reputation for being tough negotiators.

Figure 2.1: Operating Costs Per Vehicle Miles Traveled (VMT) in States with Strong vs. Weak Bargaining Rights

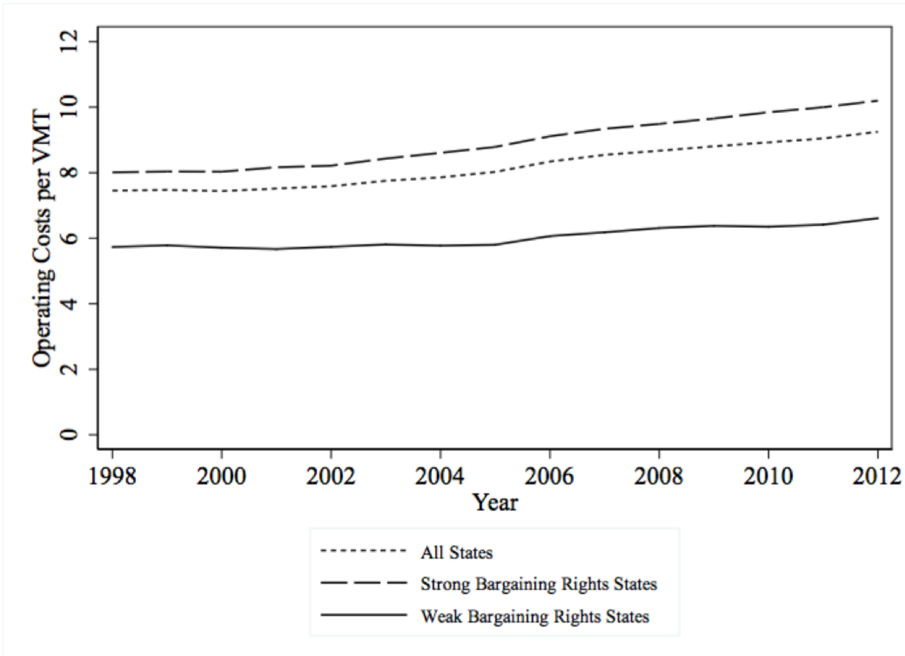

Note: Operating costs are deflated by CPI. Operating cost transit agency averages are weighted by transit agency annual VMT. States with strong bargaining rights laws are those where legislative mandate either implicitly or explicitly dictates public employers and union employees must come to an agreement on contract negotiations. States with weak bargaining rights laws either prohibit collective bargaining all together, or do not mandate that the public employer bargain with unions.

Figure 2.2: Distribution of Labor Costs per Vehicle Miles Traveled (VMT) in States with Strong vs. Weak Bargaining Rights

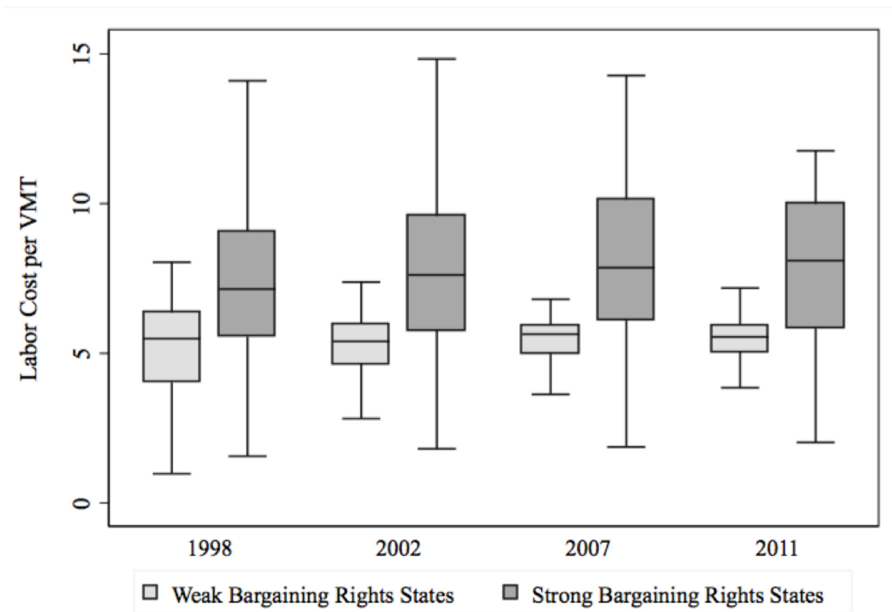

Note: Cost values reported in terms of 2011 dollars. Values are weighted by transit agency annual VMT. Figure does not show outside values.

Figure 2.3: Share of Privately Supplied Vehicle Miles (Privatization Share) in States with Strong vs. Weak Bargaining Rights

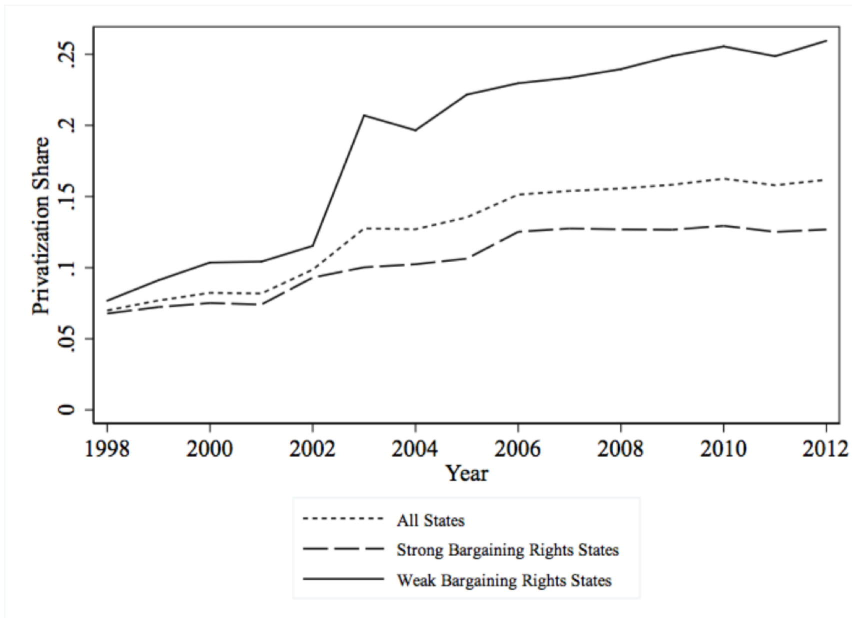

Note: This figure shows the privatization share among states with strong and weak bargaining rights laws. States with strong bargaining rights laws are those where legislative mandate either implicitly or explicitly dictates public employers and union employees must come to an agreement on contract negotiations. States with weak bargaining rights laws either prohibit collective bargaining all together, or do not mandate that the public employer bargain with unions. Designation of “strong” vs “weak” bargaining rights states discussed in more detail in footnote 10. Privatization shares are calculated as total annual VMT operated by private contractors divided by the total annual VMT. Privatization shares averages for strong and weak bargaining rights states are weighted by total annual VMT. The jump from 2002 to 2003 in weak bargaining-state privatization share is caused by three transit agencies that increased their privatization share by over 30% between 2002 and 2003. These include agencies in Dallas, TX (DART), College station, TX (Brazos Transit District), and Phoenix, AZ (Valley Metro).

Figure 2.4: Selected baseline UZA characteristics at the Democrat winning margin threshold. Sample includes 1444 agency-year pairs. Bin size=0.04.

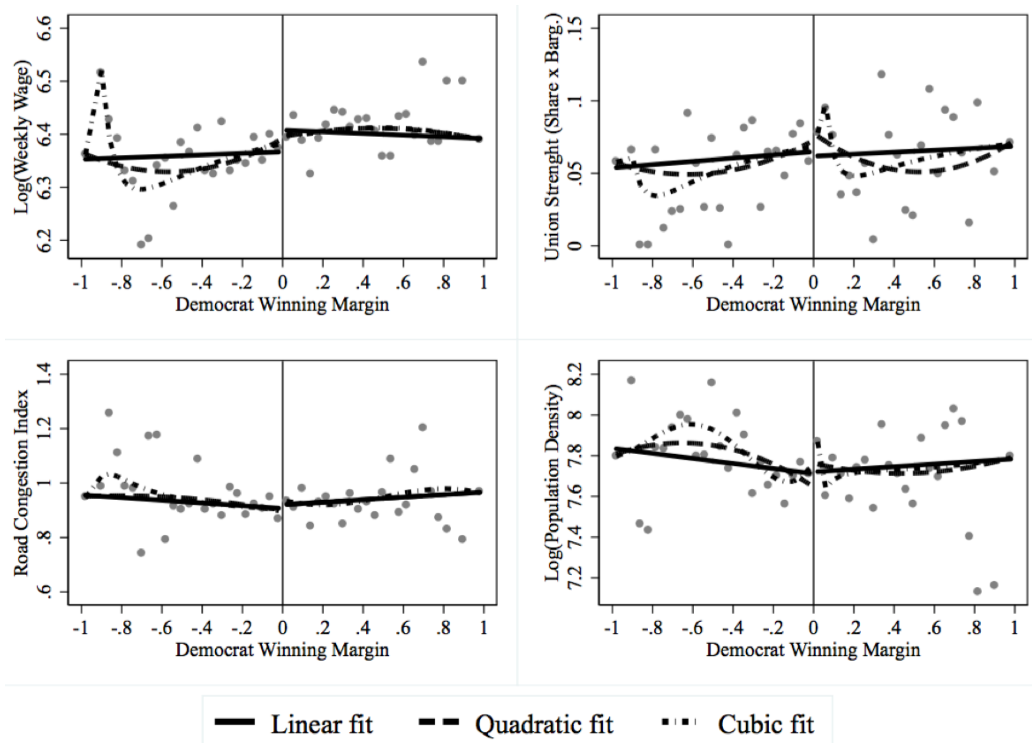

Note: Each dot corresponds to the average noted characteristic of a transit agency's UZA following mayoral election t, given the margin of victory obtained by Democrats in election t. Solid and dashed lines each represent the predicted values of the noted characteristic from linear, quadratic, and cubic polynomial control functions of the winning margin.

Figure 2.5: The Distribution of the Democrat Winning Margin Running Variable.

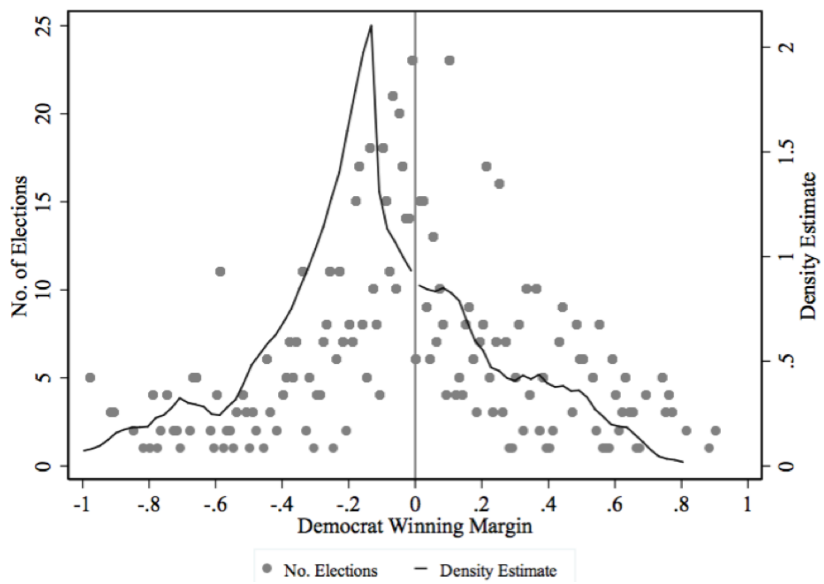

Note: Figure plots the raw number of elections (as dots) in 1-percentage point bins of the Democrat winning margin. The curve was estimated using a smoothed local linear density estimator following McCrary (2008) where the midpoint height of each bin of size 0.02 are the regressors, and the normalized counts of the number of observations falling into each bin are treated as the outcome variable. Figure excludes elections with Democrat winning margins of -1 (Republican won all votes) or 1 (Democrat won all votes).

Figure 2.6: Privatization Share and Winning Margin. Bin size=0.04.

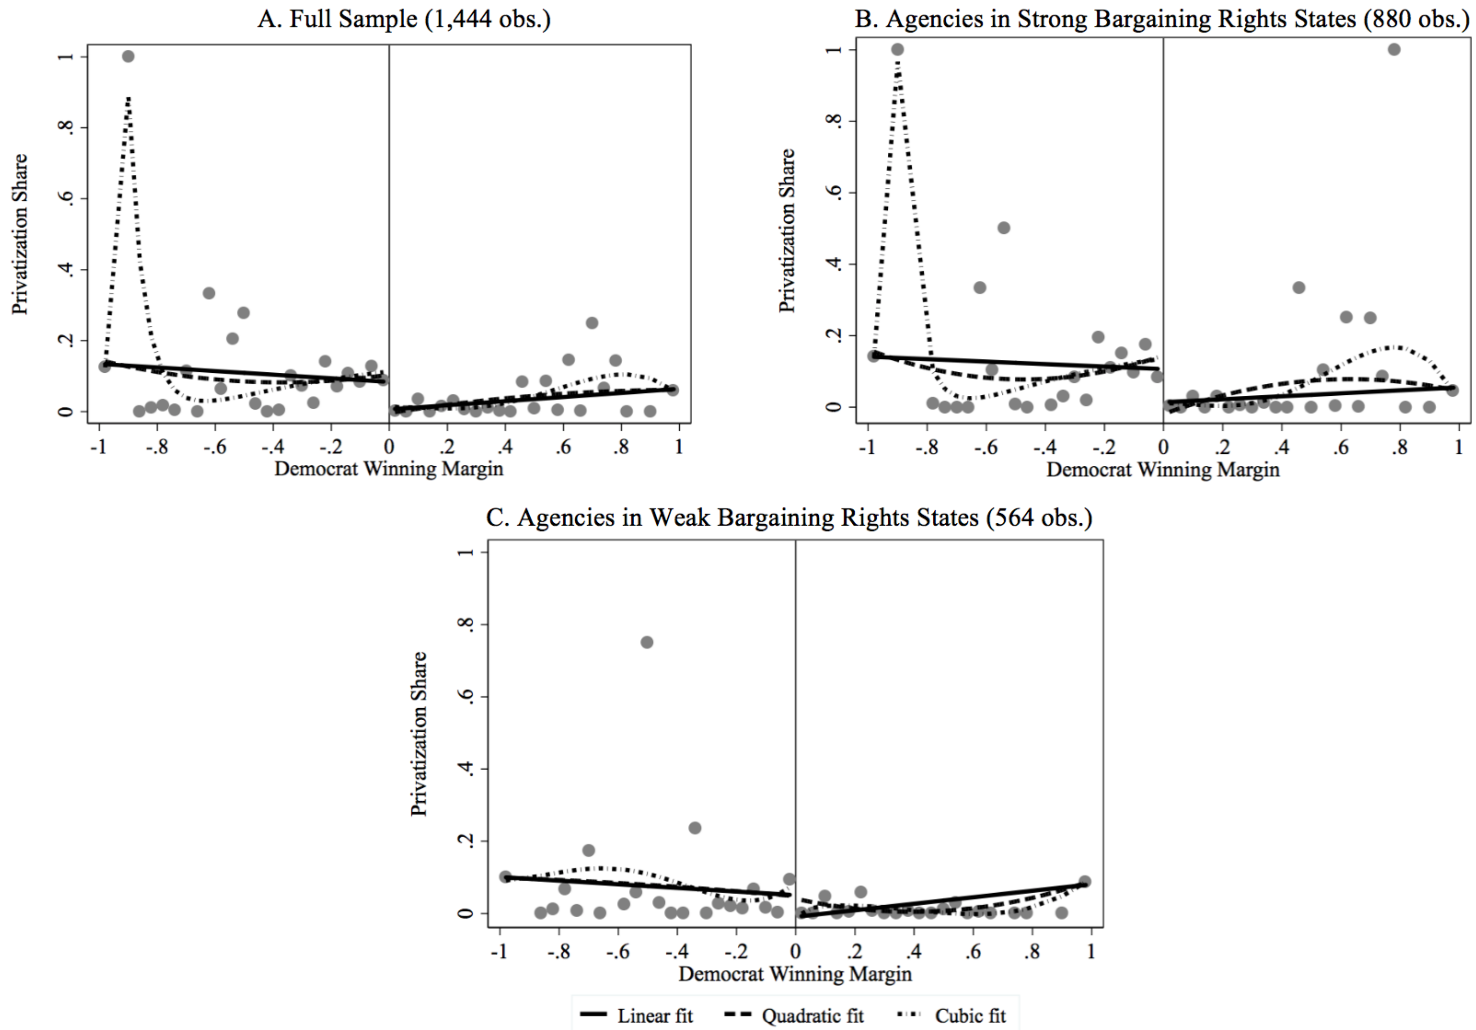

Note: Each dot corresponds to the average transit agency privatization share following mayoral election t, given the margin of victory obtained by Democrats in election t. Solid and dashed lines each represent the predicted values of privatization share from linear, quadratic, and cubic polynomial control functions of the winning margin

Figure 2.7: Placebo Estimates of Discontinuity in Bus Transit Privatization Share

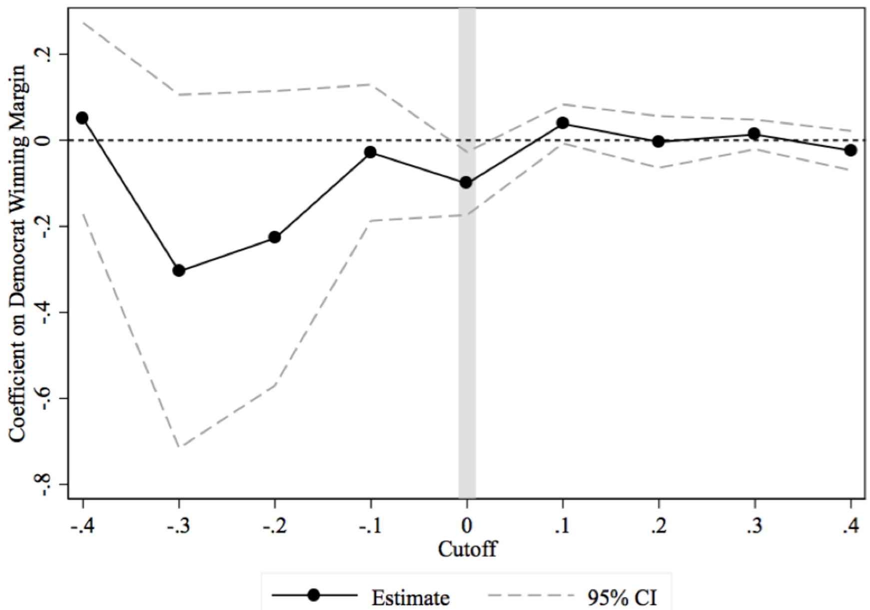

Note: Estimates based on local polynomial regressions of the form $d_{it} = P(M_{it}; \kappa) + \varepsilon_{it}$ with robust confidence intervals developed by Calonico et al. (2014). Polynomial regressions are of order two ($\kappa = 2$) in winning margin (M_{it}). Kernel type is triangular. Bandwidth, 0.294, is chosen using CCT mean-squared error optimal bandwidth selector.

Figure 2.8: Ln(Labor Costs per VMT) at the Democrat winning margin threshold. Sample includes 880 agency-year pairs. Bin size=0.04.

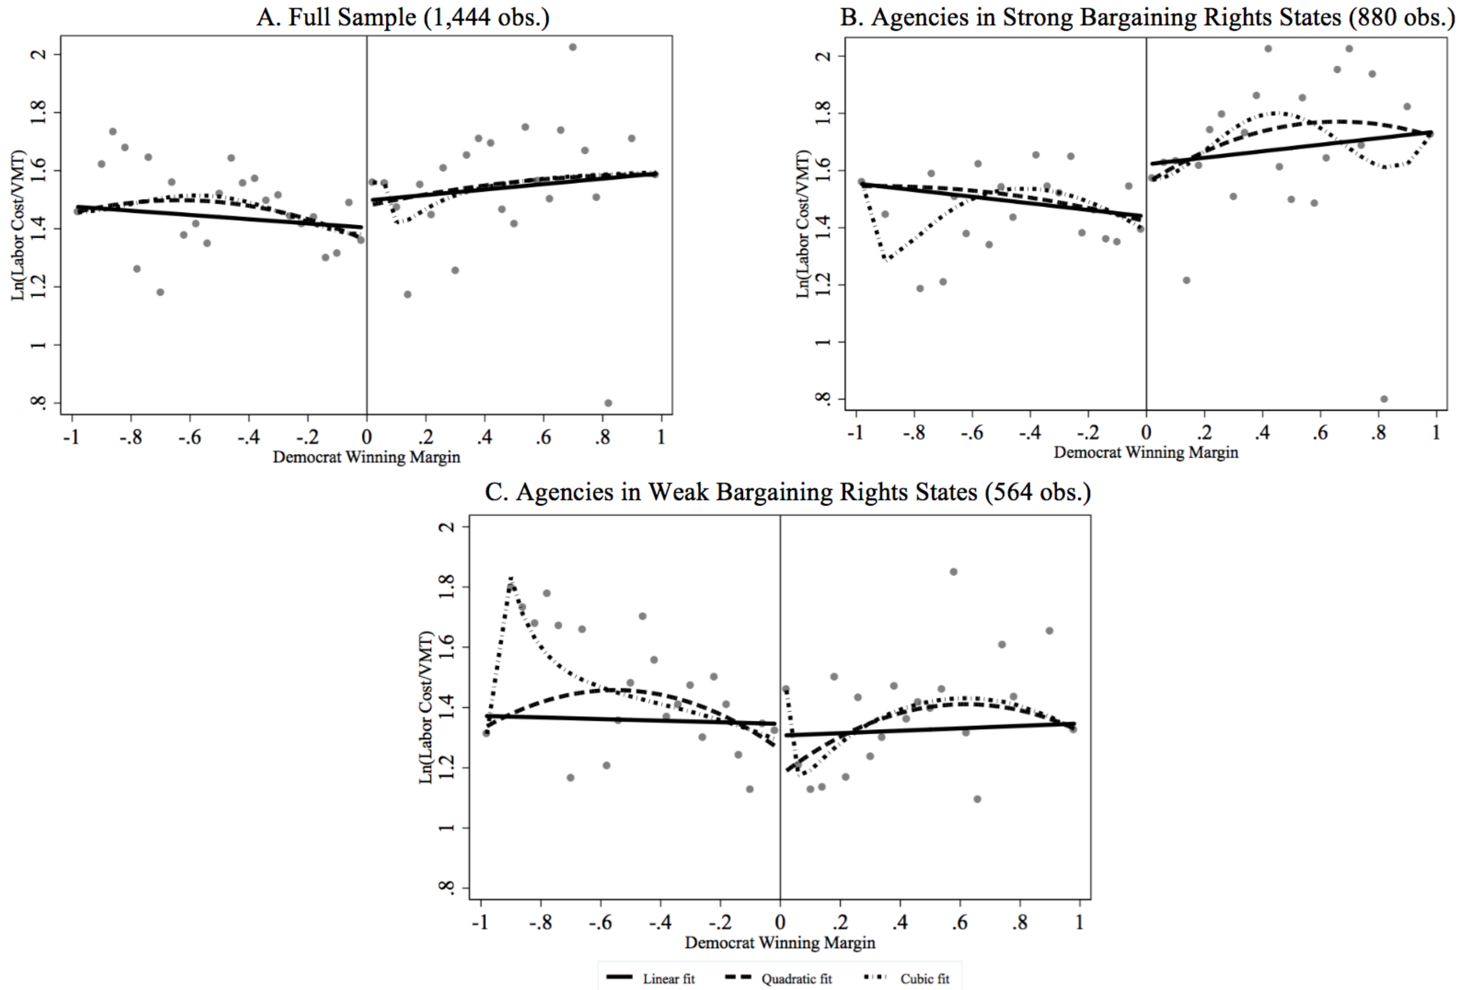

Note: Each dot corresponds to the average transit agency labor cost per vehicle miles traveled (in logs) that follows a mayoral election t, given the margin of victory obtained by Democrats in election t. Solid and dotted lines each represent the predicted values from linear, quadratic and cubic polynomial control functions of the winning margin.

Figure 2.9: Loss in US Aggregate Consumer Welfare from Inefficient Bus Transit

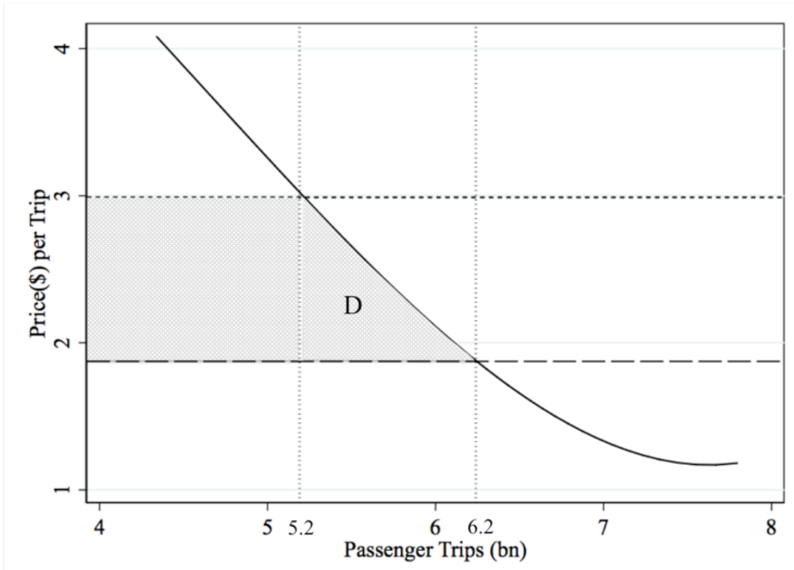

The horizontal dotted line is the supply of passenger trips under 2011 levels of privatization, where $g(\alpha_1, \hat{p}, u) = \$2.98$. At this current level of privatization, total passenger trips in 2011 were 5.2 billion. The horizontal dashed line is the supply of passenger trips under the counterfactual scenario of complete privatization, where $g(p_1, \hat{p}, u) = \$1.87$. Under the hypothetical scenario, annual passenger trips total 6.2 billion. The solid black curved line is consumer demand for passenger trips. The shaded area represents the change in consumer welfare when the price of bus transit is at its current inefficiently high level of \$2.98 per trip. The shaded area corresponds to A in Eq. B.3 of Appendix B.2. Change in US aggregate consumer surplus amounts to \$6.3 billion. Area D represents the deadweight loss. Aggregate deadweight loss amounts to approximately \$524.3 million.

Table 2.1: Summary Statistics for Top 20 Urbanized Areas (UZA)

Rank	City	Total Vehicle Miles Traveled (in '000s)	Cost Per Mile	Privatization Rate	Avg. Cost Per Mile in 1990's	Avg. Cost Per Mile in 2000's	No. Bus Transit Agencies
1	New York, NY	166,033	$18.67	8%	$13.38	$16.17	13
2	Los Angeles, CA	148,022	$9.20	21%	$8.89	$8.21	13
3	Chicago, IL	83,713	$11.17	3%	$8.93	$10.36	6
4	Houston, TX	38,796	$7.49	22%	$6.57	$6.44	2
5	Philadelphia, PA	52,569	$12.35	2%	$10.20	$10.89	3
6	Phoenix, AZ	34,945	$6.30	100%	$6.58	$6.49	5
7	San Antonio, TX	22,424	$5.91	-	$4.33	$5.26	1
8	San Diego, CA	25,415	$7.09	63%	$5.53	$6.64	3
9	Dallas-Fort Worth, TX	35,649	$7.72	0%	$8.06	$7.50	2
10	San Jose, CA	17,277	$12.63	1%	$9.57	$11.73	1
11	Austin, TX	14,323	$7.70	41%	$5.34	$6.97	1
12	Jacksonville, FL	9,366	$6.70	15%	$5.09	$5.82	1
13	San Francisco, CA	53,308	$13.58	12%	$9.29	$11.65	7
14	Indianapolis, IN	7,369	$6.06	5%	$4.79	$5.70	1
15	Columbus, OH	11,859	$7.15	-	$6.86	$7.61	1
16	Charlotte, NC	12,526	$6.31	-	$5.51	$6.07	1
17	Detroit, MI	25,582	$8.82	1%	$7.41	$8.33	2
18	El Paso, TX	7,561	$6.12	-	$4.60	$5.53	1
19	Seattle, WA	55,728	$10.06	2%	$7.06	$8.52	5
20	Denver, CO	40,644	$7.42	45%	$6.46	$6.24	1
	US Average	28,069	$9.25	16%	$7.37	$8.27	

Table 2.2: The Distribution of Hourly Pay for Full-time Bus Operators

	Boston[1]	Chicago[2]	Houston[3]	Denver[4]
Hourly Wage Distribution				
25th percentile	\$37.76	\$24.19	\$20.12	\$16.45
50th percentile	\$44.29	\$32.25	\$20.12	\$20.00
75th percentile	\$48.38	\$32.25	\$22.60	\$20.00
Mean	\$42.89	\$29.16	\$20.21	\$18.56
No. Employees	1325	4074	1343	980
Employees per 1000 VMT	0.05	0.07	0.03	0.02
Transit Agency	MBTA	CTA	MTA	RTD
Urbanized Area Statistics				
Share of Unionized Workers	8.60%	8.10%	3.20%	5.0%
Right to Work State?	No	No	Yes	No
Mean Home Price	\$459,744	\$322,764	\$225,332	\$340,703

Notes: Hourly Pay imputed from reported base pay, assuming 40 hour work weeks for 50 weeks per year. Home prices sourced from the ACCRA Cost of Living Index. Share of Unionized workers sourced from CPS.

[1] Data based on 2014 values. Source: `http://www.mbta.com/uploadedfiles/Smart_Forms/News,_Events_and_Press_Releases/Wages2014.pdf`

[2] Data based on 2014 values. Source: `http://www.transitchicago.com/foia/`

[3] Data based on 2014 values.Source: `http://salaries.texastribune.org/metropolitan-transit-authority/`

[4] Data based on 2012 values. Source: `http://www.nataliementen.com/documents/RTD\%20Salaries\%202012.pdf`

Table 2.3: Descriptive statistics for transit agencies with publicly- vs. privately-run operations

	Full Sample		RD Sample	
	Public	Private[1]	Public	Private[1]
A. UZA Characteristics				
Avg. weekly (low skill) wage ($)	595.28	597.11	598.3	594.08
Avg. road congestion index2	1.059	1.231	1.084	1.213
Avg. home price ($ '000)	320.33	381.56	321.44	385.66
Share unionized workers3	0.08	0.07	0.08	0.06
Avg. population density (pop/sq. mi)	3218.17	4197.68	3245.24	4138.58
B. Agency Characteristics				
Share Independent agency	0.59	0.8	0.57	0.8
Share City agency	0.25	0.15	0.23	0.12
Share State DOT agency4	0.15	0.05	0.2	0.08
Avg. fleet size (no. buses)	941.5	1025.36	1261.52	1233.49
Avg. fleet age (years)	7.47	7.37	7.41	7.63
Share of CNG buses5	0.1	0.2	0.09	0.19
Share of hybrid buses	0.01	0.01	0.01	0.01
C. Agency Outcomes of Interest				
Avg. share privatized VMT	0	0.2	0	0.16
Avg. labor costs per VMT ($)	6.47	6.09	7.01	6.56
No. transit agencies	251	144	105	50
No. observations	2751	955	1100	344

Notes: All values weighted by transit agency annual VMT. The same transit agency may be included in averages for public and private operations if the transit changed their privatization level on the extensive margin from one year to the next (ie from 0 to >0 or vice versa).

[1] Includes all transit agencies with at least some VMT operated by a private entity in a given year (ie privatization share>0)

[2] Road congestion index measures density of traffic across urban areas. An index greater than 1 indicates an undesirable level of areawide congestion. (Source: Texas A&M Transportation Institute)

[3] Means are conditional on share of unionized workers in strong bargaining rights states.

[4] State DOT category includes subsidiary agencies.

[5] CNG is compressed natural gas, a cleaner type of fuel than gasoline, diesel, or propane.

Table 2.4: Transit Agency Privatization Share Estimates

	(1) OLS	(2) OLS	(3) OLS	(4) LPoly
Democrat (d)[1]	-0.070**	-0.060**	-0.01	-0.100**
	(0.034)	(0.029)	(0.039)	(0.041)
Democrat * StrongBarg (d)[2]			-0.074*	
			(0.043)	
Year & State FE				
UZA & Agency Controls				
Observations	1444	1444	1444	521

Notes: Dependent variable is privatization share. Includes dominant transit agencies. (d) denotes a binary variable. All regressions include the second-order polynomials of Democrat winning margin and their interactions with Democratic mayor dummy. UZA & Agency Controls listed in Table 2.3. "LPoly" employs local polynomial RD estimation with robust confidence intervals developed by Calonico et al. (2014). LPoly estimate uses a triangular kernel and mean-squared error optimal CCT bandwidth selector of 0.264. Clustered standard errors at the UZA level are in parentheses. $^{*}p < 0.10,^{**} p < 0.05,^{***} p < 0.01$.

[1]Omitted category is Republican mayoral winner.

[2] Omitted category is weak bargaining rights state.

Table 2.5: Transit Agency Labor Cost Per Mile Estimates: Regression Discontinuity Approach

	(1) OLS	(2) OLS	(3) FRD	(4) FRD	(5) FRD	(6) FRD	(7) LPoly
Privatization share	-0.365***	-0.159**	-1.919**	-0.447	-2.163**	-1.004*	-2.265**
	(0.095)	(0.078)	(0.886)	(0.593)	(0.857)	(0.525)	(1.121)
Controls							
Year & State FE	Y	Y	Y	Y	Y	Y	
UZA & Agency Controls		Y		Y		Y	
Excluded instruments							
Democrat (d)			Y	Y	Y	Y	
Democrat * StrongBarg (d)					Y	Y	
Observations	1444	1444	1444	1444	1444	1444	549
Cost reduction from 100% privatization*	30.9%	15.0%	90.1%	46.4%	90.2%	68.1%	94.0%

Note: Dependent variable is log(labor cost per VMT). Sample includes dominant transit agencies only. All regressions control for quadratic polynomial of winning margin and their interaction with Democrat mayor indicator. (d) denotes a binary variable. "LPoly" employs local polynomial Fuzzy RD estimation with robust confidence intervals developed by Calonico et al. (2014). LPoly estimate uses a triangular kernel and mean-squared error optimal CCT bandwidth of 0.294. UZA & Agency Controls listed in Table 2.3. Standard errors clustered at UZA level. $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$
*Because Eq. 2.9 takes the semi-logarithmic functional form and the privatization rate ranges from 0 to 1, a consistent estimator for the percentage impact from full privatization (e.g., privatization going from 0 to 1) on the labor cost is $100 * \left[\exp(\hat{\gamma} - \frac{var(\hat{\gamma})}{2}) - 1\right]$. The values in the final row represent the expected cost reduction from an agency going from 0% privatized operations to 100%.

Table 2.6: Full Time Employee Headcount and Union Strength

	(1)	(2)	(3)
Right to Work (RtW) State (d)[1]	-0.073***	-0.069***	-0.01
	(0.01)	(0.013)	(0.019)
Strong Barg. State (d)[2]		0.006	0.056***
		(0.013)	(0.017)
RtW State x Strong Barg State (d)			-0.108***
			(0.025)
Year FE	Y	Y	Y
Observations	3395	3395	3395
R2	0.012	0.012	0.017

Notes: The dependent variable is the log of full time employee headcount per publicly operated VMT. All regressions are OLS. (d) denotes a binary variable. Sample includes subset of transit agencies which reported employee work hours and counts. Standard errors in parentheses. $^{*}p < 0.10,^{**} p < 0.05,^{***} p < 0.01$.
[1] Omitted category is a non-right-to-work state.
[2] Omitted category is a weak bargaining rights state.

Table 2.7: Effect of Right to Work Laws on Labor Costs using State Borders

	(1)	(2)	(3)	(4)
Right to Work State (d)1	-0.536***	-0.397***	-0.680***	-0.298***
	(0.052)	(0.048)	(0.017)	(0.054)
County Pair FE	Y	Y	Y	Y
Year FE		Y	Y	Y
State FE			Y	Y
UZA & Agency controls				Y
Observations	206	206	206	206
R2	0.53	0.709	0.709	0.902

Notes: Notes: The dependent variable is log(labor cost per VMT). All regressions are OLS. Observations consist of 21 transit agencies in 17 counties that border another state with differing Right to Work legislation. UZA & Agency Controls listed in Table 2.3. (d) denotes a binary variable. Standard errors are clustered by transit agency. $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

[1] Omitted category is a non-right-to-work state.

Table 2.8: Dominant Transit Agencies Effects on Operating Costs with City-Year Fixed Effects

	(1)	(2)	(3)	(4)
Dominant TA ("Dom") (d)[1]	0.275***	0.362***	0.356***	0.268***
	(0.051)	(0.053)	(0.06)	(0.077)
Dom x RtW State (d)[2]		-0.338***	-0.335***	-0.278***
		(0.059)	(0.058)	(0.055)
Dom (d) x UnionShare			0.052	-0.145
			(0.369)	(0.367)
Dom x StrongBarg (d)[3]				0.116**
				-0.055
UZA x Year FE	Y	Y	Y	Y
Observations	3706	3706	3706	3706
R2	0.538	0.579	0.579	0.582

Notes: Dependent variable is log(operating costs per VMT). All regressions are OLS. The "Dominant" TA has the highest average annual VMT in a given UZA. (d) denotes a binary variable. Standard errors clustered at UZA level in parentheses. $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

[1] Omitted category is a non-right-to-work state.

[2] Omitted category is a weak bargaining rights state.

[3] Omitted category is a weak bargaining rights state.

Table 2.9: Welfare Loss Estimates for Three US Cities in 2011

City (Agency)	Current (Predicted) Cost / VMT	Privatized (Predicted) Cost / VMT	Consumer Surplus ($mm)	DWL ($mm)	DWL % of Public Funding*	Strong Bargaining Rights?
Boston (MBTA)	$13.59	$8.13	$163.68	$14.77	1.7%	Yes
Chicago (CTA)	$11.29	$6.64	$297.82	$26.31	3.7%	Yes
San Antonio (VIA)	$7.18	$4.22	$59.61	$4.19	2.5%	No
US Total	$6.63	$4.16	$6,319.34	$524.35		

Notes: Predicted values for Current Cost/VMT and Privatized Cost/VMT generated using the RD regression specification. The dependent variable is total operating costs per VMT. Regression controls include all those listed in Table 3 with the addition of city-average diesel and natural gas prices. Privatized Cost/VMT counterfactual assumes privatization share is 100%. Consumer Surplus and DWL correspond to shaded area and area D in Figure 2.9, respectively.

*Sources: MBTA Fiscal 2012 Audited Financial Statements Pg 5, Net Nonoperating revenue (`http://mbta.com/uploadedfiles/About_the_T/Financials/113990_12_MBTA_FS-FINAL.pdf`); CTA President's 2012 Budget Recommendations, 2011 Forecasted Total Public Funding Pg 22, (`http://www.transitchicago.com/assets/1/finance_budget/2012_Budget_Book_for_Web.pdf`); VIA 2012 Comprehensive Annual Financial Report; Sept 30, 2012 Pg. 45, Sales Taxes + Grants Reimbursement (`http://www.viainfo.net/Organization/Docs/2012CAFR.pdf`)

CHAPTER 3

ROAD RATIONING POLICIES AND HOUSING MARKETS

3.1 Introduction

Several major cities around the world have implemented road rationing policies to address increasingly harmful levels of air pollution and traffic congestion.[1] These policies generally restrict people's ability to drive within city limits on particular days of the week and at peak commute times. Figure 3.1 shows how permanent adoption of alternate-day-based road rationing policies has accelerated since the 2000's among some the world's largest urban centers. Cities are a physical manifestation of people's desire to eliminate transportation costs. Thus, driving restrictions will likely affect not only pollution and congestion, but the ways in which people sort in space. This paper demonstrates how a shock to transit costs served to alter the spatial distribution of different income groups in the context of Beijing, China. Our results are the first to demonstrate, empirically, that policies aimed at mitigating air pollution through driving restrictions can induce equity concerns of access to economically important centers within cities.

Rapid urban growth in China since 1990 has increased automobile use within its largest cities at a staggering rate. Between 2005 and 2015, the rate of car ownership

[1] Air pollution is a primary concern for many urban areas. Over 80% of the world's major cities fail to meet World Health Organization guidelines on air quality (World Health Organization 2019). Emissions from vehicles is a particular concern in several developing nations, contributing to nearly 50% percent of particulate matter in Beijing (Viard and Fu 2015) and nearly 20% in Delhi (Jain 2018).

outpaced population growth by 3 to 1 in Beijing. The Chinese government responded to degrading air quality and road congestion through several transit-oriented policies, including heavy investment in its subway network and a license plate-based road rationing policy in July of 2008 (hereinafter, CDR or "car driving restriction"). The CDR limited car owners from driving their car one day per week according to the last digit of their license plate. We exploit this shock to commute costs to test predictions of urban land use models developed by Alonso, Muth, and Mills (hereinafter, "AMM") (Alonso et al. 1964; Mills 1967; Muth 1967), and extended by LeRoy and Sonstelie (1983). We show that the CDR increased the premium for both subway station proximity, and proximity toward Beijing's central business districts. Both of these price effects served to decentralize lower income groups away from subway stations and the central business districts.

Our analysis uses detailed, micro-level data on real estate transactions and buyer demographics of residents of Beijing. Using changes in housing prices as a reflection of changes in the demand for location amenities (Rosen 1974), we test how the driving restriction policy affected demand for both subway proximity and central business district proximity. We further explore how these housing market effects impacted the relative likelihood of income groups to sort near subway stations and Beijing's central business districts. These data allow us to include a comprehensive set of controls on housing characteristics in addition to neighborhood-level fixed effects to account for unobserved differences in neighborhood desirability. Our housing price estimates effectively compare observationally similar housing units within the same neighborhood that vary in their relative distance to the nearest subway station and the central business district. We exploit the driving restriction as a city-wide unexpected increase in the cost of driving to test how the premium for subway and central business district proximity changes over time. Our preferred estimates show that the driving restriction increased demand for subway proximity by approximately 3.6%, and the demand for central business district proximity by 1.3% per kilometer relative to pre-policy time periods. Pre-trends shows that the premium for proximity is stable in the periods prior to the driving restriction, but shift significantly in the aftermath of the policy. The shift in the demand for proximity does not appear to be driven by unobserved correlated shocks or broader investment in areas that receive new subway stations. We additionally utilize novel micro data on household income and housing locations to explore how the driving restriction impacted the residential location choices of income groups relative to one another. We find modest, but significant reductions in the composition of lower income households close to subway stations as well as Beijing's central business districts following the driving restriction.

This paper contributes to existing work on the housing market responses to driving restrictions in two important ways: First, this paper shows that the increase in housing prices in transit accessible areas found in prior literature impose additional unintended impacts on the accessibility of public transit for lower income individuals. Second, this paper exploits a large, highly detailed data set that represents 17% of all housing transactions in Beijing from 2005-2011. The wide variation available in these data allow us to focus on effects of the driving restriction within a fine geographic scale, and thereby remove any fixed differences across city neighborhoods that relate to the housing market or sorting decisions. With this novel data set, we show that prior estimates on housing price responses are attenuated by at least 50% because they do not account for unobserved differences in amenities that vary with subway locations. After accounting for such confounds, we find that the bid-rent gradient for subway proximity is much more responsive to driving restriction policies than prior literature suggests. This is of first order concern when assessing the cost benefit of investing in public transit.

While tests of the monocentric city model are not, themselves, a novel line of inquiry in the urban economics literature, their application to the setting of a rapidly developing city such as Beijing is important for the following reasons. First, in the setting of Beijing, with over 19 million inhabitants, the assumptions of competitive bidding for housing that are required by the monocentric city model are not so extreme an assumption as in several US housing markets, where the bids for any particular housing unit may be relatively thin. Consequently, quantitative predictions from the price-distance gradient may be more appropriate to the massive urban markets of China. Second, transit mode choice is largely stratified by income group in Beijing, unlike most US or European cities. Personal automobiles are costly in Beijing, highly taxed, and restricted to only those who can afford to enter a lottery system for purchase of a license plate. Lower income individuals rely on bus, subway, or cycling. This feature makes Beijing a particularly useful setting for testing predictions of LeRoy and Sonstelie (1983) because low income groups are very likely to utilize different transit modes from high income groups. The driving restriction provides variation in the transit mode choice utilized by the high income group. Consequently, we are effectively able to analyze two “natural” experiments—the subway expansion, and the car driving restriction—to test how these countervailing effects impact sorting and housing choice across the income distribution. Lastly, urban driving restrictions are enforced in more and more cities around the world, particularly in the developing world where air pollution from automobile exhaust is a particularly salient issue. The effects of these omni-present transit policies on housing markets and location choice is not well understood, but highly policy relevant.

Urban driving restrictions have been the focus of a small, but growing body of literature. Most of these studies have focused on the effectiveness of driving restriction at meeting their primary goals of reducing congestion and air pollution, with mixed results. High compliance locations appear to experience improvements to air quality following a driving restriction (Carrillo et al. 2016; Viard and Fu 2015; Wolff 2014; Chen et al. 2013); whereas locations with low compliance or the possibility of substituting behaviors experience little change in air pollution (Davis 2008; Zhang et al. 2017; Eskeland and Feyzioglu 1997). The consensus in Beijing is that the CDR was very effective at reducing air pollution and traffic congestion (Lu 2016; Viard and Fu 2015) due to high compliance and rigorous enforcement.

In contrast to the body of literature above, our paper explores how an urban driving restriction impacts the demand for housing and the subsequent re-sorting of high relative to low income groups. Our paper is most similar to Xu et al. (2015) which finds that demand for housing near subway stations increased in the six months following Beijing's driving restriction. However, our estimates suggest the magnitude of this effect is more than double that found by Xu et al. (2015). We attribute this difference to application of a more representative dataset. We exploit detailed information on over 82,000 transactions—roughly 13 times that of the previous study—that allows for a robust set of controls and geographic fixed effects.

Lu (2018) also investigates how a second driving-based policy, Beijing's 2012 license plate lottery, impacts the housing market. The license plate lottery substantially reduced the quota of available license plates, thereby restricting the ability of people to purchase cars. The author finds that this lottery policy increased housing prices within Beijing's fourth ring road, and at locations proximate to public transit stations. While the housing market impacts of Lu (2018) are similar to those we find for the CDR, we argue that the CDR provides a context better suited to test predictions of the AMM model because the CDR altered transit costs for individuals *already* driving cars; whereas the license lottery did not. This policy distinction is important because the marginal individual impacted by the CDR is someone who already owned a car and will, therefore, have to change their mode of travel *ex post*. In contrast, the marginal person impacted by the license plate lottery is one who would have liked to own a car, but cannot following the policy. Thus, there is no transit mode switch as a consequence of the lottery. In short, the assumptions required to apply the AMM to the CDR policy are much weaker. The CDR, thus, provides a credible setting to predict income-based sorting as consequence of a city-wide driving restriction.

This paper also relates to prior studies concerned with the positive relationships between inter-city location choices among the poor as a function of public infrastructure access (Glaeser et al. 2008; Brueckner and Rosenthal 2009; Brueckner et al. 1999; Baum-Snow and Kahn 2000; Waxman 2017). In general, our results underscore the importance of transportation infrastructure and transit technology as a determinant of urban spatial structure. Beijing's investment in public transportation throughout the city allowed lower income individuals to decentralize, and sort near subway stations, rather than cluster near the city center where public transit was formerly concentrated. These data patterns are consistent with findings of Glaeser et al. 2008, who find public transit investment in US cities attracts lower income residents and is an important determinant of historic centralization of the urban poor. The Beijing context provides a novel natural experiment that empirically re-affirms these prior finding through a new channel. While investment in public transit infrastructure attracts the poor, we demonstrate that increasing commute costs for the wealthy can actually disperse the poor as they are outbid by the rich.

3.2 Data & Preliminary Facts

Our empirical analysis requires information on housing prices, residential locations, and household income. We assembled this information using individual real estate transactions and mortgage loan applications, sourced from two major Beijing real estate firms and a government-backed loan system known as the Housing Provident Fund,[2] respectively. The real estate data comprises approximately 17% of all housing

[2]The Housing Provident Fund (HPF) operates as a government-backed credit market to encourage home ownership. In absence of a formal credit market, people can access loans through their employers (Tang and Coulson 2017). At participating employment firms, each employee and his/her

purchases within Beijing from 2006 through 2012. Our sample of mortgage contracts comprises the majority of the mortgage market in Beijing during the sample period, though the precise market share is difficult to quantify. While our mortgage data include all mortgage loan applications administered through the Housing Provident Fund, mortgage loans made to individuals at non-salaried jobs, individuals working part-time or unemployed, or the very wealthy will not be represented in our data.

The Beijing government enacted several city-wide policies throughout our sample period that likely impacted the housing market and demand for automobiles.[3] Our main analysis uses the transactions occurring between July of 2007 and July of 2009, one year before and after the CDR policy began, to mitigate spurious correlations with these other policies. Each housing transaction contains the sale or rental price, as well as descriptive information on the housing unit (including number of bedrooms, floor level, decoration level, types of appliances, etc.), and information on the housing complex (including geographic location, total size, parking availability, green space, proximity to key schools, etc.). Figure 3.2 shows the spatial distribution

employer have to contribute a specific percentage of his/her monthly income from the employer to the HPF account. The employees can then obtain a mortgage loan with a subsidized interest rate for home purchase (about 1.5 percentage points, or nearly 30%, lower than the commercial banks' mortgage rate; this interest rate is determined by the Ministry of Housing and Urban-Rural Development and the same fixed rate applies to all borrowers). Virtually all eligible home buyers would apply for this mortgage first before going to other sources of funding. The data cover the universe of all home purchases in the city that made use of an HPF mortgage loan (209,861 in total) from 2006 to 2013. Mortgage refinancing is uncommon in China, and in any case, there are no refinancing observations in the sample. As a result, each mortgage contract refers to a housing transaction.

[3]For example, in January of 2012 the government implemented a lottery system for purchases of automobile license plates in order to limit the total vehicle fleet on Beijing's roadways (see Lu (2018) for an analysis of the license plate restriction on Beijing's housing market). In April of 2011, the government enacted an anti-speculative policy that restricted home purchases for natives, and prohibited home purchases for non-natives of Beijing (see Sun et al. (2017) for analysis of the home purchase restriction on Beijing's housing market).

of these real estate transactions throughout the city, and relative to the subway network. The mortgage loan data also provide information on housing unit price and descriptive information on the housing complex building. Importantly, the mortgage data provide detailed information on the loan applicant's demographic characteristics, including their income, education levels, and place of employment.[4]

Table 3.1 provides summary statistics on a subset of the variables available within each of our data sets. The purchase price varies considerably more in the real estate data as compared to the mortgage data set. The mortgage data housing units are also further from the city center and further from subway stations, on average. This is partially due to the selection of individuals represented by the mortgage data, which over represents the middle class demographic. Home buyers able to utilize the HPF loan system are less likely to be entrepreneurs or independently wealthy.

Transit mode choice is stratified by income in Beijing. While we do not observe transit mode choice among the home buyers in our data, the Beijing Transport Annual Report provides aggregated statistics on mode choice by income group which supports our prediction that the CDR binds mainly for the wealthy who can afford cars. As of 2010, over 40% of earners in the top tenth percentile relied on cars to commute to work, while approximately 25% of earners in the bottom 90th percentile relied on cars to commute. For public transit, just 14% of the top earnings decile relied on subway or bus as of 2010, whereas 20-21% of individuals in the bottom 90th percentile chose to commute by subway or bus. A large portion of individuals

[4] In a subsequent analysis with Li and Zhang, we exploit these data to understand commute differences by gender.

in Beijing (20-40%) who are mostly comprised of lower-income earners walk to work (Institute 2010). Subway transit is heavily subsidized by the Beijing government and is quite inexpensive. Prior to 2014, subway fares were fixed at 2 yuan (or about $0.30) per trip.[5]

Record high pollution levels and the on-coming 2008 Olympic games spurred Beijing to introduce a driving restriction on July 20, 2008, which banned personal use of cars within the 5th Ring Road one weekday per week from 7am until 8pm on the basis of the last digit of the vehicle's license plate.[6] Viard and Fu (2015) found that Beijing's driving restriction reduced particulate matter (PM_{10}) by 21%, with strong compliance and little evidence of inter-temporal substitution (i.e., driving more during non-restricted hours) or an increase in total number of vehicles in circulation, as in Davis (2008). In Beijing, it is difficult to evade detection because cameras throughout the city (as opposed to police) monitor the plates on vehicles. If an individual violates the restriction, they are fined roughly $30 per violation. Further, Beijing restricts people's ability to purchase a second car, and purchase of a first car is regulated by a lottery system. The limited scope for noncompliance or substituting behavior were key to the success of Beijing's CDR in improving air

[5]Fares became distance-based after 2014, however trips within the 4th Ring Road (under 22 kilometers) were still relatively cheap, averaging about 5 yuan per trip.

[6]Terms of the driving restriction changed three times after its introduction in July of 2008. From July through September of 2008, the policy was more strict, allowing cars to drive only every other day during the week, following an odd-even policy. These restrictions ended temporarily in September 20th, but were reinstated on October 11th. The new policy restricted vehicle use one day per week, extending from 6am through 9pm. On April 11, 2009, the daily restriction period narrowed to 7am to 8pm. The one-day-per week policy is structured such that two out of ten possible plate numbers ((0,5 (1,6), (2,7), (3,8), (4,9)) are restricted each weekday. Assignment of the pairs rotates every thirteen weeks. Viard and Fu (2015) provides a review of this policy and prior work that have investigated the effects of driving restrictions on air pollution.

quality. Given the strong compliance, how did this policy impact the distribution of wealth through Beijing?

We first review patterns of economic activity and location sorting within Beijing. Figure 3.3 shows the location of Beijing's major employment centers. Beijing has at least six major employment centers located at various quadrants of the city. The city's geographic center is not a traditional business district, which is generally an assumption of the monocentric city model, but mainly a cultural and consumer-oriented district. For this reason, we define the Beijing "central business district" (CBD) as a relative measure and assign the closest major employment center as a given housing unit's CBD. In subsequent robustness checks, we find our results are generally insensitive to alternative definitions of Beijing's CBD.

Figure 3.4 shows the relationship between income and distance to the central business district in Panel A and the nearest subway station in Panel B, respectively. The decentralization of the lower-income relative to the higher-income is prevalent in the raw data. Both gradients have become steeper over time, particularly for subway proximity. In the most recent period, Beijing appears similar to older US and European cities in its demographic spatial structure as average income declines with distance. However, these patterns are a marked shift from the pre-driving restriction period, when the relationship between income and CBD proximity is less pronounced. Income and subway proximity appear almost uncorrelated in the pre-period in Panel B. The slight positive relationship is consistent with the theoretic predictions of Glaeser et al. (2008), whereby lower income individuals tend to sort near transit oriented locations.[7]

3.3 Theoretical Framework

We predict how a city-wide road rationing policy will affect the sorting of low relative to high income individuals through the stylized AMM monocentric city model (Alonso et al. 1964; Mills 1967; Muth 1967), and its extension by LeRoy and Sonstelie (1983).

3.3.1 A model of sorting near the CBD

Consider a monocentric city with a central business district (CBD) where all residents work, supplying l units of labor, to earn wl income. Residents live outside the CBD in the residential district of the city. Commute time increases monotonically with distance from the CBD. Consumers decide where to live, choosing δ so as to maximize consumption of housing h with price P, the numeraire good z with price one, and time spent outside of work, s. Households commute distance δ to work and pay t, the unit time cost of commuting, and τ, the fixed cost of commuting. Our empirical approach focuses on price and sorting effects over a relatively short time horizon of two to four years, consequently we assume that the city boundary is fixed at $\bar{r}$, housing supply is fixed, and the city is closed, without in or out migration. Prior

[7] Appendix Figure C.2 plots the income-distance gradient for alternative definitions of Beijing's CBD. In each definition, the post-period gradient becomes steeper.

to any policy intervention on behalf of the government, residents face the following maximization problem:

$$\max_{h,\delta} U(h(\delta), z(\delta), s(\delta)) \text{ s.t.} \tag{3.1}$$

$$wl(\delta) - \tau = P(\delta)h(\delta) + z(\delta) \tag{3.2}$$

$$l(\delta) + s(\delta) + t\delta = 1 \tag{3.3}$$

The first constraint is an individual's budget constraint, while the second is their time budget constraint. Let $wt\delta$ be the time cost of commuting distance δ, and τ be the fixed, pecuniary cost. In subsequent analyses, we will be interested in how the time cost of commuting is affected by a driving restriction, therefore we can denote t as a function of days in which driving is restricted, $t(n)$. In equilibrium, residents choose a housing location δ such that:

$$\frac{\mathrm{d}P(\delta)}{\mathrm{d}\delta} = -\frac{wt(n)}{h(\delta)} \tag{3.4}$$

Thus, individuals trade off consumption of housing at further distances from the CBD for shorter commutes and more leisure time. The price of housing falls with distance enough to compensate individuals for their longer commutes. Groups facing different preferences for housing consumption h and time costs wt will have their own distinct bid-rent gradient (i.e., Eq. 3.4). LeRoy and Sonstelie (1983) develop the case where a high-income group, possessing greater opportunity cost of time utilizes a more expensive, but faster transit mode relative to the lower income group. This setting is quite applicable to Beijing where wealthier people use cars and the lower income rely on public transit.

In a simplified version of the Beijing scenario where the poor use public transit and the rich use cars, the poorer group will have a steeper bid-rent gradient than the rich and, therefore, live closer to the city center if and only if:

$$\frac{h_r(\delta)}{h_p(\delta)} > \frac{t_{car}(n)}{t_{pub}(n)} \times \frac{w_r}{w_p} \tag{3.5}$$

where subscript p denotes the housing consumption and wage of the poor and r the same for the rich.[8] Beijing's transit policies during our study period serve to change the ratio of transit time costs for the car users relative to the public transit users, $\frac{t_{car}(n)}{t_{pub}(n)}$. The driving restriction increases the cost of commuting for the rich by increasing t_{car}, whereas the subway expansion reduces the cost of commuting for the poor by decreasing t_{pub}. Both of these policies make the right-hand side larger, thus Eq. 3.5 is less likely to hold following Beijing's driving restriction. This logic tells us that higher-income individuals will find housing close to the city center more attractive, and will consequently move closer to the city center following the driving restriction. Similarly, the subway expansion will reduce the comparative advantage of central-city housing for the poor, and they will move further from the CBD relative to the rich.

[8]We can also interpret Equation 3.5 in terms of elasticities. Let $y = wl$ denote income. Let the income elasticity of housing demand be $\epsilon_{h,y} = \frac{y}{h}\frac{\partial h}{\partial y}$ and let the income elasticity of the time cost of commuting be $\epsilon_{tw,y} = \frac{y}{tw}\frac{\partial tw}{\partial y}$. If Equation 3.5 holds, then $\epsilon_{h,y} > \epsilon_{tw,y}$. In other words, the change in one's demand for housing as wealth increases must be greater than the change in one's relative value of time as wealth increases.

3.3.2 A model of sorting near subway stations

Next, the government implements a driving restriction such that residents must commute to work via subway at least one day per week. Now, time spent away from work $s(\delta)$ is a function of both leisure time $\lambda(\delta)$ as well as the time it takes to walk to the nearest subway station. Let n define the number of times that residents would prefer to drive, but must take the subway. n includes both government restrictions on driving, as well as heavy traffic days. Let ω be walking speed; and x be the distance from place of residence to the nearest subway station. Then we can define time spent away from work as:

$$s(\delta) = \lambda(\delta) + n\frac{x}{\omega} \tag{3.6}$$

If $n = 0$, then time spent away from work is exactly equal to leisure time. However, if $n > 0$, then $s(\delta)$ depends additionally on the walking distance to the nearest subway stop. Under a driving restriction regime, the resident's choice is again to maximize consumption of housing, the numeraire good, and leisure time, subject to income and time constraints. We assume the amount of housing consumed h does not vary once the consumer selects δ, their distance from the CBD. However, conditional on δ, the unit price of housing also depends upon n and x, so that the price of housing is $P(x, n|\delta)$. To simplify notation, we present the consumer utility maximization decision *conditional* on the choice of δ as follows:

$$\max u(h, z; \lambda) \text{ s.t.} \tag{3.7}$$

$$wl - \tau = P(x, n)h + z \tag{3.8}$$

$$l + \lambda + n\frac{x}{\omega} + t\delta = 1 \tag{3.9}$$

By substituting the budget constraints into equation (3.7), the optimization problem can be written as a choice of only x:

$$\max_x u(wl - \tau - P(x, n)h; \;\; 1 - (t\delta + l) - \frac{nx}{\omega}) \tag{3.10}$$

Within a small radius of δ, equilibrium is achieved through the spatial variation in P and consumers substituting consumption x and more leisure time λ through a shorter walk to the station.

Denote μ_i as the derivative of Eq. (3.10) with respect to the ith argument, and denote the value of time as the marginal utility of substitution between leisure and the numeraire good: VOT= $\frac{\mu_2}{\mu_1}$. The first order condition of (3.10) can be written as:

$$\frac{\partial p}{\partial x} = -\mathrm{VOT}\frac{n}{\omega h} \tag{3.11}$$

The poor will live closer to stations if:

$$\mathrm{VOT}_p\frac{n_p}{\omega h_p} > \mathrm{VOT}_r\frac{n_r}{\omega h_r}. \tag{3.12}$$

We assume the rich and poor have similar walking speeds, but face different effective driving restriction days. The driving restriction is more likely to bind for the rich relative to the poor because the rich are more likely to rely on cars to commute as a group relative to the poor. As n_r increases relative to n_p, the driving restriction makes Eq. 3.12 less likely to hold. This logic tells us that higher-income individuals will move closer to subway stations following the city-wide driving restriction.

An advantage of our setting is that we isolate shifts in the bid-rent gradients, as opposed to their actual slopes. Several factors can confound empirical estimates of $\frac{\mathrm{d}P(\delta)}{\mathrm{d}\delta}$ in Eq. 3.4 or $\frac{\partial P}{\partial x}$ in Eq. 3.11, such as proximity to high quality schools or proximity to recreational amenities (i.e., parks or restaurants). Omission of such unobserved factors will bias the price-distance gradients, either downward in the case of positive correlation of x and desirable amenities; or upward in the case of a positive correlation of x and disamenities, such as the prevalence of noise and congestion near the city center or subway stations.[9] While use of spatial fixed effects can correct for some of the omitted variable bias, there is a problematic trade off between employing fixed effects at a fine-enough scale to remove all spatially-dependent amenity variation, and identification of the bid-rent gradient.

Beijing's driving restriction policy exogenously shifted both gradients. Consequently, we can identify changes in demographic composition through estimating the second derivative of both equations with respect to n as follows:

$$\frac{\mathrm{d}^2 p}{\mathrm{d}x\mathrm{d}n} = -\frac{wt'(n)}{h(\delta)} \tag{3.13}$$

and

$$\frac{\partial^2 p}{\partial x \partial n} = -\mathrm{VOT}\frac{1}{\omega h} \tag{3.14}$$

Our empirical approach will estimate changes in both bid-rent gradients, and then test predictions of the model on sorting of lower relative to higher income households near the CBD and subway stations.

[9]Beijing's subway network is entirely underground, thus the infrastructure itself does not alter the visual appeal of a particular location. However, there may be some dis-amenity from living directly beside a subway station due to heavier foot traffic or noise.

3.4 Empirical Approach

Our goal is to estimate how Beijing's transit policies affect the locations of lower relative to higher income individuals. Expansion of Beijing's subway network will lower the cost of commuting by public transit and the driving restriction will increase the cost of commuting by car. Both of these effects will serve to de-centralize the poor relative to the rich.

We first explore how Beijing's transit policies impact demand for housing near the city's business districts and near subway stations. The residential locations of the rich relative to the poor will be driven by changes in their willingness to pay for proximity. Figure 3.5 plots the bid rent gradients before and after the driving restriction. The price premium per square meter for subway access clearly increases following the driving restriction in Panel B. The premium for access to the central business district in Panel A is less striking, but still demonstrates a slight tilt from the pre-CDR period. These price dynamics are intuitive if both the rich and poor are competing for housing proximate to subway stations, but less so for proximity to the central business district. The poor are not impacted by the driving restriction, thus do not compete for housing close to the city center. Both rich and poor, however, prefer housing closer to subway stations, all else equal.

Our identification strategy exploits variation in the housing price-distance gradient before and after the CDR policy within a neighborhood (a "jiedao") of Beijing. Figure 3.6 demonstrates our within-neighborhood variation. Within our data sample, each jiedao contains an average of 45 housing complexes, or about 2,500 transactions from 2007 through 2009. A jiedao is the smallest administrative unit at which socioeconomic data are collected in China, outlined by the thick black lines. The average size of a jiedao in Beijing is about 43,000 people, spanning three square miles (Ding et al. 2005). Our sample covers roughly two-thirds of the 300 jiedao within Beijing. By identifying the gradient change estimates off of time-wise variation within a small geographic unit, we are able to control for unobserved differences in amenities across Beijing neighborhoods, such as the existence of high quality shops or restaurants.

Let i index housing units, c index housing complexes, j index jiedao neighborhoods, and t index time. We estimate the effect of the CDR policy on the distance-to-subway (and distance-to-CBD) price premium through the following specification:

$$\ln(p_{ijt}) = \beta(Km_{it} \times R_t) + \delta Km_{it} + \rho R_t + \mathbf{X}_{ijt}\Theta + \gamma_j + \tau_t + \varepsilon_{ijt} \tag{3.15}$$

where p is the price of housing per square meter, Km measures the distance to the nearest subway station or business district, and R is a binary variable equal to one in periods after July 20th, 2008 when the CDR went into effect. $\mathbf{X}$ is a vector of controls for housing unit and building complex attributes and γ_j and τ_t are jiedao and time fixed effects, respectively (we alter the time unit in various specifications). The parameter β provides the additional price premium demanded for moving one kilometer closer to either a subway station or the central business district as a consequence of the driving restriction.[10] The identifying assumption for β to provide a causal effect is that housing prices would have trended similarly

[10] We choose a log-linear specification for ease of interpretation, however we show in Appendix Table C.2 that our results are insensitive to a log-log specification used in Xu et al. (2015).

for close relative to far housing units in absence of the driving restriction. We test this assumption in the following section by examining trends in the distance-price premium in periods prior to the driving restriction.

For this approach, we do *not* need to assume that location characteristics, like quality of restaurants and shops, are uncorrelated with either distance to the city center, or the nearest subway. In other words, we do not need to assume that the subway network or location amenities are randomly assigned across neighborhoods. The driving restriction policy allows us to credibly identify β, the *change* in the price gradient because the driving restriction was unanticipated.

3.5 Results

We now present evidence on the connection between housing demand and transit policies in Beijing. Expansion of Beijing's subway network will reduce transit costs for those relying on public transit, while the driving restriction will increase transit costs for drivers. The model of LeRoy and Sonstelie (1983) predicts that changes in t_{pub} and t_{car} should shift the price-distance gradient, as well as the residential location choices of those relying on public transit relative to those relying on personal car travel.

3.5.1 Transit Policies and the Housing Market

First, we test the assumption that housing prices for proximate relative to distant housing units would exhibit common trends absent the driving restriction. We estimate a dynamic version of Eq. 3.15 where the effect of subway proximity is allowed to vary in each quarter of our sample, and thereby allows for visual examination of pre-trends in the data. Let i index housing units, j index jiedao neighborhoods, and t index one of 28 quarters.

$$\ln(p_{ijt}) = \sum_{t=q1(\ni q15)}^{q28} \delta_t (Km_{it} \times D_t) + \alpha Km_{it} + \mathbf{X_{ijt}}\theta + \gamma_j + \tau_t + \varepsilon_{ijt}. \tag{3.16}$$

The coefficient δ_t measures the difference in the housing price-distance gradient in quarter t relative to the third quarter of 2008 (the 15th quarter in our sample). Figure 3.7 presents evidence that the difference in price per square meter across proximate relative to distant housing units did not differ significantly from the quarter when the driving restriction was enacted. However, starting in the first quarter following the CDR in July of 2008, the premium for subway proximity increased to approximately 5% per kilometer. Around April of 2011, there is an additional uptick on the price premium, which corresponds to a substantial restriction in the government's circulation of license plates and a lower likelihood of winning the license plate lottery (Yang et al. 2014).

In Figure 3.8, we show these patterns hold after allowing for a non-parametric relationship between subway proximity and price. Instead of imposing a linear relationship between distance and price as in Eq. 3.16, we test for differences between housing units within 3 kilometers and those over 3 kilometers of a subway in each quarter of the study period. Results of the non-parametric specification demonstrate, again, that the premium for subway proximity increased approximately 2-5% following the CDR.

Our regression estimates reaffirm that the price-subway proximity gradient became steeper following the road rationing policy. Results are generally robust to alternative controls, and finer time-location fixed effects. In Table 3.2, subway proximity commands a 1 to 5% price premium per kilometer prior to the CDR policy. This level estimate may be biased if stations are sited endogenously, however we interpret the positive sign as consistent with subway access being valuable to city residents. The interacted term shows that after the driving restriction, housing units one kilometer closer to a subway station sell for approximately 3% more than a comparable housing unit one kilometer further from a subway station. Comparing column (1) to column (2), jiedao fixed effects increase the precision of the policy effect substantially and serve to absorb nearly all cross-section variation in subway proximity and prices. Column 4 uses year-by-month fixed effects, such that variation is driven by transactions within the month of July 2008. Lastly, Column 5 applies district by quarter-of-year fixed effects to account for district-level growth trends that may be spatially correlated with subway proximity. While these additional fixed effects mask important relevant variation, the point estimates attenuate only slightly.

The city-wide driving restriction policy increases the premium for subway access only within walking distance of subway stations. Figure 3.9 displays the non-parametric relationship between subway proximity and price following the CDR by plotting the mean price change for each half-kilometer bin relative to housing units outside of five kilometers from subway stations. The premium falls to zero after approximately three kilometers, or about 1.8 miles. Intuitively, subway proximity has no impact on housing prices outside of a reasonable walking distance from the station. Figures 3.9 supports our interpretation that demand for the subway network—as opposed to correlated, unobserved amenities—is the mechanism driving the price premium for proximity following the 2008 driving restriction.

We next consider the effect of the driving restriction on demand for living close to one of Beijing's central business districts. Figure 3.10 presents a similar pattern for distance to the central business district as did Figure 3.7; prior to the driving restriction, the price premium for CBD proximity is not significantly different from the driving restriction quarter. However there is a small increase of about 1% in the premium for proximity to the CBD following the driving restriction. Regression results in Table 3.3 show that CBD proximity commands approximately 1.5% higher price per kilometer after the driving restriction. Use of jiedao fixed effects in column (2) significantly increases the magnitude of the policy effect, implying unobserved dis-amenities can attenuate estimates of the CBD-proximity gradient. Estimates of the price gradient change under alternative CBD definitions in Appendix Table C.1 are closer to 1%, but within the range of 1.5%.

Our results show that subway station access came to be in much higher demand relative to CBD access following the driving restriction. Expansion of the subway network makes transit-oriented housing locations more desirable for all marginal subway riders. However, the driving restriction makes housing closer to the CBD more desirable only for marginal drivers with a high value of time. Those wealthy enough to drive their cars are a relatively smaller portion of Beijing's population compared to those reliant on public transit, thus it is not surprising that the subway price effect is much stronger than the CBD price effect. Both results, however, are consistent with the LeRoy and Sonstelie (1983) prediction that changes to transit costs will alter the distance-price gradient.

3.5.2 Comparing the Subway-Distance Premium with Prior Work

Our results on the subway-price gradient are significantly larger than those of Xu et al. (2015), who find that the elasticity of price with respect to subway distance is -0.02%. We compare results directly in Appendix Table C.2 where we employ the same 6-month time period, log-log specification, and start date for the CDR (October 11, 2008) as in Xu et al. (2015). As in Xu et al. (2015), we also exclude housing units located near newly-built subway stations, leaving approximately 19,000 observations. In the first column, we attempt to replicate their main result by including controls for location attributes, such as distance to the city center and dummies for whether the housing complex is located within a "key" school district.[11] Our point estimate

[11] Specifically, we replicate column 8 of Table 2 in Xu et al. (2015)

is within the range of their main result, but is three times larger in magnitude. We attribute the difference in our estimates to having substantially more spatial variation and a more representative data set of the broader Beijing housing market. After introducing location-specific fixed effects, our estimate becomes larger and is highly significant. The jiedao fixed effects are an important improvement upon prior work because they account for unobserved heterogeneity in location amenities. Noticeably, the addition of detailed controls on the housing attributes reduce the estimates, but the effect is still more than double that of Xu et al. (2015).

Additionally, Xu et al. (2015) found that the subway distance premium after the driving restriction ranged between 36-60% of the pre-restriction premium. We estimate the pre-restriction premium for subway accessibility in columns (4) and (8) following Xu et al. (2015) where we limit the sample period to 6 months prior to the driving restriction. Several confounding factors are likely correlated with subway proximity and housing desirability, thus we interpret columns (4) and (8) with caution. For comparison purposes, our results suggest that the post-restriction premium is double the pre-restriction premium, substantially greater than Xu et al. (2015). This underscores that not only was the gradient shift much larger than prior estimates suggest; but the magnitude of this gradient shift is of first order economic significance. This substantial difference in our results implies that subway proximity may be highly correlated with certain disamenities, like noise or congestion, that put downward bias on the housing price-subway distance gradient. Failure to control for these unobserved attributes can significantly underestimate the benefits of transit infrastructure.

3.5.3 Transit Policies and the Demographic Sorting

Given that the price-distance gradients for both subway access and CBD access became steeper following the driving restriction, how did this impact the location choices of different income groups? From LeRoy and Sonstelie (1983), we can predict how demographic groups will locate relative to one another by isolating changes in their commute costs. Inequalities in Eqs. 3.5 and 3.12 provide conditions for lower income groups to live closer to the city center and near subway stations, respectively. Eq. 3.5 is more likely to hold if the transit speed of the rich is much faster than that of the poor (i.e., the time units per mile of t_{car} is much smaller than that of t_{pub}) or if the value that the rich and poor place on their time do not differ substantially.

In the early part of our study period, the difference between t_{car} and t_{pub} is likely to have been large relative to the post period. The city had no restrictions on auto use, and Beijing's subway network was limited to four lines (out of 22 today) leaving areas outside of the Second Ring Road largely bypassed by the subway network. Indeed, Figure 3.4 shows that prior to Beijing's driving restriction, the composition of wealth by distance (either to the CBD or to subway stations) was more evenly distributed relative to the post-policy years.

Beijing's CDR will cause t_{car} and t_{pub} to converge; the subway expansion will reduce t_{pub} while the driving restriction will increase t_{car}. Both mechanisms serve to reduce the likelihood that Eq. 3.5 holds. In other words, both the subway investment and the driving restriction should alter the demographic composition such that higher income people move toward the city center, and lower income people move away from the city center. Similarly, the CDR will increase n_r relative to n_p in Eq. 3.12, making the inequality less likely to hold. The result should be that lower income people move further away from subway stations relative to higher income people.

We formally test predictions of LeRoy and Sonstelie (1983) by using data on household income and location from the mortgage application data. Our analysis spans two years before and two years after the driving restriction (July 20, 2006 through July 20, 2010) to gain statistical power and to allow for a time lag in household's adjustment to the driving restriction policy. We estimate the following in order to understand how individuals choose where to live relative to both subway stations and the central business district:

$$\ln(Km_{izt}) = \alpha(\ln I_{izt} \times R_t) + \delta \ln I_{izt} + \rho R_t + \mathbf{X}_{izt}\theta + \zeta_z + \tau_t + \mu_{izt} \tag{3.17}$$

where i indexes households, z indexes zip codes, and t indexes time. α provides the elasticity of income with respect to distance following the driving restriction policy.

Our results in Tables 3.4 and 3.5 are generally consistent with the model's prediction: wealthier (poorer) households move closer to (further from) both subway stations and central business districts relative to lower (higher) income households after the driving restriction. Zip code level fixed effects allow us to examine how income changes within a neighborhood over time. Variation comes from compositional changes in the type of households selecting to live in a particular neighborhood over time, as well as from some zip codes being "treated" by the expanding subway network.

In columns (4) and (5) of Table 3.4 where we employ the most restrictive controls, the results show that a 10% reduction in household income corresponds to being located 0.80% further - or nearly one kilometer - from a subway station relative to prior to the CDR policy.

Table 3.5 also shows that following the driving restriction, wealthier individuals are more likely to move proximate to the central business district. Again considering the most restrictive specifications in columns (4) and (5), the results show that a 10% reduction in household income corresponds to being located 0.17% further from the nearest business district following the driving restriction - or approximately one third a kilometer. Conditioning on variation within a particular subway line in columns (4) and (5) reduces the estimate magnitude substantially, and the coefficient is no longer statistically significant. These results are generally consistent with alternate definitions of the Beijing CBD, as shown in Appendix Table C.3.

While these results on demographic shifts are modest, they suggest that city-wide policies aimed at reducing traffic and air pollution can be potentially regressive because they not only increase the premium for center-city locations, but they also increase competition for housing near public transit, the mode choice disproportionately utilized by lower income groups. In absence of Beijing's aggressive subway investments, the housing market and demographic sorting responses would likely be much stronger.

The potential regressivity of a city-wide driving restriction policy will depend upon the level of enforcement and the potential for behavioral adjustments. For example, if the purchase of new cars in Beijing were unregulated, wealthy car drivers could have circumvented the license-plate based road rationing by purchasing second car, as in the case of Mexico City (Davis 2008). Such offsetting behavior would likely mitigate the driving restriction's effect on the housing market and demographic sorting. Beijing's strong compliance rate means that car owners can only adjust through a combination of using public transit and relocating within the city to reduce their total commute time.

3.6 Robustness Checks

The government of Beijing has invested heavily in expanding its subway network since 2000. As of 2000, Beijing had two subway lines with 31 stations while today the city has 21 lines with over 370 stations. Appendix Figure C.1 shows how the subway network has expanded substantially over the last two decades, particularly since 2010. The placement of new lines and stations is unlikely to be random. To the extent that increased subway proximity over time is correlated with other location attributes that affect housing prices, such as expectations on commercial development, our results on the subway proximity-price premium may be spurious or driven by reverse causality.

To address these concerns, we first exclude housing units that experienced subway station development nearby. For this sensitivity check, our sample includes only housing units that maintained the same distance to their nearest subway station from 2005 through 2016. Approximately 1,200 building complexes in our sample met this criteria, leaving about 50% of all transaction observations. Relying purely on time-wise variation, rather than spatial variation in subway expansion, results of the bid-rent gradient shift for subway proximity are consistent with our main results. The bid-rent gradient in Appendix Figure C.3 demonstrates a substantial tilt following the driving restriction. Table C.4 shows estimates of the subway proximity-price gradient for this restricted sample. The gradient change is attenuated but within a standard deviation of our main results in Table 3.2. For housing units with no change in their proximity to the nearest subway station, the driving restriction increased the premium for these units by an additional 1.5% per kilometer compared to 3.2% estimated from our full sample based on column (5) of both tables. We interpret these results as a lower bound on the increased demand for subway proximity. The subway stations proximate to this housing subsample are some of the oldest lines and stations in the system network (Beijing's oldest lines were built in 1969 and 1971, without substantial additions until the mid-2000's) thus proximity to these areas may be less desirable to the extent that these lines offer less network advantages and may run less efficiently compared to the newer lines.

As an alternative strategy, we test whether areas that received future subway development after our study period experienced differential price trends over time. In Figure 3.11, we estimate the effect of subway proximity among a sample of housing units that were outside of walking distance from a subway station up through 2013, but came to be within 3 kilometers after new stations were built after 2013. This sample is a subset of housing units located in areas that received development in the future, but should not be affected by the driving restriction during our study period because they are not within reasonable walking distance of subway stations. Any price effects from the driving restriction would raise concern that correlated shocks stemming from unobserved economic investment or growth caused the subway proximity-price gradient shift, as opposed to increased demand for the subway itself. Figure 3.11 is suggestive that the driving restriction did not significantly increase the price premium for subway proximity among this group of housing units. The point estimates are imprecisely estimated due to smaller sample sizes in the post-CDR period. However, the quarterly estimates do not show a clear upward trend, as in Figure 3.7.

3.7 Conclusion

Road rationing policies are an increasingly common policy instrument used among major cities to reduce traffic congestion and air pollution. While prior work has investigated the effectiveness of these policies at improving air quality and congestion, less is known about the ramifications of these policies on the residential location decisions of those who drive relative to those that do not.

Urban land use theory provides clear predictions on how such policies will impact the housing market and the sorting of demographic groups relative to one another. This paper uses detailed, micro-level data on home purchases and buyer demographics in combination with a city-wide driving restriction to test these theories in the context of Beijing, China. Our analysis relies on fine-scale spatial fixed effects to control for location-specific correlates of public transit. We find that the driving restriction increased demand for subway proximity by twice that found by prior literature. We interpret this difference as evidence of the non-random nature of transit infrastructure siting decisions. We additionally utilize novel micro data on household income and housing locations to explore how the driving restriction impacted the residential location choices of income groups relative to one another. Following the policy-induced shocks to housing prices, we find the composition of individuals living proximate to subway stations as well as proximate to Beijing's CBDs shifted toward wealthier households. Each of these results is consistent with the AMM model and its extension by LeRoy and Sonstelie (1983).

These results provide some suggestive evidence that city-wide road rationing policies can have the unintended consequence of limiting access to public transit for lower income individuals. Such effects are likely to be stronger in markets where car ownership is cost-prohibitive to the poor and when enforcement of the driving restriction is strict. Future work should explore the welfare implications of such policies by testing how actual commute times change among the wealthy relative to the poor following a driving restriction. Exposure to pollution and congestion are additionally important pieces to consider. Such analysis can inform whether rent stabilization, or a welfare transfer process may be necessary to offset impacts of driving restrictions on housing affordability.

Figure 3.1: Global Growth in Urban Road Rationing Policies

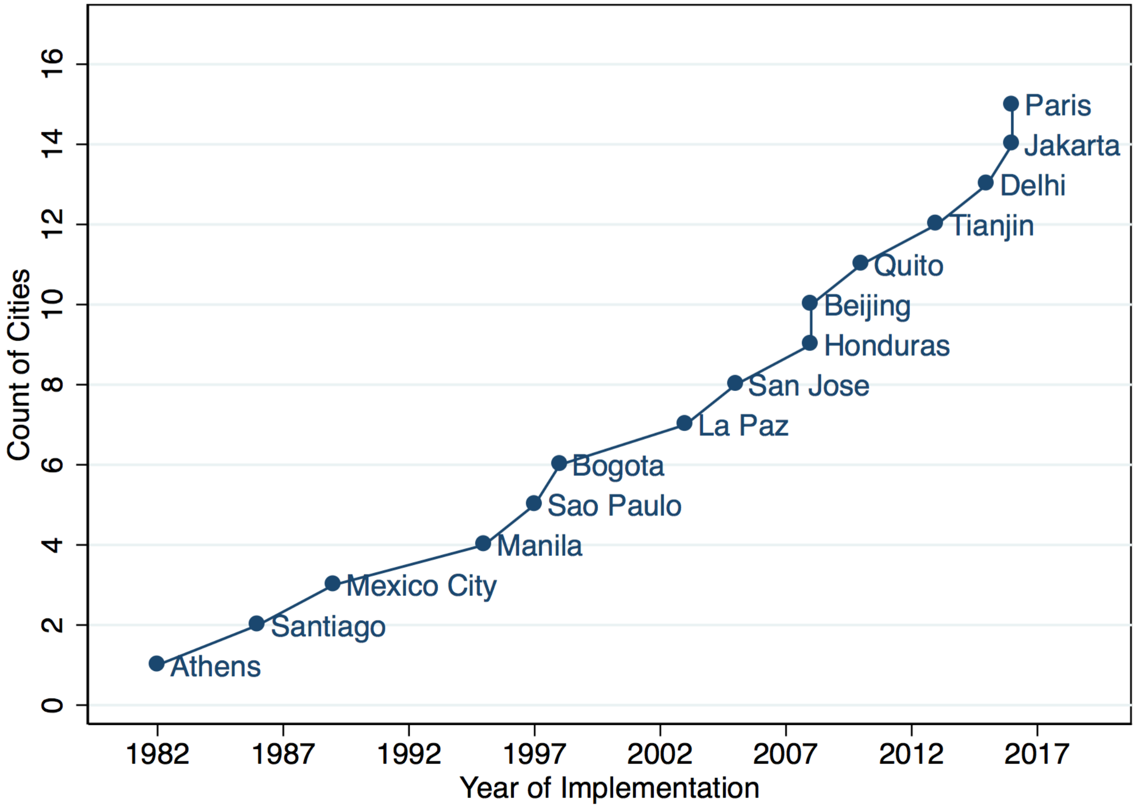

Sources: `http://www.chinadaily.com.cn/china/2013-12/16/content_17175846.htm`; `https://www.bbc.com/news/world-europe-38236926`; `https://en.wikipedia.org/wiki/Road_space_rationing`

Figure 3.2: Housing Units & Subway Stations in Beijing

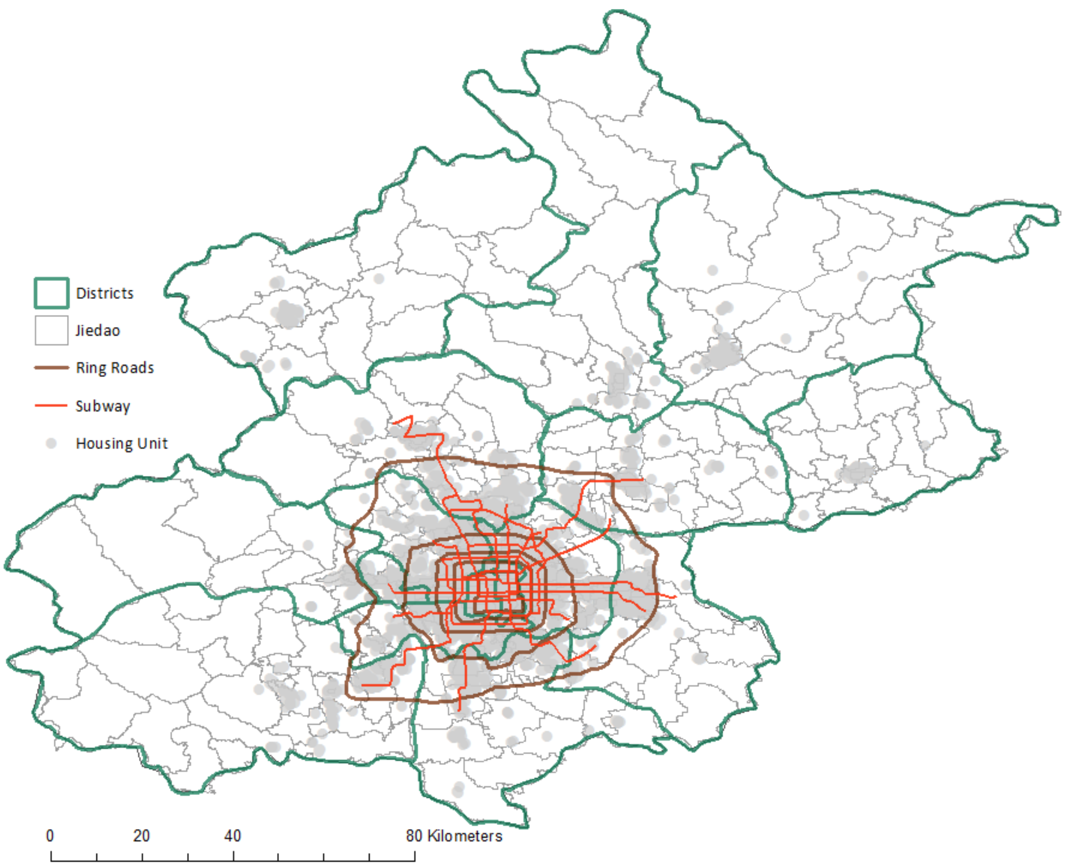

Sources: Beijing Real estate data.

Figure 3.3: Business Districts of Beijing

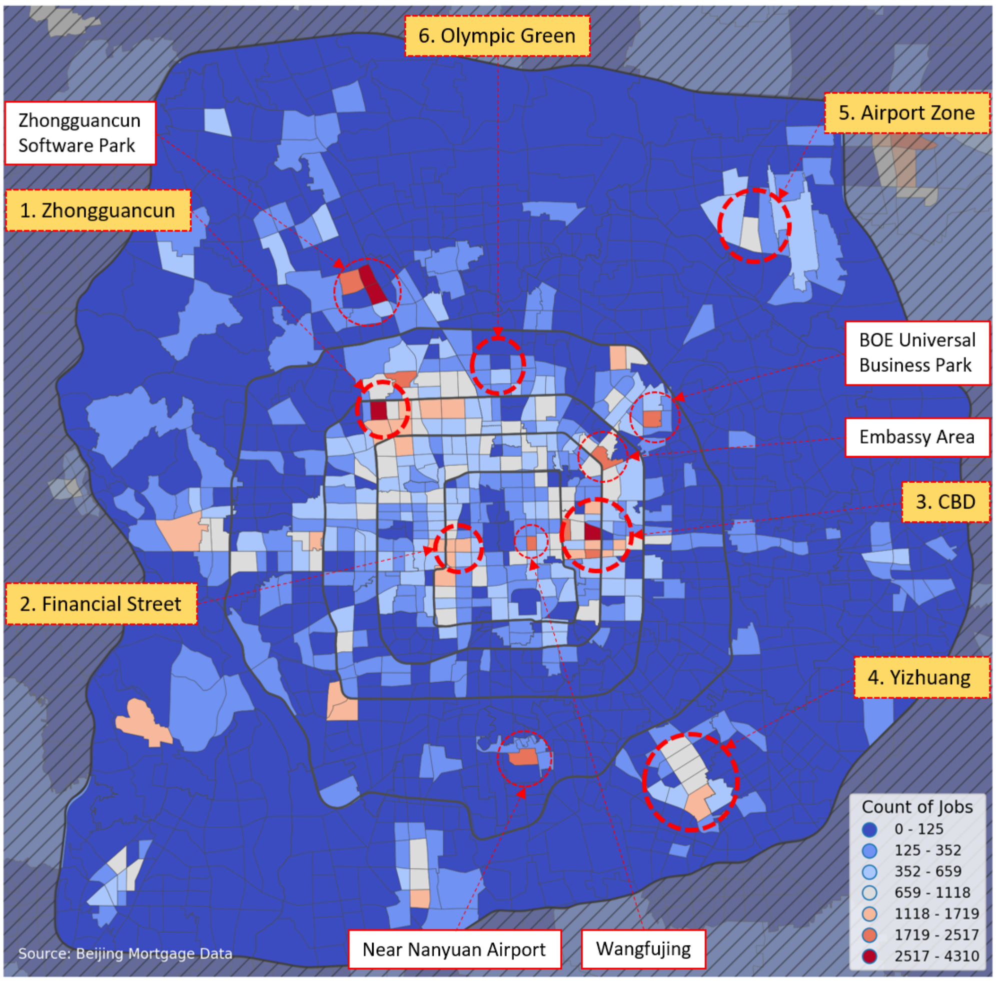

Sources: Ziye Zhang, HPF Mortgage Data. Figure shows employment concentration by transportation analysis zone. Employment location data sourced from HPF Mortgage Data.

Figure 3.4: Income Proximity Gradient

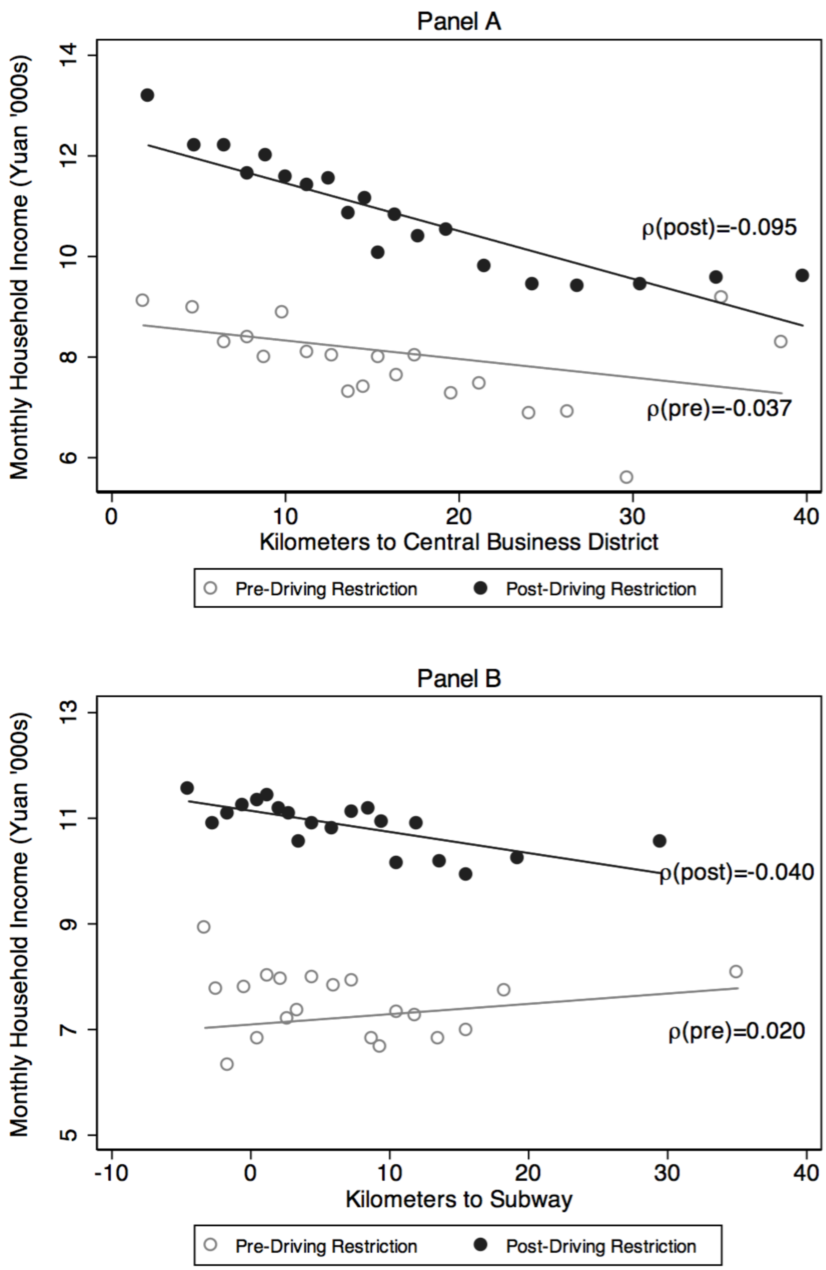

Note: Figure plots mean household income for each of 20 distance bins. Each dot represents 850 and 3670 obs per bin in pre and post, respectively. Panel A means are residualized by distance to the nearest subway. Panel B means are residualized by distance to central business district. "Central Business District" defined as the closest of 7 main business districts. ρ(pre) and ρ(post) are regression coefficients. Includes years 2005-2014. Source: Housing Provident Fund mortgage application dataset.

Figure 3.5: Price Proximity Gradient

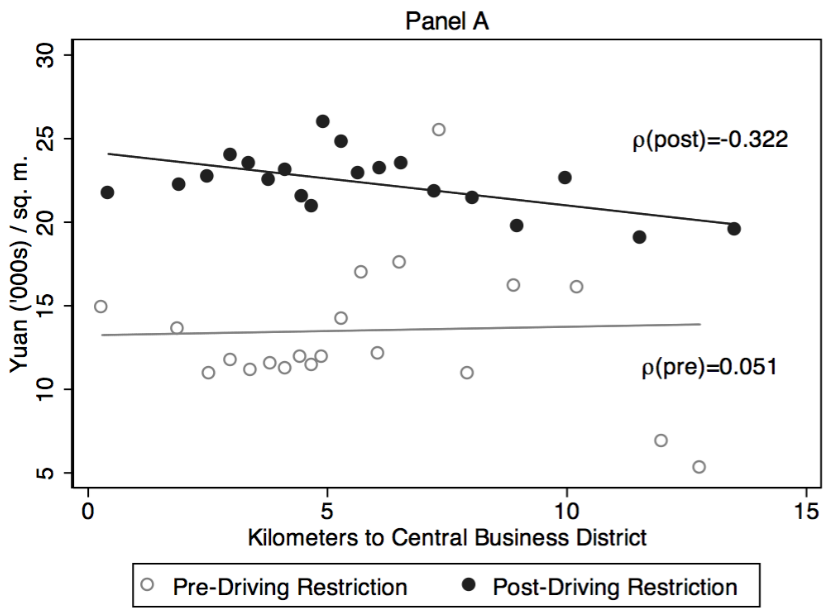

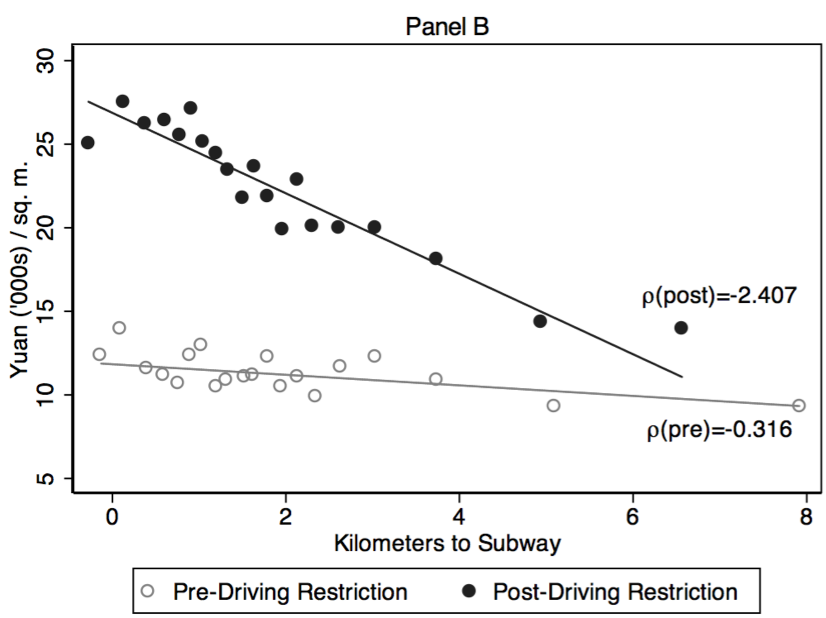

Note: Figures plot mean house price for each of 20 distance bins. Each dot represents 6,300 and 12,900 obs per bin in pre and post, respectively. Panel A means are residualized by distance to the nearest subway. Panel B means are residualized by distance to the central business district. "Central Business District" defined as the closest of 7 main business districts. ρ(pre) and ρ(post) are regression coefficients. Includes years 2005-2014. Source: Real estate transaction dataset.

Figure 3.6: Neighborhood Variation

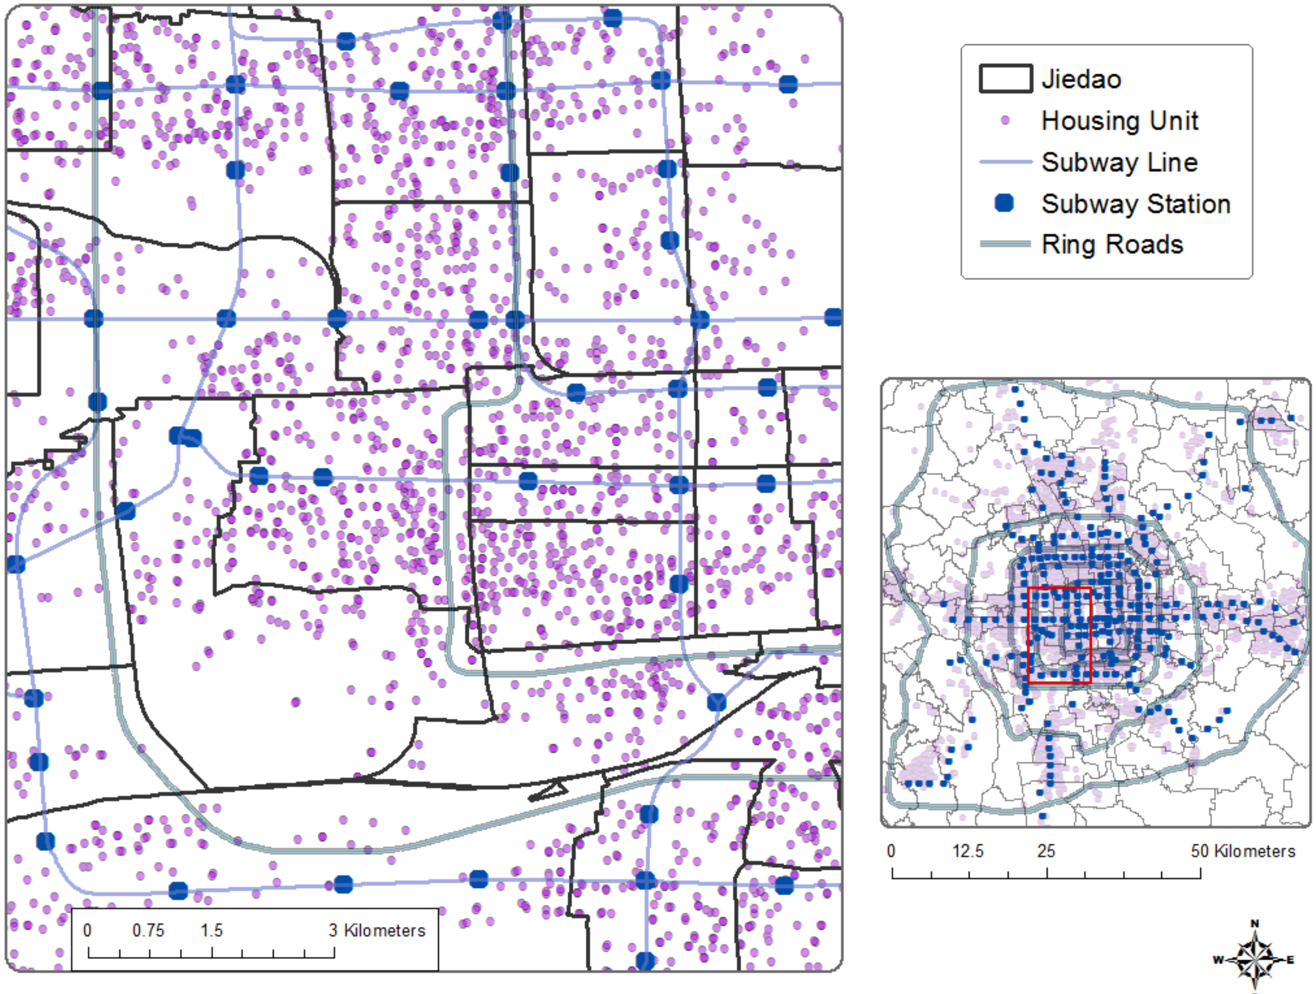

Sources: Beijing Real estate data; Housing Provident Fund data. Figure shows a southwest section of central Beijing, between the second and third ring roads. CDR effects are identified based off of variation in distance to subway stations, or distance to the nearest central business districts across housing units *within* a jiedao.

Figure 3.7: Event Study of the Subway-Housing Price Premium

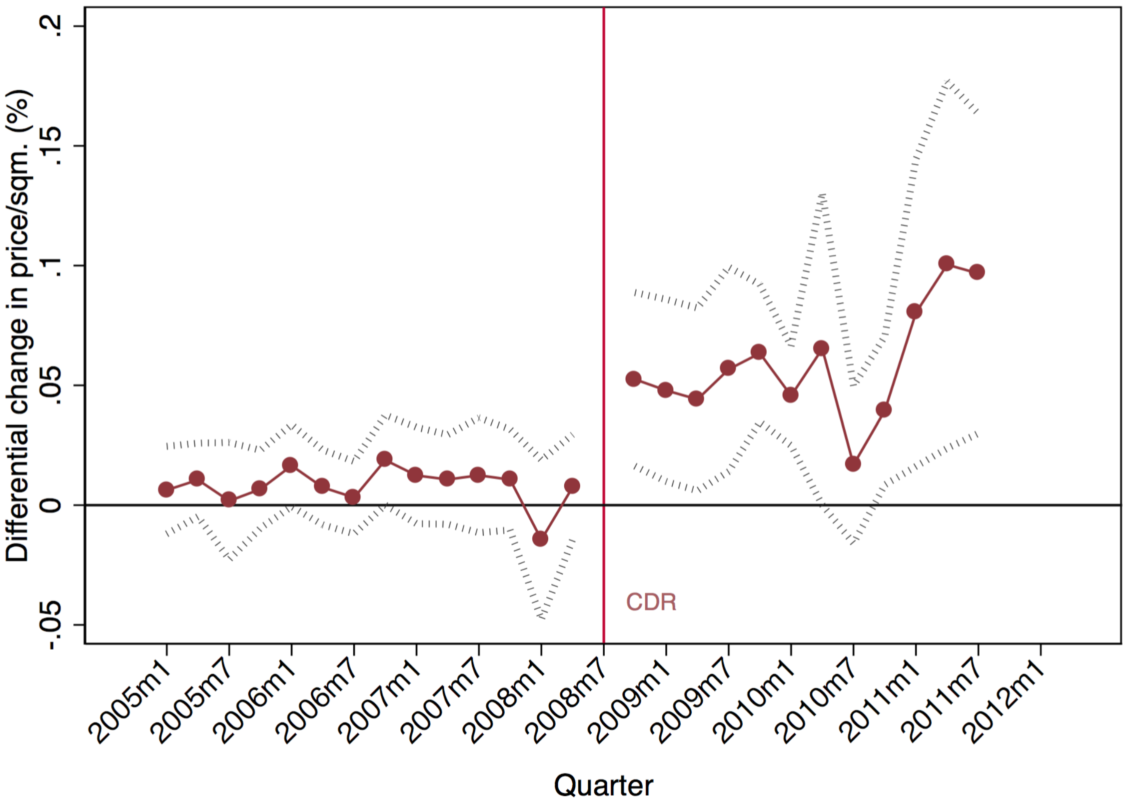

Note: Figure shows the partial effect of subway distance on housing price ln(total price/sqm in ¥2007) at each quarter between Jan 2005 and Dec 2011. The omitted quarter is July-October 2008. Sample includes 256,149 transactions. Controls include district, unit type (resale or newsale), and jiedao fixed effects; complex controls include age, age^2, size, floor-area ratio, green space, no. total floors. Unit controls include unit size, decoration level, whether at top floor and facing direction. Standard errors clustered at jiedao level.

Figure 3.8: Price premium for housing units within vs outside a 3km radius from subway stations

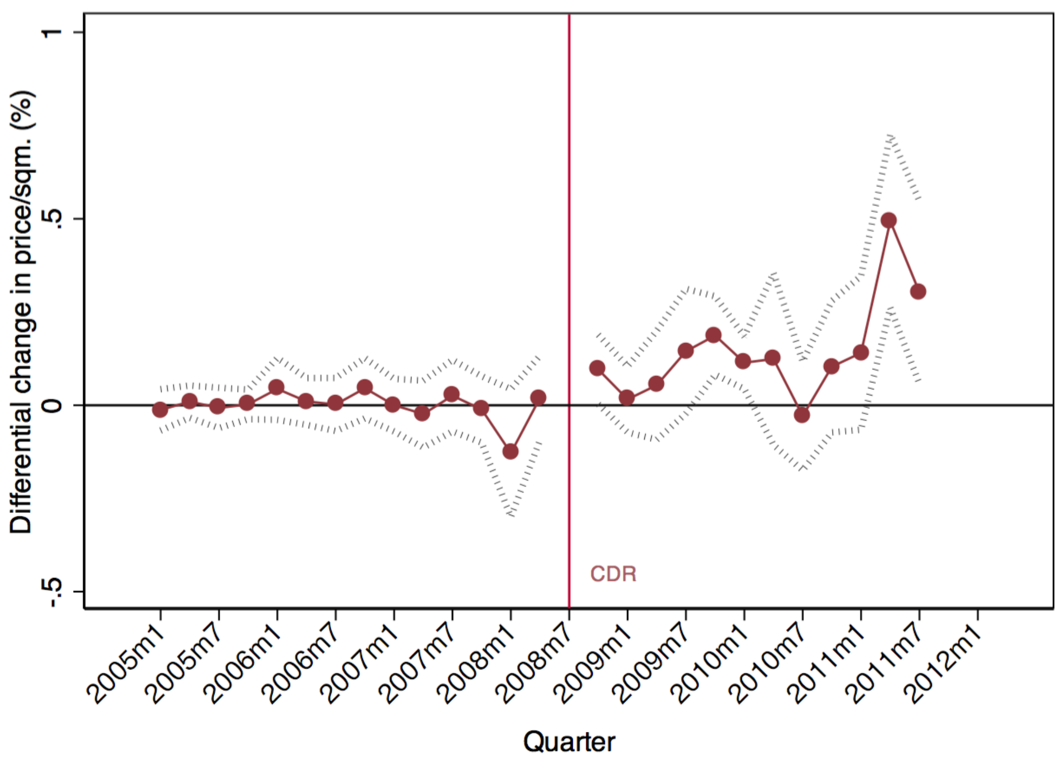

Note: Figure plots estimates of the average change in price per square meter for housing units within vs over 3km of a subway station by quarter. The reference quarter is July-October of 2008. Dashed lines represent 95% confidence intervals. Controls include district, unit type (resale or newsale), and jiedao fixed effects; complex controls include age, age^2, size, floor-area ratio, green space, no. total floors. Unit controls include unit size, decoration level, whether at top floor, and facing direction. Standard errors clustered at jiedao level.

Figure 3.9: Average Effect of Driving Restriction on Housing Prices by Subway Distance Bin

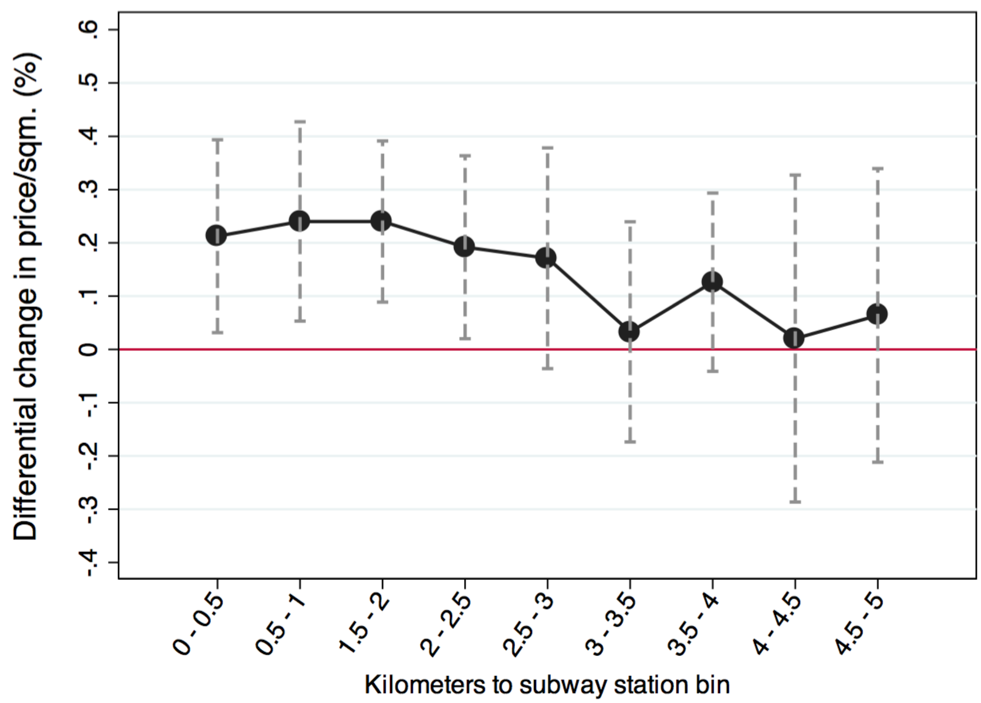

Note: Figure plots estimates of the average change in housing price per square meter following the CDR by half-mile distance bins to subway stations. The reference bin includes housing units over 5 kilometers from subway stations. Dashed lines represent 95% confidence intervals. Includes transactions from July 2006 through July 2010. Controls include district, unit type (resale or newsale), and jiedao fixed effects; complex controls include age, age^2, size, floor-area ratio, green space, no. total floors. Unit controls include unit size, decoration level, whether at top floor, and facing direction. Standard errors clustered at jiedao level.

Figure 3.10: Event Study of the CBD-Housing Price Premium

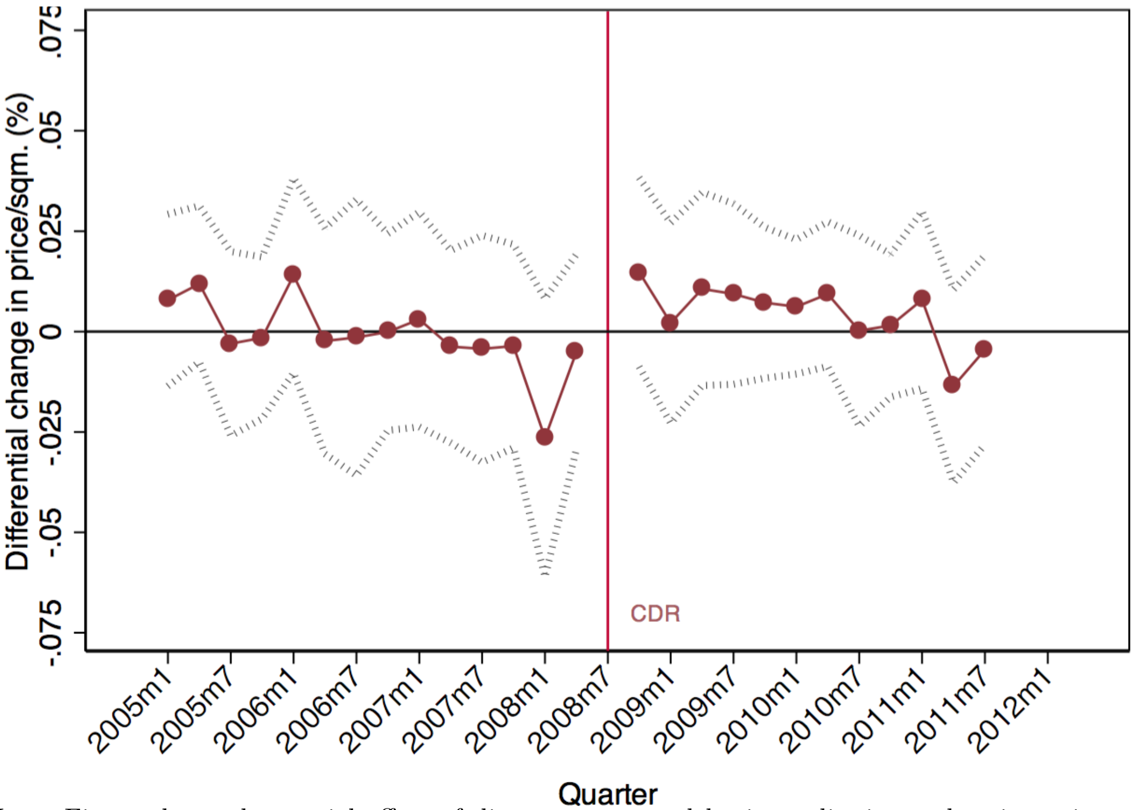

Note: Figure shows the partial effect of distance to central business district on housing price ln(total price/sqm in ¥2007) at each quarter between Jan 2005 and Dec 2011. The omitted quarter is July-October 2008. Sample includes 256,149 transactions. Controls include distance to nearest subway station, district, unit type (resale or newsale), and jiedao fixed effects; complex controls include age, age^2, size, floor-area ratio, green space, no. total floors. Unit controls include unit size, decoration level, whether at top floor, and facing direction. Standard errors clustered at jiedao level. "Central Business District" defined as the closest of 7 main business districts.

Figure 3.11: Placebo Test of Driving Restriction

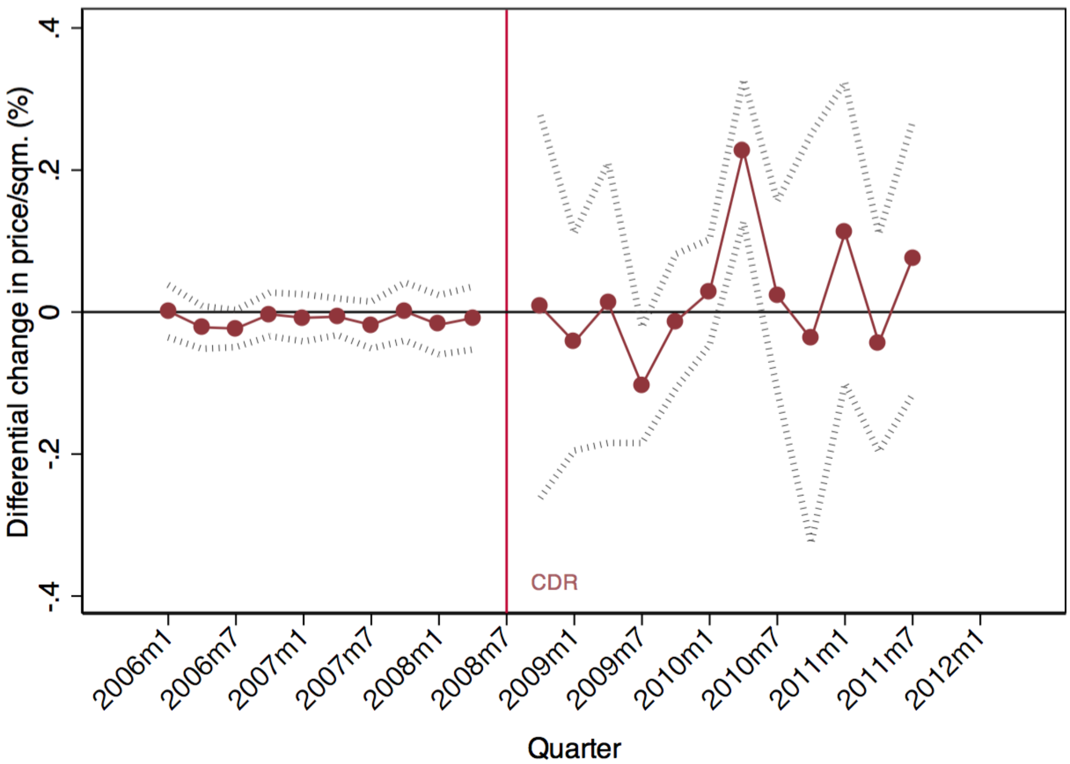

Note: Figure shows the partial effect of subway distance on housing price ln(total price/sqm in ¥2007) at each quarter between Jan 2006 and Dec 2011. The omitted quarter is July-October 2008. Sample includes housing units located in building complexes that are over 3km from the nearest subway station through the event period; but are under 3km from a station after the event period ends (beginning in 2013). The sample includes 65,758 transactions. Controls include district, unit type (resale or newsale), and jiedao fixed effects; complex controls include age, age^2, size, floor-area ratio, green space, no. total floors. Unit controls include unit size, decoration level, whether at top floor, and facing direction. Standard errors clustered at jiedao level.

Table 3.1: Descriptive Statistics

	Real Estate Transactions		Mortgage Data	
	Mean	(St. Dev)	Mean	(St. Dev)
Total purchase price ('07 Yuan)	1,195,255.8	866,141.1	543,993.8	269,514.8
Price per sq.m ('07 Yuan)	10,586.6	4,182.1	5,447.2	2,467.3
Unitsize (sq.m.)	109.3	46.4	103.0	32.1
Km to subway	3.2	2.6	12.9	13.5
Km to nearest CBD	5.7	3.5	18.0	13.7
Km to City Center	9.3	3.4	24.2	15.9
Building Age	10.7	6.8	5.7	7.0
Building Floor-to-Area Ratio	2.9	1.2	2.2	1.0
Building Green space ratio	0.3	0.1	33.4	6.8
Household Monthly Income ('07 Yuan)			7,193	5,002
Age of household head			35.8	7.3
Years of work experience of household head			6.7	8.7
Education level of household head (mode)			Bachlor's	
No. Complexes	4,114		3,971	
No. Neighborhoods	185		183	
Observations	237,140		46,471	
Observations by Year				
2005	1,928		4,858	
2006	62,497		5,273	
2007	48,637		4,410	
2008	25,003		4,225	
2009	66,783		12,513	
2010	34,927		8,321	
2011	16,374		6,871	

Note: The unit of observation for the Real Estate Transaction data is a housing purchase transaction. The unit of observation for the Mortgage Data is a mortgage loan application, or a household, equivalently. The means are calculated using pre-policy years 2005 through July 2008.

Table 3.2: Effects of CDR policy on Subway-Price gradient

	(1)	(2)	(3)	(4)	(5)
Subway Proximity (km) x CDR	0.040*	0.040***	0.036***	0.034***	0.032***
	(0.023)	(0.014)	(0.011)	(0.011)	(0.012)
Subway Proximimty (km)	0.049***	0.005	0.017*	0.015*	0.015
	(0.015)	(0.014)	(0.009)	(0.009)	(0.011)
Avg Price Premium / Km	$405.34	$403.95	$364.65	$348.76	$328.23
Jiedao FE		Y	Y	Y	Y
Controls			Y	Y	Y
Year-Month FE				Y	
DistrictxYear-Quarter FE					Y
Observations	82002	82002	82002	82002	82002
Adjusted R^2	0.125	0.495	0.611	0.621	0.623

Note: Dependent variable is ln(total price per square meter in 2007 real Yuan). Standard errors clustered at jiedao level. Sample spans 12 mos. before and after CDR. Average price premium evaluated at the mean unit size (100 sqm for owner-occupied) within 1 and 1.5 km of a subway station. Controls include year, month, and district fixed effects; controls for complex age, age2, size, floor-area ratio, green space, no. total floors; controls for housing unit size, decoration level, floor level, facing direction, and no. bedrooms. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.3: Effects of CDR policy on Central Business District-Price gradient

	(1)	(2)	(3)	(4)	(5)
CBD Proximity (km) x CDR	0.006	0.016**	0.013**	0.012**	0.016*
	(0.011)	(0.008)	(0.005)	(0.005)	(0.009)
CBD Proximity (km)	0.035***	0.003	-0.010	-0.008	-0.011
	(0.010)	(0.018)	(0.017)	(0.016)	(0.016)
Avg Price Premium / Km	$55.74	$163.82	$126.64	$119.76	$157.00
Jiedao FE		Y	Y	Y	Y
Controls			Y	Y	Y
Year-Month FE				Y	
DistrictxYear-Quarter FE					Y
Observations	82002	82002	82002	82002	82002
Adjusted R^2	0.148	0.488	0.611	0.621	0.624

Note: Dependent variable is ln(total price per square meter in 2007 real Yuan). Standard errors clustered at jiedao level. Sample spans 12 mos. before and after CDR. Average price premium evaluated at the mean unit size (100 sqm for owner-occupied) within 4 and 6 km of a the nearest central business district. Controls include distance to nearest subway station, year, month, and district fixed effects; controls for complex age, age2, size, floor-area ratio, green space, no. total floors; controls for housing unit size, decoration level, floor level, facing direction, and no. bedrooms. "Central Business District" defined as the closest of 7 main business districts. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.4: CDR Policy and Income sorting near Subways

	(1)	(2)	(3)	(4)	(5)
Ln(Household Income)× CDR	-0.046	-0.040	-0.015	-0.080***	-0.081***
	(0.077)	(0.057)	(0.050)	(0.024)	(0.025)
Ln(Household Income)	-0.139*	-0.017	-0.002	0.051**	0.048**
	(0.073)	(0.046)	(0.036)	(0.023)	(0.022)
Year FE	Y	Y	Y	Y	Y
District FE		Y			
Zip FE			Y	Y	Y
Subway Line FE				Y	Y
Controls					Y
Observations	18135	18135	18135	18135	18135
Adjusted R^2	0.636	0.811	0.911	0.922	0.923

Note: Dependent variable is ln(Distance to Subway (km)). Income is household monthly income ('000 yuan). CDR equals 1 after July 20 2008. Standard errors clustered by zip code. Sample spans July 20, 2006-July 20, 2010. Controls for distance to nearest CBD. Demographic controls include husband and wife age, employment rank, education, employer type, and tenure.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.5: CDR Policy and Income sorting near the Central Business District

	(1)	(2)	(3)	(4)	(5)
Ln(Household Income)$\times$ CDR	-0.031	-0.053*	-0.051***	-0.017	-0.017
	(0.081)	(0.030)	(0.016)	(0.015)	(0.015)
Ln(Household Income)	-0.752***	-0.041	0.008	-0.012	-0.008
	(0.107)	(0.026)	(0.014)	(0.012)	(0.012)
Year FE	Y	Y	Y	Y	Y
District FE		Y			
Zip FE			Y	Y	Y
Subway Line FE				Y	Y
Controls					Y
Observations	18135	18135	18135	18135	18135
Adjusted R^2	0.196	0.841	0.944	0.956	0.956

Note: Dependent variable is ln(Distance to CBD (km). Income is household monthly income ('000 yuan). CDR equals 1 after July 20 2008. Standard errors clustered by zip code. Sample spans July 20, 2006-July 20, 2010. Controls for distance to subway. Demographic controls include husband and wife age, employment rank, education, employer type, and tenure.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

APPENDIX A

APPENDIX FOR CHAPTER 1

A.1 Figures & Tables

Figure A.1: Population Distribution of Local Government Sample

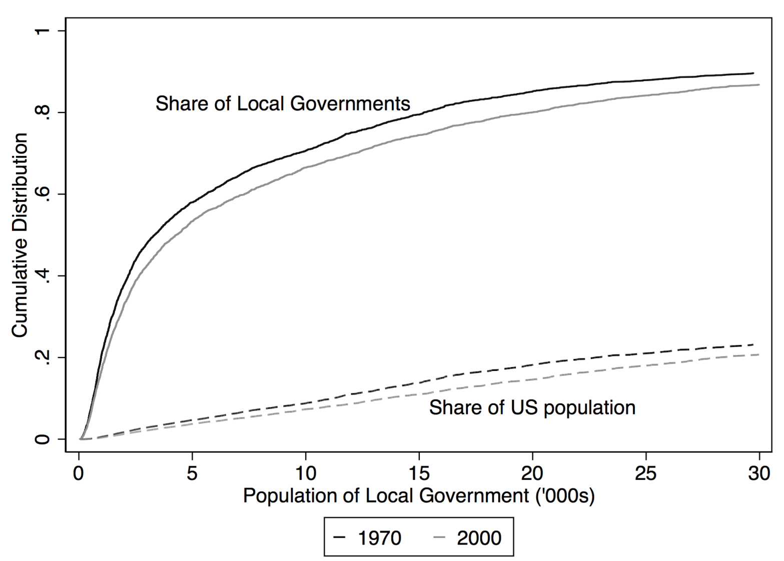

Source: US Census of Governments. This plot shows the cumulative distribution function of the sample local government population size, and the corresponding share of the US population as of 1970 and 2012. Includes 3,334 governments categorized as "municipalities" and "townships" by the Census of Governments, which operate a wastewater treatment plant. See Section 1.4 for further details on sample selection.

Figure A.2: Wastewater Expenditures per capita by Size of Downstream Population Using State Compliant Plant Share Variation

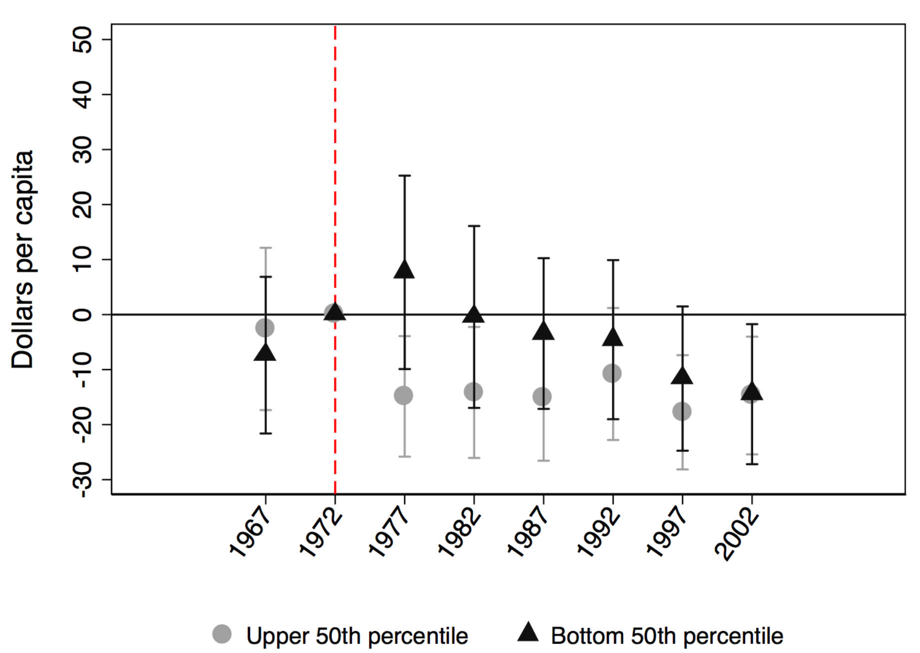

Source: USGS, Census of Governments, author's calculations. Figure plots $\delta_{t\times\underline{50}} + \delta_t$ and δ_t from equation: $y_{it} = \sum_t \delta_{t\times\underline{50}}(I_{\underline{50}} \times S_s \times D_t) + \sum_t \delta_t(S_s \times D_t) + (D_i \times D_t) + \gamma_t + \nu_i + \varepsilon_{it}$ where the dependent variable is wastewater expenditures per capita for city i in year t; $I_{\underline{50}}$ is an indicator for a city having downstream population in the bottom 50th percentile, S_s is state share of compliant plants as of 1972, D_t is an indicator for year t, and D_i is downstream population. Bands show 95% confidence intervals. The reference year is t=1972 and the reference city is one in a low compliant state. Black triangles show the estimated difference in expenditures in year t relative to 1972, relative to high downstream-high compliant state cities, and all low compliant state cities, for cities with downstream population size in the bottom 50th percentile. Gray estimates show the difference in expenditures in year t relative to 1972 for cities with downstream population size in the top 50th percentile, relative to low downstream-high compliant state cities, and all low compliant state cities. Robust standard errors are clustered at the city level. The gray bars show that cities in the top 50th percentile of downstream population within high compliant states have significantly less wastewater expenditures after 1972 relative to all other cities. This is due to the fact that downstream population is *less* predictive of *ex ante* compliance in low compliant states. In other words, all cities in low compliant states, regardless of their downstream population size, were more likely to be noncompliant. Consequently, their expenditures are not significantly different from low downstream population cities in high compliant states.

Figure A.3: Plant Service Population vs Local Government Population

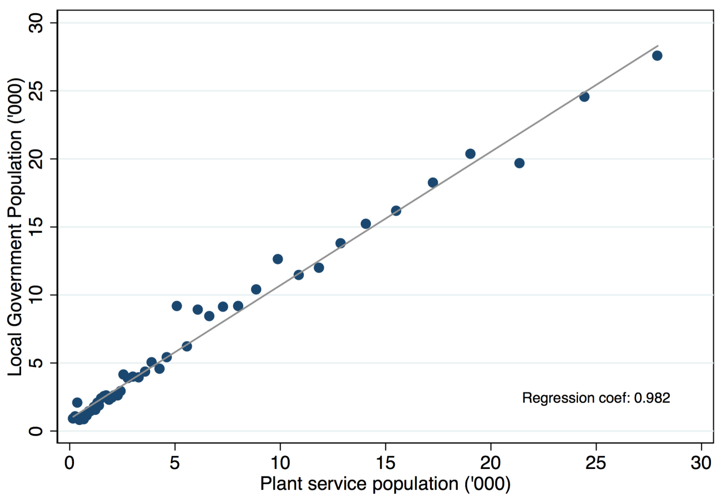

Source: US Census of Governments; EPA CWNS. Plot shows relationship between local government population as reported by Census of Governments, and service population of its corresponding plant as reported by EPA. Includes 3,226 governments with population less than 30,000. Regression coefficient (r) estimated from Census population =r Plant population + e. Each dot represents approximately 60 cities.

Table A.1: Specification Sensitivity of CWA NonCompliance and Wastewater Expenditures per capita

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Difference-in-Differences						
Primary'72xPost	12.713**	64.239***	65.585***	64.948***	64.053***	59.114***	58.715***
	(6.329)	(10.010)	(9.730)	(9.808)	(9.917)	(9.406)	(9.417)
	Instrumented Difference-in-Differences						
Primary'72xPost	47.825**	101.599***	140.820***	143.433***	129.552**	155.228***	151.768***
	(23.241)	(35.054)	(52.438)	(49.797)	(50.268)	(50.067)	(50.576)
YearFE		Y	Y	Y	Y	Y	Y
CityFE		Y	Y	Y	Y	Y	Y
RegionTrend			Y	Y	Y	Y	Y
RiverPopulation x YearFE				Y	Y	Y	Y
Industry Composition x YearFE					Y	Y	Y
City Controls x YearFE						Y	Y
County Income Trend							Y
Observations	14866	14866	14866	14866	14866	14866	14866

Note: Dependent variable is wastewater expenditures per capita in 2012 dollars. Difference-in-differences reports estimates of β from Equation 1.2. IV reports estimates of β_{IV} from Eq. 1.3. Industry composition includes mean from 1967 to 1972 of county share of employment in manufacturing and water-polluting manufacturing. City controls include means from 1967 to 1972 of intergovernmental grants and distance from coastline. Region trend includes 8 US census regions, based on the Bureau of Economic Analysis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.2: Specification Sensitivity of CWA NonCompliance and Local Government Growth

	Difference-in-Differences											
	Dissolved Oxygen			Ln(Population)			Ln(Median house price)			High Skill Share		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
Primary'72xPost	-0.077	-0.030	0.039	-0.264***	-0.024*	-0.030***	-0.152***	-0.057***	-0.006	-0.022***	-0.009***	-0.005**
	(0.058)	(0.050)	(0.047)	(0.067)	(0.012)	(0.011)	(0.020)	(0.010)	(0.008)	(0.005)	(0.002)	(0.002)
	Instrumented Difference-in-Differences											
	Dissolved Oxygen			Ln(Population)			Ln(Median house price)			High Skill Share		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
Primary'72xPost	-0.918***	1.357***	1.543***	-1.682***	0.181***	0.121*	-0.294***	-0.227***	0.013	0.034	0.051***	0.027**
	(0.330)	(0.286)	(0.327)	(0.395)	(0.070)	(0.071)	(0.101)	(0.057)	(0.054)	(0.026)	(0.014)	(0.013)
Observations	14177	14177	14177	8264	8264	8264	8264	8264	8264	6366	6366	6366
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
City FE		Y	Y		Y	Y		Y	Y		Y	Y
All Controls			Y			Y			Y			Y

Note: Difference-in-differences reports estimates of β from Eq. 1.2. IV reports estimates of β_{IV} from Eq. 1.3. Includes decade interval years only (1972, 1982, 1992). High Skill Share is the share of city population with 4 or more years of college (1972) or bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment. Standard errors clustered by city. Controls includes all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.3: Capitalization of CWA Infrastructure Mandate

	Wastewater User Fee	
Per Capita Annual Fee Increase, τ	\$62	
Interest Rate, (r)	10%	16%
Net present value over 18 years of τ, $NPV_{(r,18)}$	\$1,014	\$641
Mean increase in housing value, H:	\$1,069	\$1,069
Value of Mandate, $(1 + H/NPV)$	2.05	2.67

Note: Value of Mandate provides the implied value of the CWA mandated infrastructure, interpreted as a multiplier on the NPV of user fees. Estimates of per capita annual fee increase (δ_{POST}) sourced from Panel B of Table 1.4. 10% and 16% rate of return based on average return from 1975-1992 on 10-year treasury bond and S&P 500, respectively. Net present value calculation is: $NPV_{(r,18)} = \sum_{t=0}^{18} R_t/(1+r)^t$ where R_t is total user fee per household, calculated as 2.5 people per household $\times$ the per capita tax increase. 18 years covers 1975 (the first year of significant plant infrastructure construction, as well as EPA grant distribution (Copeland 2015)) through 1992. Mean loss in housing value calculated as the average city median housing value as of 1970 $\times$ the housing price effect in Panel B of Table 1.5.

Table A.4: Effect of CWA Compliance on Local Government Budgets Using Count of Big Cities Downstream as IV

PANEL A: Expenditures Per Capita

	Total	Wastewater			Other					
		Total	Capital	Operating	Total	PublicSafety	PublicWorks	GenAdmin	Welfare	Rec
Primary'72 x Post	486.405	208.348***	154.527***	48.357***	146.727	79.770**	29.046	6.915	15.468	15.528
	(347.818)	(61.432)	(57.122)	(16.102)	(317.553)	(34.074)	(311.407)	(20.883)	(23.263)	(15.160)
Baseline mean	1032.09	66.67	38.65	27.69	650.54	132.39	384.35	61.36	29.43	43.02

PANEL B: Revenues Per Capita

	TotalRevenues			UserFees		Taxes			Debt	
	Total	InterGovt	Own	Total	Wastewater	Total	Property	Sales&License	LongTerm	ShortTerm
Primary'72 x Post	758.405***	480.875***	277.303*	58.865	61.988***	-3.027	13.417	-16.536	-103.882	205.286***
	(186.078)	(93.602)	(150.281)	(65.654)	(20.851)	(55.876)	(46.782)	(29.244)	(760.276)	(76.353)
Baseline mean	1019.02	170.60	848.45	111.91	31.41	389.85	300.87	88.99	1394.71	87.83
First Stage F-statistic	19.49	19.49	19.49	19.49	19.49	19.49	19.49	19.49	19.49	19.49
Clusters	2975	2975	2975	2975	2975	2975	2975	2975	2975	2975
Observations	14866	14866	14866	14866	14866	14866	14866	14866	14866	14866

Note: Dependent variables are in 2012 dollars per capita. Table reports estimates of β from Eq. 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$, where the instrument for $(P_i \times POST_t)$ is Eq. 1.4 and D_i is count of cities downstream with population greater than 30,000. Standard errors clustered by city. $\mathbf{X}$ includes all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.5: Effect of CWA Noncompliance on Water Quality and Municipal Growth Using Count of Big Cities Downstream as IV

	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary'72 x Post	1.622***	0.135*	0.073	0.028**
	(0.344)	(0.071)	(0.057)	(0.014)
Observations	14177	8264	8264	6366
First Stage F-statistic	18.92	19.36	19.36	14.00
Baseline mean	8.11	36710	95461	0.11
Marginal effect (%)	20%	14.11%	7.43%	26%

Note: Excluded instrument is count of cities downstream with 1970 population greater than 30,000. Includes decade interval years only (1972, 1982, 1992). "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Table reports estimates of β from Eq. 1.3: $y_{ist} = \beta_{IV}(P_i \times POST_t)(P_i \times \widehat{POST}_t) + \mathbf{X_i}\theta_t + D_i\sigma_t + (R \times t) + \tau_t + \kappa_i + \epsilon_{ist}$ Standard errors clustered by city. Includes all controls listed in Table 1.2, column(4). $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Table A.6: Effect of CWA Noncompliance on Local Government Budgets, Excluding Coastal Cities

PANEL A: Expenditures Per Capita

	Total	Wastewater			Other					
		Total	Capital	Operating	Total	PublicSafety	PublicWorks	GenAdmin	Welfare	Rec
Primary'72xPost	307.685	98.679**	54.010	43.229***	150.128	36.733**	57.346	9.098	30.364	16.588
	(282.468)	(47.450)	(44.502)	(15.408)	(246.815)	(16.202)	(240.076)	(19.243)	(24.552)	(14.539)
Baseline mean	964.84	66.44	38.36	27.82	630.86	122.92	389.20	57.26	22.80	38.67

PANEL B: Revenues Per Capita

	TotalRevenues			UserFees		Taxes			Debt	
	Total	InterGovt	Own	Total	Wastewater	Total	Property	Sales&License	LongTerm	ShortTerm
Primary'72xPost	509.782***	285.569***	223.963	28.472	49.381***	73.299	7.867	65.290**	-403.944	29.705
	(173.024)	(72.001)	(149.417)	(66.914)	(18.602)	(47.849)	(38.332)	(26.720)	(844.964)	(67.254)
Baseline mean	955.27	152.90	802.42	114.26	31.87	329.56	247.72	81.85	1351.27	76.53
First Stage F-statistic	18.74	18.74	18.74	18.74	18.74	18.74	18.74	18.74	18.74	18.74
Clusters	2553	2553	2553	2553	2553	2553	2553	2553	2553	2553
Observations	12564	12564	12564	12564	12564	12564	12564	12564	12564	12564

Note: Excludes cities within 50 kilometers of an ocean coastline. Dependent variables are in 2012 dollars per capita. Table reports estimates of β from Eq. 1.3. Standard errors clustered by city. Includes all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.7: Effect of CWA Noncompliance on Water Quality and Municipal Growth, Excluding Coastal Cities

	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary'72xPost	1.227***	0.148**	-0.012	0.027**
	(0.321)	(0.074)	(0.056)	(0.012)
First Stage F-statistic	18.08	18.14	18.14	14.41
Baseline mean	8.18(mg/l)	24,703	$91,492	11%
Marginal effect (%)	15%	15.61%	-1.35%	24%
Clusters	2515	2543	2543	1951
Observations	11979	7043	7043	5276

Note: Excludes cities within 50km of an ocean coastline. Includes decade interval years only (1972, 1982, 1992). "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Standard errors clustered by city. Includes all controls listed in Table 1.2, column(4). $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$

Table A.8: Effect of CWA Noncompliance on Local Government Budgets, Excluding Hydrologic Regions with Largest Downstream Population

PANEL A: Expenditures Per Capita

	Total	Wastewater			Other					
		Total	Capital	Operating	Total	PublicSafety	PublicWorks	GenAdmin	Welfare	Rec
Primary'72 x Post	462.934	255.223***	184.799**	68.690***	201.458	109.054***	-37.104	24.321	38.879	66.307**
	(433.562)	(85.569)	(75.785)	(25.637)	(361.814)	(38.851)	(347.725)	(39.213)	(30.447)	(26.854)
Baseline mean	1073.49	67.03	40.59	26.17	669.88	138.12	387.12	65.16	33.13	46.34

PANEL B: Revenues Per Capita

	TotalRevenues			UserFees		Taxes			Debt	
	Total	InterGovt	Own	Total	Wastewater	Total	Property	Sales&License	LongTerm	ShortTerm
Primary'72 x Post	904.032**	530.981***	372.943	36.135	98.215***	286.303	138.638	148.220***	-684.323	-86.256
	(359.807)	(144.491)	(306.569)	(94.560)	(31.533)	(185.367)	(173.392)	(49.224)	(1589.046)	(124.778)
Baseline mean	1066.22	181.68	884.59	107.23	29.36	425.96	335.41	90.57	1428.26	97.71
First Stage F-statistic	9.53	9.53	9.53	9.53	9.53	9.53	9.53	9.53	9.53	9.53
Clusters	2183	2183	2183	2183	2183	2183	2183	2183	2183	2183
Observations	11084	11084	11084	11084	11084	11084	11084	11084	11084	11084

Note: Excludes cities within the upper Mississippi and Ohio river hydrologic regions. Dependent variables are in 2012 dollars per capita. Table reports estimates of β from Equation 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + R_\tau + \nu_i + \tau_t + \varepsilon_{it}$. The instrument for $(P_i \times POST_t)$ is Equation 1.4. Standard errors clustered by city. $\mathbf{X}$ includes all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.9: Effect of CWA NonCompliance on Water Quality and Municipal Growth, Excluding Hydrological Regions with Largest Downstream Population

	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary'72 x Post	1.244***	0.282**	-0.097	0.027
	(0.442)	(0.119)	(0.083)	(0.021)
First Stage F-statistic	9.49	9.71	9.71	8.10
Baseline mean	8.07 mg/l	42,119	$96,036	11%
Marginal effect (%)	15%	31.67%	-9.60%	25%
Clusters	2152	2179	2179	1741
Observations	10545	6120	6120	4815

Note: Include decade interval years only (1972, 1982, 1992). "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Standard errors clustered by city. Includes all controls listed in Table 1.2, column(4). $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Table A.10: CWA Noncompliance and Local Government Budgets using Full Sample

	Expenditures			Revenues						
Panel A: Difference-in-Differences	Total	Wastewater	Other	Total	InterGovt	UserFees	WWUserFees	Taxes	LongTermDebt	ShortTermDebt
Primary'72xPost	-32.969	33.285***	-45.425*	-24.609	20.279**	0.160	-3.923*	-6.555	-155.893	16.527*
	(34.508)	(6.813)	(27.105)	(26.316)	(8.474)	(8.590)	(2.219)	(8.944)	(102.693)	(8.929)
Marginal effect (%)	-3%	58%	-8%	-3%	13%	0%	-17%	-2%	-11%	24%
Panel B: Instrumented Difference-in-Differences	Total	Wastewater	Other	Total	InterGovt	UserFees	WWUserFees	Taxes	LongTermDebt	ShortTermDebt
Primary'72xPost	969.663***	239.874***	757.749***	1282.497***	744.439***	67.010	69.263***	53.144	-525.122	146.464**
	(276.258)	(66.598)	(201.701)	(297.309)	(104.000)	(124.439)	(17.544)	(161.111)	(990.956)	(61.659)
Marginal effect (%)	103%	419%	130%	140%	465%	84%	307%	15%	-38%	217%
First Stage F-statistic	33.11	33.11	33.11	33.11	33.11	33.11	33.11	33.11	33.11	33.11
Baseline mean	944.49	57.20	580.90	915.86	159.94	80.13	22.54	350.15	1383.99	67.57
Clusters	10820	10820	10820	10820	10820	10820	10820	10820	10820	10820
Observations	48348	48348	48348	48348	48348	48348	48348	48348	48348	48348

Note: Includes full sample of wastewater treatment plants, including those that were built after 1972. Dependent variables are in 2012 dollars per capita. Standard errors clustered by city. Table reports estimates of β from Eq. 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$. The instrument for $(P_i \times POST_t)$ in Panel B is Eq. 1.4. Standard errors clustered by city. $\mathbf{X}$ includes all controls listed in Table 1.2, column(4). $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$

Table A.11: Effect of CWA Noncompliance on Local Government Growth using Full Sample

	DO2	Ln(Population)	Ln(Median house price)	High Skill Share
		Panel A: Difference-in-Differences		
Primary'72xPost	0.035	-0.026***	-0.012*	-0.005***
	(0.032)	(0.009)	(0.006)	(0.002)
Marginal effect (%)	0%	-2.58%	-1.15%	-5.17%
		Panel B: Two Stage Least Squares		
Primary'72xPost	2.517***	0.189**	-0.118**	0.036**
	(0.366)	(0.086)	(0.057)	(0.015)
Marginal effect (%)	14 %	22.72 %	-3.37 %	76 %
First Stage F-statistic	30.98	30.44	30.44	17.76
Baseline mean	8.15 mg/l	19,716	$87,447	10%
Clusters	10487	10807	10807	8458
Observations	45456	28270	28270	21227

Note: Includes full sample of wastewater treatment plants, including those that were built after 1972. Includes decade interval years only (1972, 1982, 1992). "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Table reports estimates of β from Eq. 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$. The instrument for $(P_i \times POST_t)$ in Panel B is Eq. 1.4. Standard errors clustered by city. **X** includes all controls listed in Table 1.2, column(4). $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

Table A.12: Ex Ante Noncompliance & Downstream Population Using Full Sample

	Cross Section (1972)		Panel	
	(1)	(2)	(3)	(4)
Downstream Population	-0.028***	-0.029***		
	(0.006)	(0.007)		
Downstream PopulationxStateShare'72xPost			0.033	0.000
			(0.185)	(0.190)
Downstream PopulationxPost			-0.019*	-0.019*
			(0.010)	(0.011)
StateShare'72xPost			-1.671***	-1.761***
			(0.199)	(0.211)
Observations	6155	6124	48248	48242
F Statistic	20.998	7.454	33.532	33.109
River FE	Y	Y		
GeographyControls		Y		
CityControls		Y		
City,YearFE			Y	
RiverPopxYearFE			Y	Y
CityControlsxYearFE				Y

Note: Includes full sample of wastewater treatment plants, including those that were built after 1972. The dependent variable in (1) and (2) is an indicator for primary treatment as of 1972. In (3) and (4), this is interacted with a post CWA indicator. Standard errors clustered by city. Geography controls include distance to river edge and distance to river mouth. CityControls include pre-CWA averages from 1967-1972 of: share of county-level employment in manufacturing and water-polluting industries; and annual federal, state, and local intergovernmental grants. Specifications (3) and (4) interact these baseline controls with year fixed effects. Downstream population is normalized by its standard deviation.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.13: Effect of CWA Compliance on Local Government Budgets Controlling for Time-Varying Effects of State "League of Conservation Vote" Score

PANEL A: Expenditures Per Capita

	Total	Wastewater			Other					
		Total	Capital	Operating	Total	PublicSafety	PublicWorks	GenAdmin	Welfare	Rec
Primary'72xPost	530.013*	191.267***	132.704**	54.239***	160.475	79.671***	63.927	-17.322	29.630	4.570
	(291.471)	(60.049)	(54.243)	(17.957)	(240.720)	(27.203)	(227.514)	(28.482)	(30.320)	(17.804)
Baseline mean	1032.09	66.67	38.65	27.69	650.54	132.39	384.35	61.36	29.43	43.02

PANEL B: Revenues Per Capita

	TotalRevenues			UserFees		Taxes			Debt	
	Total	InterGovt	Own	Total	Wastewater	Total	Property	Sales&License	LongTerm	ShortTerm
Primary'72xPost	718.755***	540.827***	177.776	41.854	76.931***	-108.775	-77.303	-31.501	641.202	136.233
	(219.460)	(107.964)	(180.545)	(72.573)	(23.476)	(78.935)	(69.412)	(32.924)	(925.774)	(88.805)
Baseline mean	1019.02	170.60	848.45	111.91	31.41	389.85	300.87	88.99	1394.71	87.83
First Stage F-statistic	14.36	14.36	14.36	14.36	14.36	14.36	14.36	14.36	14.36	14.36
Clusters	2975	2975	2975	2975	297577	2975	2975	2975	2975	2975
Observations	14866	14866	14866	14866	14866	14866	14866	14866	14866	14866

Note: Dependent variables are in 2012 dollars per capita. Table reports estimates of β from Eq. 1.3: $y_{ist} = \beta_{IV}\widehat{(P_i \times POST_t)} + \mathbf{X_i}\theta_t + D_i\sigma_t + (R \times t) + \tau_t + \kappa_i + \epsilon_{ist}$, where the instrument for $(P_i \times POST_t)$ is Eq. 1.4. $\mathbf{X}$ includes LCV score interacted with year fixed effects, as well as all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.14: Effect of CWA compliance on Water Quality & Municipal Growth, Controlling for Effects of State "League of Conservation Vote" Score

	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary'72xPost	1.543***	0.149*	0.106	0.030**
	(0.327)	(0.080)	(0.065)	(0.014)
First Stage F-statistic	19.56	15.14	15.14	11.39
Baseline mean	8.11	36710	95461	0.11
Marginal effect (%)	19%	15.71%	10.97%	27%
Observations	14177	8264	8264	6366

Note: Population, housing price, and high skill regressions include decade interval years only (1972, 1982, 1992). Dollars in USD 2012 values. "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Table reports estimates of β from Eq. 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + R_\tau + \gamma_t + \nu_i + \varepsilon_{it}$. The instrument for $(P_i \times POST_t)$ is Equation 1.4. Standard errors clustered by city. Includes all controls listed in Table 1.2, column(4). Marginal effect for ln(population) and ln(median house price) calculated as $(\exp(\beta - var(\beta)/2) - 1) \times 100$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.15: CWA Noncompliance and Local Government Budgets by Manufacturing Concentration

	Expenditures			Revenues						
	Total	Wastewater	Other	Total	InterGovt	UserFees	WWUserFees	Taxes	LongTermDebt	ShortTermDebt
Primary '72 x Post	455.598	126.603**	287.612	744.135***	383.492***	13.220	52.456***	26.975	-433.957	101.135
	(288.266)	(50.421)	(257.522)	(178.520)	(80.307)	(57.451)	(19.158)	(56.173)	(909.631)	(64.894)
Primary '72 x Post x	46.408	38.580	-86.285	-81.365	59.987	-18.173	4.610	-42.874	201.536	-49.466
AboveMedianManuf	(154.054)	(27.799)	(123.680)	(132.747)	(48.719)	(49.014)	(10.861)	(42.918)	(656.818)	(48.229)
First Stage F-statistic	7.28	7.28	7.28	7.28	7.28	7.28	7.28	7.28	7.28	7.28
Baseline mean	944.49	57.20	580.90	915.86	159.94	80.13	22.54	350.15	1383.99	67.57
Clusters	2965	2965	2965	2965	2965	2965	2965	2965	2965	2965
Observations	14866	14866	14866	14866	14866	14866	14866	14866	14866	14866

Note: Above Median Manu equals 1 for cities with county-level employment shares in manufacturing above the population median as of 1974. Dependent variables are in 2012 dollars per capita. Standard errors clustered by city. Table reports estimates of δ_{POST} and $\delta_{POST \times \omega}$ where ω equals 1 if a city has above median employment share in manufacturing: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST \times \omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + D_i\sigma_t + (R \times t) + \gamma_t + \nu_i + varepsilon_{it}$. Equation 1.4 instruments for $(P_i \times POST_t)$. Standard errors clustered by city. **X** includes all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.16: Effect of CWA compliance on Water Quality & Municipal Growth by Manufacturing Concentration

	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary '72 x Post	1.777***	0.173***	0.067	0.031***
	(0.308)	(0.064)	(0.052)	(0.011)
Primary '72 x Post x AboveMedianManuf	-0.404*	-0.077	-0.088***	-0.001
	(0.223)	(0.049)	(0.033)	(0.008)
Observations	14177	8264	8264	7618
First Stage F-statistic	10.51	10.78	10.78	8.56

Note: Population, housing price, and high skill share regressions include decade interval years only (1972, 1982, 1992). Dollars in USD 2012 values. "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-1992); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Table reports estimates of δ_{POST} and $\delta_{POST\times\omega}$ where ω equals 1 if a city has above median employment share in manufacturing: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST\times\omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + D_i\sigma_t + (R \times t) + \gamma_t + \nu_i +$ $varepsilon_{it}$. Equation 1.4 instruments for $(P_i \times POST_t)$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.17: Heterogeneity in CWA Noncompliance in Growth Outcomes

	Ln(Population)	Ln(House price)	High Skill Share
Primary'72 x Post	-0.483***	-0.106	0.136***
	(0.150)	(0.133)	(0.027)
Primary'72xPostxJulyTemp	0.019***	0.005	-0.003***
	(0.004)	(0.003)	(0.001)
Primary'72 x Post	0.113	-0.002	0.025**
	(0.069)	(0.053)	(0.011)
Primary'72xPostxDisttoWater	-0.035	-0.040**	-0.004*
	(0.027)	(0.016)	(0.002)
Primary'72 x Post	0.155**	0.065	0.011
	(0.069)	(0.057)	(0.013)
Primary'72xPostxDisttoMouth	-0.002***	-0.003***	0.000
	(0.001)	(0.001)	(0.000)
Observations	8261	8261	7615
First Stage F-statistic	10.49	10.49	8.86

Note: Table reports estimates of δ_{POST} and $\delta_{POST\times\omega}$ from Eq. 1.7: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST\times\omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + (\overline{N}_i \times \gamma_t) + D_i\sigma_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$, where ω is one of three variables: (i) "July Temp" is average maximum July temperature from 1970-2000; (ii) "Distance to Water" is kilometers to nearest large water body, where "large" is a water feature with a stream order of 6 or larger as defined on a scale of 1-13 by USGS; and (iii) "Distance to Mouth" is distance to river mouth in percentiles. Equation 1.4 instruments for $(P_i \times POST_t)$. Standard errors clustered by city. $\mathbf{X}$ includes all controls listed in Table 1.2, column(4).

A.2 Data Sampling Restrictions

First, I restrict my analysis to the roughly 8,318 facilities (out of 20,361 total municipal plants) that were listed as operational wastewater treatment plants as of 1972.

Thus, my analysis does *not* include cities that built a wastewater treatment plant after the CWA came into effect. This restriction increases the likelihood that compliant and noncompliant cities shared important *ex ante* unobservable characteristics that determine economic growth such as willingness of their taxpayer base to invest in long-lasting infrastructure projects. I further exclude plants that ceased operation over time by including only facilities that appear in each decade and in at least half of the 13 surveys between 1972 and 2003. This restriction further drops 22% of the facilities that appeared in the 1972 survey, leaving 6,440 plants. To reduce potential instances of measurement error or misreporting, I exclude wastewater treatment plants that did not meet all of the following criteria: maintains facility type "wastewater treatment plant" as opposed to sewer system, septic, or other (excludes 12.6% of facilities); reports having wastewater treatment plant technology and is recorded as a wastewater treatment facility (excludes 2.3% of facilities); does not cease having a plant if had a plant in the prior year (excludes 3.7% of facilities), and does not downgrade technology type from secondary to primary (excludes 19% of facilities). These additional sample restrictions eliminate approximately 2,462 plants.

These sample restrictions serve to reduce measurement error of treatment plant technology and help to ensure that variation across my treatment and control cities is driven primarily by differences in the CWA technology standard, as opposed to cyclical infrastructure degradation, or structural municipal decline. Appendix Table A.18 shows descriptive statistics comparing my restricted sample to the full population of municipal treatment plants. By utilizing only continuously operating plants, the population size of cities in my analysis is larger, on average, than the mean plant-operating municipality. As a robustness check in Section 1.6.3, I re-estimate my main empirical specifications without imposing either the 1972 criteria or the misreporting exclusions and find qualitatively similar results to those of my restricted sample.

Table A.18: Characteristics of Sample vs Population of Cities with a Municipal Plant

	Sample	Population	P-val for difference in means
Population	35,075	10,074	0.000
Median House Price ($)	97,469	91,212	0.000
Share of population with a college degree	0.111	0.101	0.000
County income per capita	23,345	22,869	0.000
County employment share in manufacturing	0.368	0.356	0.001
County employment share in water-polluting manufacturing	0.148	0.136	0.000
Total revenues pc ($)	1,046	881	0.000
Intergovernment revenues pc ($)	180	157	0.000
Revenues from own sources pc ($)	865	724	0.000
Total taxes pc ($)	417	335	0.000
Total user fees pc ($)	108	68	0.000
Wastewater user fees pc ($)	31	19	0.000
Long-term debt outstanding pc ($)	1,387	1,372	0.924
Short-term debt outstanding pc ($)	94	58	0.000
Total expenditures pc ($)	1,069	916	0.000
Wastewater expenditures pc ($)	78	56	0.000
Total other expenditures pc ($)	654	560	0.000
Number of Cities	2,964	7,895	
Panel Frequency	5.2	4.5	
Observations	14,860	33,624	

Note: All variables measured as means in 1967 and 1972. P-value denotes significance of difference in means. See Section 1.4 for details on data sources.

A.3 Dissolved Oxygen & Water Quality

I focus on dissolved oxygen as my preferred measure of water quality for two main reasons. First, dissolved oxygen plays a crucial role in water ecosystems: insufficient levels of dissolved oxygen can cause fish, amphibians, and plant life to die off. Because the primary goal of the CWA was to restore and maintain the biological integrity of US surface waters and to make all water "fishable and swimmable," dissolved oxygen provides an holistic measure of the effectiveness of the CWA technology mandate in meeting the CWA goals. Second, dissolved oxygen is directly impacted by municipal sewerage. Secondary treatment can increase dissolved oxygen levels by removing harmful bacteria from the wastewater effluent, including fecal coliforms, and nutrients such as nitrogen and phosphorous (Minnesota Pollution Control Agency 2009). These pollutants are potentially hazardous to human health and can induce eutrophication, thereby reducing the clarity and aesthetic value of surface waters.

In summary, high levels of dissolved oxygen correlate with water quality attributes that are likely to be valued by individuals, such as visual aesthetics and the opportunities for fishing and swimming recreation. However, dissolved oxygen may not provide the most salient metric for water quality. Visual clarity of water, for example, can be high even if the water quality is inhospitable to aquatic life and dissolved oxygen levels are low. In absence of large fish kills or algal blooms, variation in dissolved oxygen, nitrogen or phosphorous may be unobservable to the eye (Leggett and Bockstael 2000). Turbidity provides a closer measure of water clarity, however turbidity is not closely related to overall ecosystem health. To the extent that the dissolved oxygen improvements from secondary treatment are largely undetected by local residents, my estimates on the local value of water quality from wastewater treatment infrastructure will be attenuated toward a null effect.

A.4 Comparison of Water Quality Results

Keiser and Shapiro (2018) focus on municipal wastewater treatment plants as the relevant treatment unit and employ a triple difference estimation strategy to show that water quality downstream of grant-receiving plants improved significantly more than that of plants that did not receive federal grants. Their study isolates changes to water quality within 25 miles downstream of a wastewater treatment plant after the plant receives an infrastructure grant. In contrast, my paper considers changes to average surface water quality within 25 miles of the city center for cities under pressure to comply with the CWA mandate. Both *ex ante* compliant and noncompliant cities could receive EPA infrastructure grants. A second major difference is that Keiser and Shapiro (2018) focus on dissolved oxygen deficit (among others) as their measure of water quality, whereas I focus on dissolved oxygen in its compound form. Because healthy levels of dissolved oxygen can differ across water bodies depending on the ambient temperature, salinity, and depth, researchers sometimes consider dissolved oxygen saturation (or dissolved oxygen deficit) as a standardized measure. I focus on the compound form to reduce potential mis-measurement from converting dissolved oxygen to dissolved oxygen saturation, and then aggregating up water quality readings to a city-level average. Because my empirical approach relies on within-city variation, any cross sectional differences across geographic locations with respect to their water chemistry is unlikely to bias my results.

A.5 Heterogeneity in Growth Responses by City Size

Despite having significantly lower costs of compliance, larger noncompliant cities experienced more adverse growth outcomes relative to smaller noncompliant cities. Appendix Table A.19 shows estimates of Eq. 1.7. Population of large cities grew approximately 10% slower, on average, between 1972 and 2002 relative to smaller cities; however, the total effect for both is statistically indistinguishable from zero. Median housing prices among larger noncompliant cities declined approximately 9% more on average over the 30–year study period relative to noncompliant smaller cities. Both classes of city size show evidence of high–skill sorting following the CWA.

A.5.1 Mechanisms for Heterogeneous Responses

Despite having lower compliance costs, larger cities appear to have substantially lower overall benefits from the infrastructure mandate relative to smaller cities. I explore two potential explanations for this surprising result. The first relates to housing supply elasticity. If larger cities generally have less elastic housing supply, then cost of living increases in larger cities will more readily manifest in price declines relative to smaller cities. I test for differences in housing supply response across larger relative to smaller cities in Appendix Table A.20. While all results are noisily estimated, point estimates for the larger cities suggest very limited response in housing quantity following the CWA infrastructure mandate, consistent with inelastic housing supply. In contrast, smaller cities have 2-10 times larger housing supply responses, indicating more elastic supply. Noticeably, housing supply in smaller cities responds positively to the CWA mandates. In the following section, I present a potential explanation for this positive response.

The second potential mechanism driving more adverse growth responses in larger cities relates to heterogeneity in water quality improvements from the CWA. I test this mechanism by exploiting the local average treatment effect of my instruments, which predicts *ex ante* noncompliance driven by low downstream population. A city may have low downstream population either because it is situated near the mouth of a river, or if it is situated at any point on a sparsely-populated river. Considering the former case, large and small cities located near the end of a populous river are likely to have different improvements to water quality following the CWA. Larger cities are more likely to have been able to pressure upstream neighbors to adopt secondary treatment prior to the CWA. Consequently, the CWA is less likely to improve water quality generated from upstream polluters for larger cities near the end of a river. Smaller cities, in contrast, may have greater benefits from mandated wastewater treatment owing both to their own investments as well as those of upstream neighbors. I explore this potential source of heterogeneity by categorizing cities according to their population size relative to the total population of their river network. Appendix Table A.21 presents estimates of Eq. 1.7 where ω equals 1 for cities whose own population relative to the total river population are in the top 50th percentile of their river network's population share distribution.[1] Because the complier group of my instrument is likely to be cities located near the end of a river mouth, differences in relative population size among this group can inform how much differences in upstream treatment adoption may drive heterogeneous local growth responses to the CWA. Results of Appendix Table A.21 are suggestive that cities occupying a larger share of their river population have *lower* increases in water quality, housing supply, housing prices, and share of high skill residents relative to cities occupying a smaller share of their river population. Growth differences across city size are most pronounced in the shorter time period presented in the top panel. Consequently, larger cities may have faced worse growth responses to the uniform CWA mandate because they were less likely to benefit from mandated upstream pollution abatement.

[1]The US Hydrography dataset organizes the US river system into 409 distinct river networks, defined by a common terminating location (i.e. river mouth). I define a river system by the "terminal path" identifier.

Table A.19: Effect of CWA Compliance on Local Government Growth by City Size

	Dissolved Oxygen	Ln(Population)	Ln(Median house price)	High Skill Share
Primary'72xPost	1.455***	0.130*	-0.007	0.013
	(0.320)	(0.067)	(0.056)	(0.014)
Primary'72xPostxAboveMedian	0.045	-0.086***	-0.090***	0.007*
	(0.088)	(0.019)	(0.013)	(0.004)
Total Above Median Effect	1.500***	0.044	-0.097*	0.02
	(0.313)	(0.067)	(0.055)	(0.013)
Baseline mean	8.108 mg/l	$36,710	95,461	11.02%
First Stage F-statistic	10.981	10.770	15.501	12.210
Observations	14177	8264	8264	6366

Note: Includes decade years only (1972, 1982, 1992). "High Skill Share" is the share of city population with 4 or more years of college (1972) or a bachelor's degree or higher (1982-2002); and includes only balanced panel of cities with annual observations for educational attainment (e.g., cities with a population over 2,500 as of 1970). Table reports estimates of δ_{POST} and $\delta_{POST\times\omega}$ from Eq. 1.7: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST\times\omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + (\overline{N}_i \times \gamma_t) + D_i\sigma_t(R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$, where Equation 1.4 instruments for $(P_i \times POST_t)$. Standard errors clustered by city. Includes all controls listed in Table 1.2, column(4). Marginal effect for ln(population) and ln(median house price) calculated as $(\exp(\beta - var(\beta)/2) - 1) \times 100$, where $var(\beta) = (se(\beta_1)^2 + se(\beta_2)^2 + 2cov(\beta_1, \beta_2)$.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.20: Housing Supply Response by City Size

	Ln(Housing Units)		
	(1)	(2)	(3)
Primary'72xPost (Below Med)	0.137	0.119	-0.056
	(0.087)	(0.085)	(0.086)
Primary'72xPost (Above Med)	0.000	-0.011	-0.033
	(0.083)	(0.081)	(0.082)
First Stage F-statistic	13.231	13.627	12.97
Controls	Y	Y	Y
PopxYearFE		Y	Y
Incl. 1972-1982			Y
Baseline mean	16816	16816	16815
Observations	10489	10489	8610

Note: Includes decade years only (1972, 1982, 1992, 2002). Column (3) includes only the first decade following the CWA. Number of housing units per municipality sourced from IPUMS NGHIS. Table reports estimates of δ_{POST} and $\delta_{POST\times\omega}$ from Equation 1.7: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST\times\omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + (\overline{N}_i \times \gamma_t) + D_i\sigma_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$, where Equation 1.4 instruments for $(P_i \times POST_t)$. Standard errors clustered by city. "Controls" include all controls listed in Table 1.2, column(4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.21: CWA Noncompliance Effect by City Size as Share of River Population

	DO_2	Ln(Housing Units)	Ln(House Price)	Ln(Population)	Share College
			Includes 1972-1992		
Primary'72xPost	1.390***	0.053	0.358***	0.086	0.039***
	(0.344)	(0.083)	(0.075)	(0.077)	(0.014)
Primary'72xPostxAboveMedian	-0.057	-0.031	-0.096***	0.006	-0.007
	(0.113)	(0.035)	(0.024)	(0.025)	(0.005)
Observations	14177	7618	8264	8264	7618
			Includes 1972-2002		
Primary'72xPost	0.979***	0.046	0.020	0.035	0.043***
	(0.320)	(0.095)	(0.057)	(0.082)	(0.015)
Primary'72xPostxAboveMedian	0.026	-0.024	-0.094***	0.017	-0.011*
	(0.105)	(0.040)	(0.019)	(0.030)	(0.006)
Observations	18998	10489	11127	11127	10489

Note: Table reports estimates o δ_{POST} and $\delta_{POST\times\omega}$ from Eq. 1.7: $y_{it} = \delta_{POST}(P_i \times POST_t) + \delta_{POST\times\omega}(P_i \times POST_t \times \omega_i) + \mathbf{X_i}\theta_\tau + (\overline{N}_i \times \gamma_t) + D_i\sigma_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$, where Eq. 1.4 instruments for $(P_i \times POST_t)$. Standard errors clustered by city. $\mathbf{X_i}$ includes all controls listed in Table 1.2, column(4). Above Median equals 1 if a city's population share of its river network is in the top 50th percentile among cities on its river network. There are 409 river networks. Aggregate median river population share is 0.003%. $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$

Table A.22: Federal Grant Receipt and City Population Size

	Population Size			
	Above Median	Below Median	Obs.	Baseline Mean
Panel A: Annual Wastewater Expenditures				
Total per capita	136.17***	166.81***	11,924	$66.7
	(51.75)	(53.34)		
Capital per capita	106.02**	121.61**	11,924	$38.68
	(48.71)	(50.58)		
Panel B: Annual Federal grant receipts				
Grant flow per capita	152.36***	144.72***	11,924	$14.63
	(37.65)	(39.60)		

Note: Table reports estimates of β from Eq. 1.2: $y_{it} = \beta(P_i \times POST_t) + \mathbf{X_i}\theta_t + (R \times t) + \gamma_t + \nu_i + \varepsilon_{it}$ where $(P_i \times POST_t)$ is using Eq. 1.4. The first stage F-statistic for Panels A and B is 14.64. Includes years 1967-1987. Standard errors clustered by city. Includes all controls listed in Table 1.2 column 4. $^{*}p < 0.10,^{**} p < 0.05,^{***} p < 0.01$. Panel B uses grant data defined as "Federal intergovernmental other" by the Census of Governments, which included municipal receipt of federal CWA infrastructure grants from 1967 through 1992. In addition to grants for sewerage, the "other" federal grants category includes grants for generally infrequent cost needs, such as disaster assistance and homeland security. Note, the post-CWA increase in federal aid more than covers total wastewater expenditures for larger cities, but falls short by more than $20 per capita for smaller cities.

APPENDIX B

APPENDIX FOR CHAPTER 2

B.1 Labor & Energy Cost Calculations

We calculated fuel costs using reported fuel-use data from NTD and estimates of retail fuel prices for each of the reported fuel inputs. Fuel prices were sourced from the U.S. Energy Information Administration. Fuel price data vary by state and year. NTD does not report labor costs separately from total operations costs. Thus, we infer labor costs as the residual operating costs net of our calculated energy costs. The following figure shows the breakdown of line items included in operating expenses as reported by transit agencies to NTD. Vehicle operating costs include: labor costs (including salaries, wages, benefits, and pensions), costs of fuel, materials and supplies, utility costs, taxes, and liability costs. With the exception of fuel inputs, we cannot separately identify these costs from total operating costs. Thus, our estimate of labor costs is an upward bound, and includes expenses for materials and miscellaneous expenses.

Figure B.1: Operating Costs Line Items as Reported to NTD

2013 Urban Reporting Manual

Operating Expenses form (F-30)

Home | e-File | Annual | Monthly Ridership | Safety & Security | Notes | Issues | Reports | Communications | Sys Admin | Search

Form Name: Operating Expenses (F-30) **Mode:** LR **Service:** DO

Line	Expense Object Class	a **Vehicle Operations 010 Total**	b **Vehicle Maintenance 041 Total**	c **Non-Vehicle Maintenance 042 Total**	d **General Administration 160 Total**	e **Total Modal Expenses**
	Labor (501)					
01	Operators' salaries and wages (01)					
02	Other salaries and wages (02)					
03	**Fringe Benefits (502)**					
04	**Services (503)**					
	Materials and Supplies (504)					
05	Fuel and lubricants (01)					
06	Tires and tubes (02)					
07	Other materials and supplies (99)					
08	**Utilities (505)**					
09	**Casualty and Liability Costs (506)**					
10	**Taxes (507)**					
13	**Miscellaneous Expenses (509)**					
15	**Total Modal Expenses**					

B.2 Details on the Welfare Analysis Calculation

The procedure we use to estimate welfare loss from high public transit cost is as follows: We begin by specifying the consumer compensating variation obtained from reducing bus transit operating costs by privatizing.

Consider the consumer's minimum cost of achieving a utility level, u, given prices p:

$$g(p,u) = \min_{x}\{p * x | u(x) \geq u\} \tag{B.1}$$

The compensated demand function for good $x_i \in x$ is given by:[1]

$$x_i = \frac{\partial g(p,u)}{\partial p_i} = g(p,u) \tag{B.2}$$

Glaister (1974). Let α_1 and p_1 be the price of bus fare under current privatization levels (where majority of bus transit is publicly operated), and under the counterfactual scenario of complete privatization, respectively. Let all other prices faced by the consumer, $\hat{p}$, be held constant. The loss to the consumer from purchasing bus fare, x_1, at α_1 instead of the lower efficient price p_1 can be expressed by:[2]

$$A = \int_{p_1}^{\alpha_1} g_1(z,\hat{p},u)dz = g(\alpha_1,\hat{p},u) - g(p_1,\hat{p},u) \tag{B.3}$$

The area A is the change in consumer surplus due to the change in transit operating costs.

[1] The difference between the compensated and Marshallian demand function is the income effect. Deaton (1974) shows that if the utility function is additively separable, and the expenditure on the good is a small fraction of income, then the compensated and uncompensated elasticities will be close. For purposes of this exercise, we posit that bus transit is additively separable to all other consumption goods and services.

[2] Small and Rosen (1981) provide a detailed outline for the theoretical justification for measuring price-induced utility changes as areas to the left of the relevant compensated demand curves.

To apply this theoretical framework to our empirical analysis, we specify $g(.)$ as the transit agency supply curve. We have posited a production function whereby transit agencies employ inputs in fixed proportions to produce a bus mile. Thus, marginal costs are equal to unit operating costs, $\frac{C}{VMT}$. Let the transit agency supply curve equal their marginal costs, thus $g(.) = \frac{C}{VMT}$. The consumer's expenditure for bus transit is proportional to the transit agency's marginal costs of providing bus service.

In addition to a supply curve, the other inputs necessary to calculate welfare loss include a demand curve for bus transit, a base unit cost of transit, and the counterfactual unit cost of transit under complete privatization. Our base unit cost of transit is the operating costs per VMT, $\frac{C}{VMT}$, predicted by our model under privatization levels observed as of 2011. We predict $\frac{C}{VMT}$ under current privatization using the RD specification presented in Section 2.6.1. To generate the estimates, we regress the log of total operating costs per VMT on predicted privatization share (city mayoral party affiliation instruments for privatization share in the fuzzy RD design), following the right-hand-side of Eq. 2.9. Since we are interested in predicting total operating costs, rather than labor operating costs, we control for energy input prices in addition to each of the covariates impacting labor unit costs presented in Table 2.3. Under 2011 observed privatization levels, the predicted average $\frac{C}{VMT}$ is \$6.63. The counterfactual $\frac{C}{VMT}$ predicted by our model if all transit agency operations are completely privatized is approximately \$4.16.[3]

[3] As our dependent variable is in log form, we calculate the predicted unit cost values using a smearing adjustment following Greene (2003): $\hat{y} = \exp\left(X'\hat{\beta} + \sigma/2\right)$

We calculate a demand curve for bus transit using an elasticity measure of bus ridership, or passenger trips, with respect to bus fare. According to Gagnepain and Ivaldi (2002),Oum et al. (1992), and TCRP (2004), this elasticity measure is approximately -0.4. Specifically, $\epsilon_R = -0.4 = (\%\Delta\text{Ridership})/(\%\Delta\text{Fare})$. Given the aggregate ridership of bus transit in 2011 was approximately 5.2 billion, we can solve for the change in ridership given an incremental change in bus transit fare. This requires that we specify a transformation from our predicted unit operating costs to bus passenger fare.[4]

To perform this transformation, we multiply the predicted $\frac{C}{VMT} = \$6.63$ by the average VMT per passenger trip observed in 2011, 0.45, to calculate predicted unit cost per passenger trip, $\frac{C}{Trip}$. By employing a constant term, $\frac{C}{Trip}$, for the transformation of predicted unit costs, we assume the number of passenger trips per VMT is unaffected by privatization levels. This static simplification is consistent with Gagnepain and Ivaldi (2002) who modeled the relationship between VMT and passenger trips as a reduced form function of population characteristics, density, and road congestion in a particular year. Since our welfare analysis is particular to the year 2011, our methodology is consistent with Gagnepain and Ivaldi (2002) in assuming that

[4]We have characterized transit agency supply in terms of VMT, while consumer demand is in terms of passenger trips. Transit agency costs and revenues are driven by two different output variables that are closely related. Outputs from the perspective of the transit agency are VMT, whereas the passenger uses VMT as inputs in the "production" of their final consumption good, which is passenger miles or passenger-trips. ("This characterization of transit output is quite useful if we construct a general equilibrium model of transit, in which transit firms supply intermediate-type outputs while passengers, who demand them, generate final-type outputs." (Berechman 2013). Prior literature has frequently employed similar transformations as those discussed here, recognizing the difference between VMT as intermediate output and passenger trips as final output in transit systems (Gagnepain and Ivaldi 2002).

the relationship between capacity and demand is static in a particular year. Our previous empirical findings on the impacts of privatization on transit agency VMT and rider occupancy further support this simplification. Appendix Table B.4 shows that privatization does not have a deterministically significant impact on either VMT, PMT, or rider occupancy rate.

Our transformation further assumes transit agencies price at marginal cost. In this way we can interpret unit cost per passenger trip as an approximation to the predicted bus fare per passenger trip. We acknowledge that marginal cost pricing is not representative of public transit pricing schemes. Public transit pricing is, in general, heavily subsidized to make up for the large fixed costs of infrastructure. Mohring (1972) gives a detailed discussion of dynamic optimal pricing schemes. We abstract away from identifying the pricing scheme of transit agencies in order to highlight general welfare effects from inefficient public transit. In summary:

$$g(\alpha_1, \hat{p}, u) = \$6.63 * 0.45 = \$2.98 \text{per trip}$$
$$g(p_1, \hat{p}, u) = \$4.16 * 0.45 = \$1.87 \text{per trip}$$
$$\%\Delta \text{Ridership} = \%\Delta \text{Cost per Trip} * \epsilon_R$$

where $\%\Delta$Cost per Trip are based on incremental value changes between \$1.87 and \$2.98. We calculate the area A in Eq. B.3 using the trapezoidal method. A in Eq. B.3 corresponds to the shaded area in Figure 2.9.

The procedure for calculating welfare loss specific to Boston, Chicago, and San Antonio mimic the above discussion with the following exceptions. First, rather than using an aggregate ridership value of 5.2 billion, the base ridership value is based on city-specific estimates. Second, the transformation calculations employ the city-specific average VMT per passenger trip observed in 2011. These values and their sources are shown in the following table:

City Transit Agency Values as of 2011

City (Agency)	Total Annual Passenger Trips (mn.)	Total Annual VMT (mn.)	VMT per Passenger Trip
Boston (MBTA)	109.9	27.2	0.25
Chicago (CTA)	310.4	58	0.19
San Antonio (VIA)	44.2	22.4	0.51

Source: NTD.gov Data Tables, Table 19: Transit Operating Statistics for 2011

B.3 Alternative Empirical Strategies as Robustness Checks

In this section, we present two alternative empirical strategies based on different identifying assumptions from that of the RD analysis.

B.3.1 Union Contract Cycle as Instrument

We carry out estimation of Eq. 2.9 and instrument for the privatization share using union contract cycle schedules for 10 transit agencies. When a union's labor contract is over and up for negotiation, the transit agency manager may find it easier to privatize a portion of their operations to private contractors. However, the status of the labor contract (whether it is active or not) should not be correlated with a transit agency's unit labor costs. Appendix Table B.5 shows the first stage results. The end of a labor contract cycle is positively correlated with privatization share, as we would expect. Appendix Table B.6 shows the impact of privatization on labor costs under OLS and GMM employing the contract cycle instrument. The cost saving estimates are very similar in magnitude to the RD estimates discussed in Section 2.6.1; a 1% increase in the share of bus miles that are privatized reduces labor costs per mile by between 0.6% and 0.9%.

B.3.2 Using Subway Cities as Controls

Due to limited within-city variation in the NTD data, we are unable to employ city-level fixed effects in our main specifications discussed in Section 2.6.1. Thus, our identification comes from variation in transit agencies within the same state with differing levels of privatized bus miles. The analyses in Section 2.6.1 include various controls for transit agency characteristics, such as fleet size, type, and the share of buses that use alternative fuel. However, it is possible that certain unobservable city-specific factors affect both privatization decisions and unit labor costs, and these unobservables generate the differences we observe across agencies within the same state.

To address this endogeneity concern, we look at the subset of agencies in our NTD data whose UZA also has a subway system. Subway systems, by their nature as massive public works projects, are public entities and cannot be privatized. Comparing the operating costs of bus transit to the operating costs of subway transit in the same city will fully control for any city-wide unobserved effects that impact transit costs, which may confound our main cross-city identification. We carry out two specifications: in column (1) of Appendix Table B.7 the dependent variable is the log of bus operating costs per VMT, and the regression controls for subway operating costs per VMT. In column (2), the dependent variable is the log of the ratio of bus to subway operating costs. Under both specifications, a 1% increase in privatized bus miles is associated with a 1% decrease in operating costs per VMT. These results are again, consistent with those found in Section 2.6.1 under the RD method.

B.4 Figures & Tables

Figure B.2: Privatization share among non-dominant agencies. Bin size=0.04. Sample includes 662 agency-year pairs.

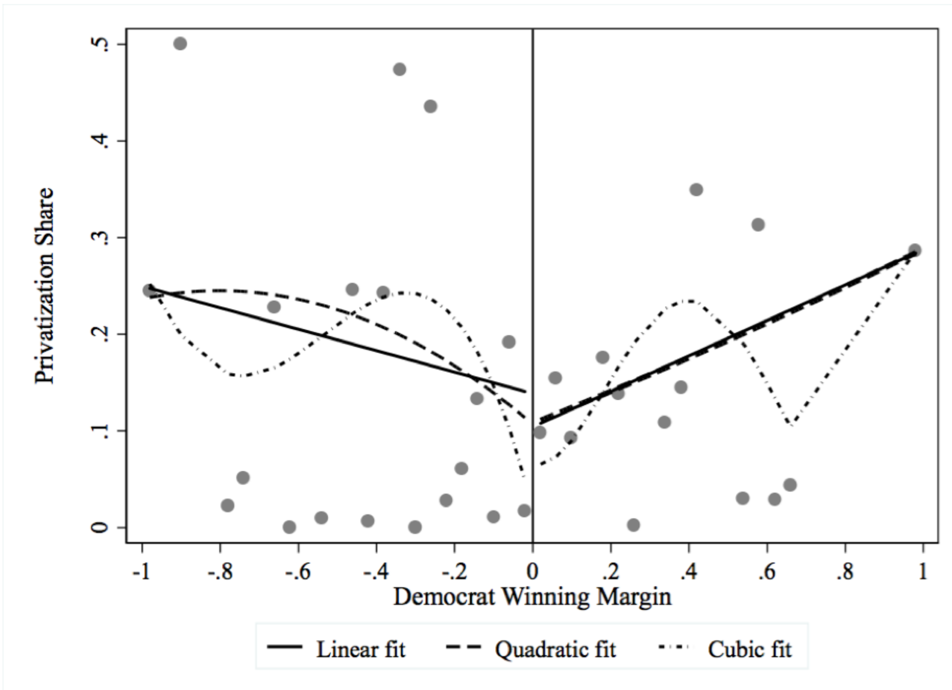

Source: "Non-dominant" agencies include those that do not have the largest average annual vehicle miles traveled from 1998-2011 in a UZA. Each dot corresponds to the average transit agency privatization share following mayoral election t, given the margin of victory obtained by Democrats in election t. The solid and dashed lines represent the predicted values of privatization share from linear, quadratic, and cubic polynomial control functions of the winning margin.

Figure B.3: Predicted Log(Labor Costs) Given Pre-Existing Agency Characteristics at the Democrat Winning Margin Threshold. Bin size=0.04. Sample includes 1,299 agency-year pairs

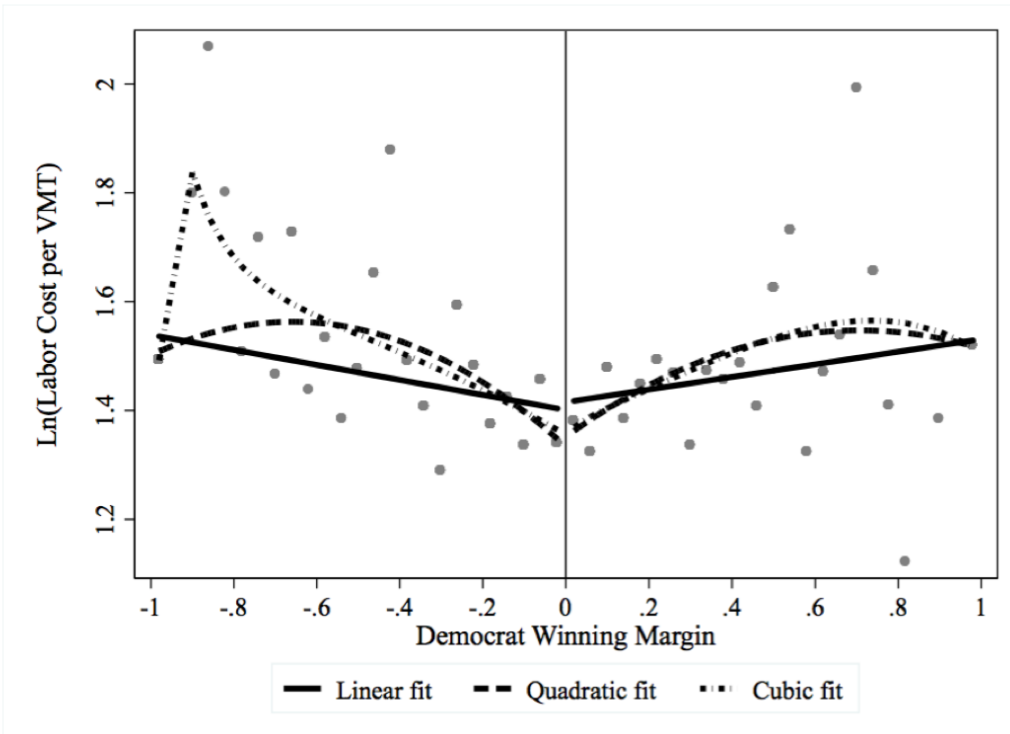

Source: Each dot corresponds to the average predicted log(labor costs) per VMT following mayoral election t, given the margin of victory obtained by Democrats in election t. The solid and dashed lines represent the linear, quadratic, and cubic polynomial control functions of best fit. Predicted values of log(labor costs) per VMT are obtained from a regression model relating observed log (labor costs) per VMT to each of the 11 predetermined covariates, x, listed in Appendix Table B.1 as follows: $y_{it} = \mathbf{x}\kappa + \varepsilon_{it}$. The vector of coefficients, κ, is estimated from a random subsample of 150 observations. Next, we generate out-of-sample predictions of log(labor costs) per VMT: $\hat{y}_{it} = \mathbf{x}\kappa$. The resulting covariate index function can be interpreted as the best linear prediction of mean log labor costs given the vector of predetermined variables (Card et al. 2015).

Figure B.4: Privatization Share at the Democrat winning margin threshold.
Bin size=0.04

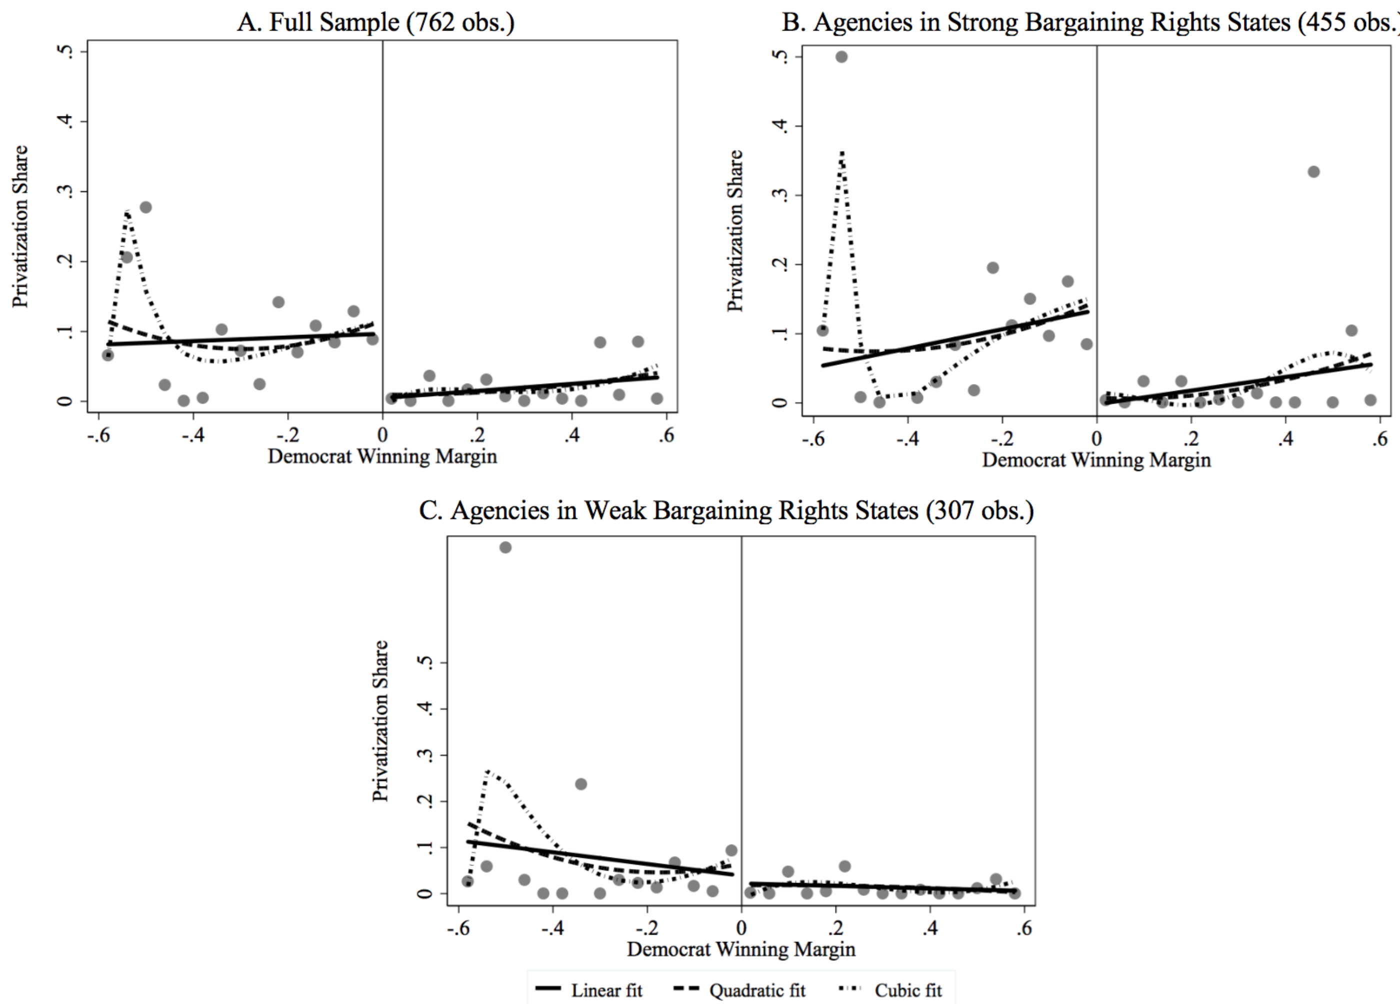

Source: Each dot corresponds to the average transit agency privatization share following mayoral election t, given the margin of victory obtained by Democrats in election t. The solid and dashed lines represent the predicted values of privatization share from linear, quadratic, and cubic polynomial control functions of the winning margin.

Figure B.5: Residual Ln(Labor Costs per VMT) at the Democrat winning margin threshold among agencies in strong bargaining rights states. Sample includes 880 agency-year pairs. Bin size=0.04

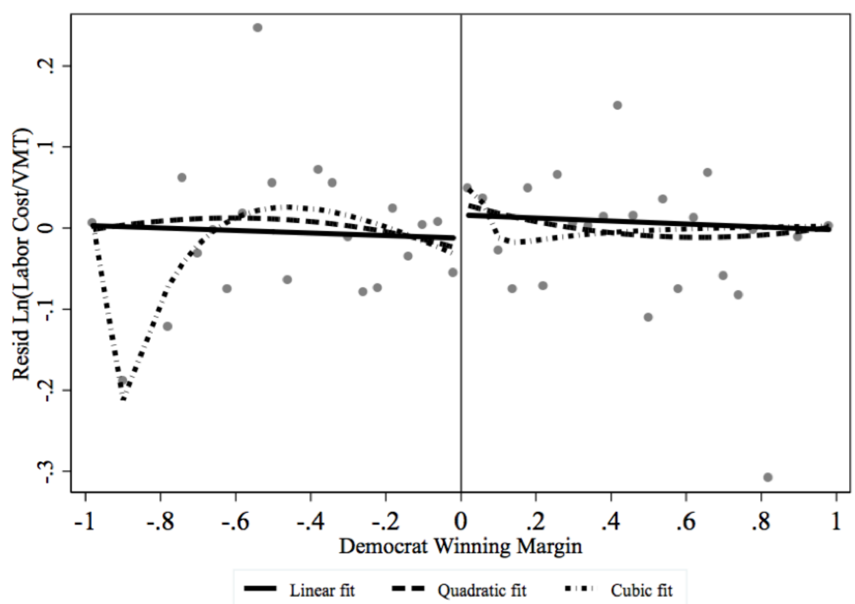

Source: Each dot corresponds to the average residual log unit labor cost following mayoral election t, given the margin of victory obtained by Democrats in election t. The solid and dashed lines represent the predicted values of residual log unit labor cost from linear, quadratic, and cubic polynomial control functions. Figures plot residuals generated from OLS regressions of log unit labor cost on all controls listed in Table 2.3, and year and state fixed effects. Standard errors clustered at the UZA level.

Figure B.6: Local Polynomial Regressions with Varying Bandwidth: Effect of Privatization on Log(Labor Cost per VMT)

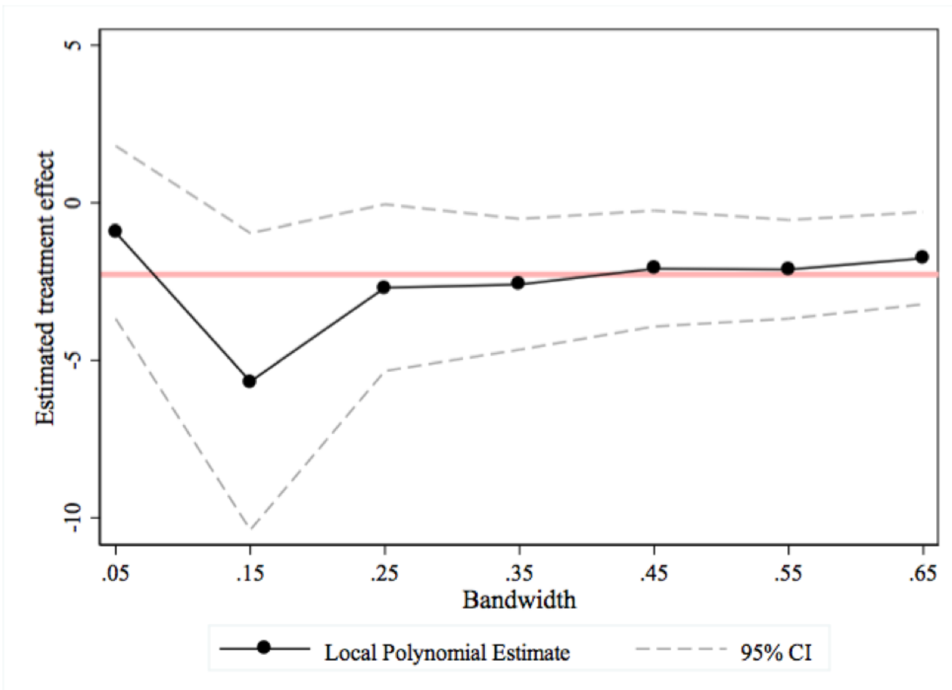

Source: This figure re-estimates the local polynomial specification in column (7) of Table 2.5 under varying bandwidths. The estimated treatment effect from (7) in Table 2.5 is represented by the shaded red line, which has an optimal CCT bandwidth of 0.294. All treatment effects estimated using local polynomial Fuzzy RD estimation with robust CCT confidence intervals, triangular kernel, and the denoted bandwidth, and are of order two in winning margin.

Figure B.7: Mayoral Regime Characteristics at the Democrat Winning Margin Threshold. Sample includes 1,444 agency-year pairs. Bin size=0.04.

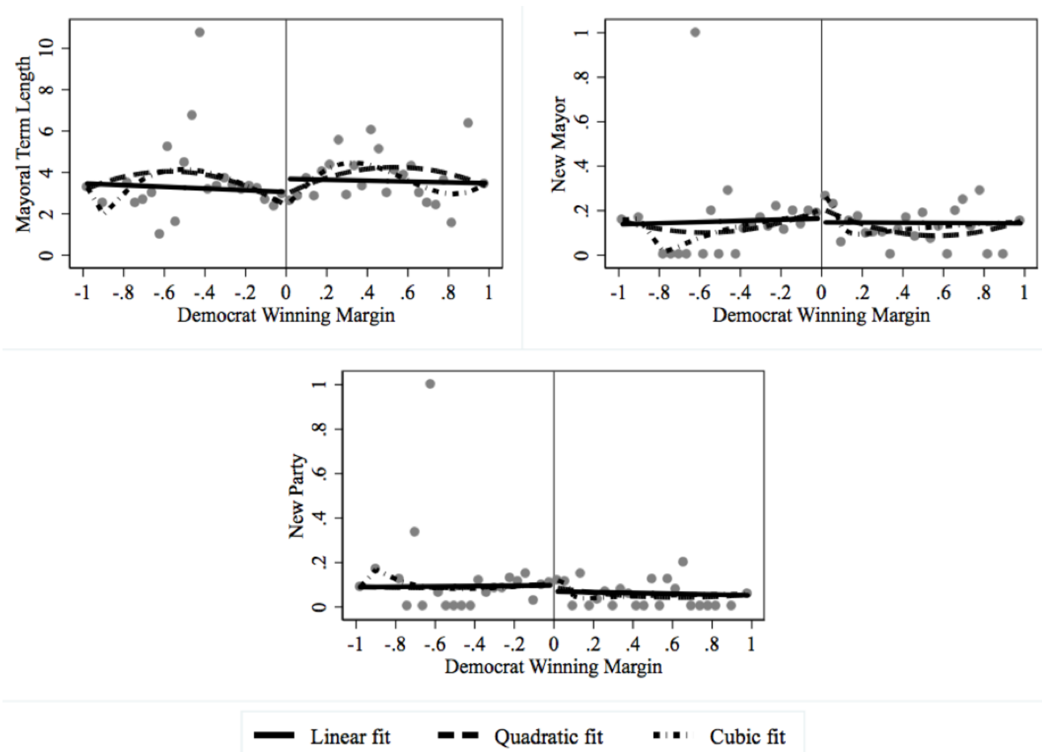

Source: Each dot corresponds to the average noted UZA mayoral regime characteristic following mayoral election t, given the margin of victory obtained by Democrats in election t. The solid and dashed lines represent the predicted values of the noted characteristic from linear, quadratic, and cubic polynomial control functions without covariates.

Figure B.8: Residual service attributes at the Democrat winning margin threshold. Bin size=0.04. Sample includes 1430 agency-year pairs.

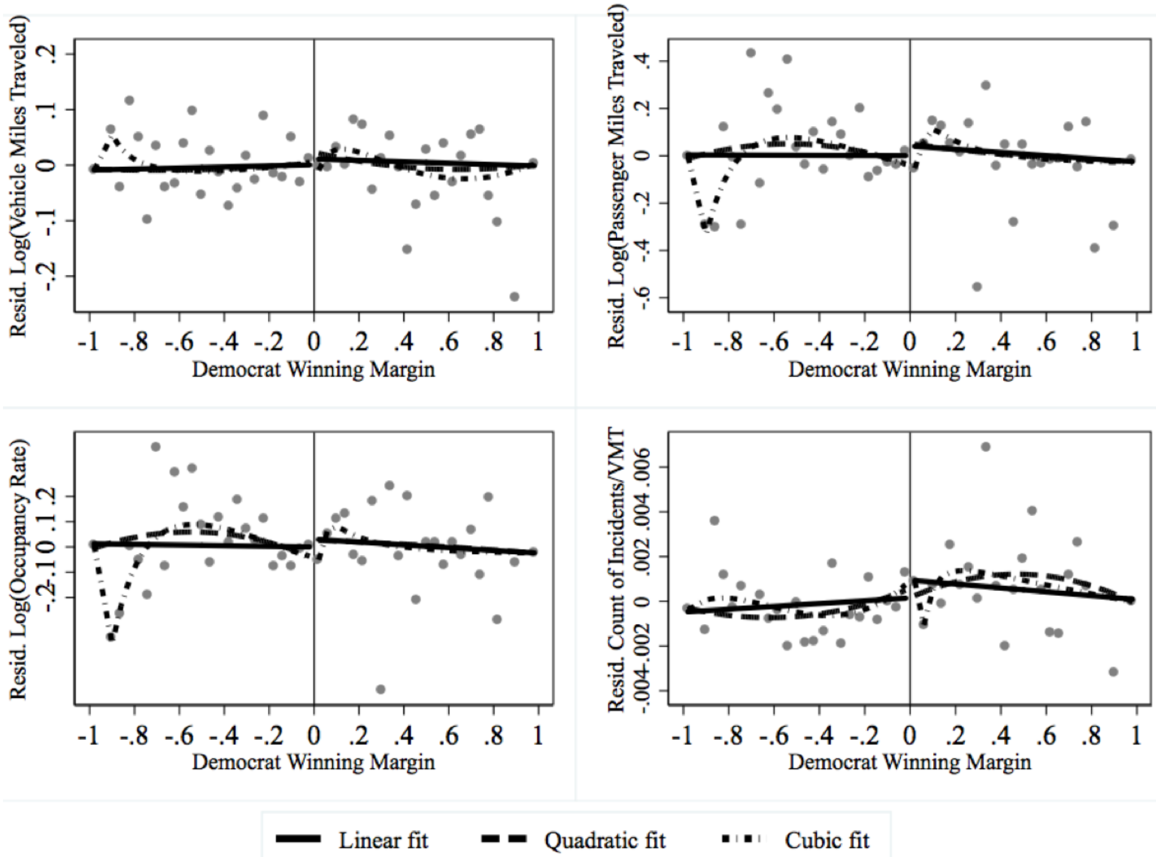

Source: Each dot corresponds to the average noted residual service attribute following mayoral election t, given the margin of victory obtained by Democrats in election t. The solid and dashed lines represent the predicted values of residual service attribute from linear, quadratic, and cubic polynomial control functions. Figures plot residuals generated from OLS regressions of each service attribute. All regressions include all controls listed in Table 2.3, as well as log(fleet size) and log(average fleet age) as of 1997 to control for initial service demand conditions, and year and state fixed effects. Standard errors clustered at the UZA level.

Table B.1: Test for Discontinuity in Covariate Means Near Winning Margin Threshold

	Covariate Mean				P-Value
	[-0.6,-0.3)	[-0.3, 0)	[0, 0.3)	[0.3, 0.6]	
A. UZA Characteristics					
Weekly (low skill) wage ($)	577.47	589.8	609.17	605.99	0.983
Road congestion index[1]	0.891	0.916	0.927	0.94	0.181
Home price ($ '000)	316.843	283.183	259.375	255.399	0.403
Share unionized workers[2]	0.057	0.064	0.064	0.059	0.41
Population Density (pop/sq. mi)	2836.4	2399.6	2397.6	2395.5	0.716
B. Agency Characteristics					
Share independent agency	0.552	0.574	0.661	0.446	0.175
Share city agency	0.385	0.367	0.239	0.384	0.082
Fleet size (no. buses)	216.1	162.3	206.6	391.2	0.449
Fleet age (years)	7.2	7.6	7.6	7.6	0.738
Share of CNG buses[3]	0.108	0.076	0.059	0.069	0.644
Share of hybrid buses	0.012	0.011	0.01	0.007	0.074
Observations	96	324	230	112	

Notes: Columns 1-4 report the variable mean for UZA's within the noted Democrat winning margin window. P-values are based on discontinuity estimates, λ, from second-order polynomial regressions in Democrat winning margin: $y_{it} = \sigma D(S)_{it} + \lambda_1 M_{it} + \lambda_2 M_{it}^2 + \lambda_3 D_{it} M_{it} + \lambda_4 D_{it} M_{it}^2 + \nu_{it}$ or ($|M| \leq 0.6$) where y is a covariate, M is the winning margin, and $D = 1$ if the winning mayor is a Democrat. Standard errors were clustered at the UZA level.

[1] Road congestion index measures density of traffic across urban areas. An index greater than 1 indicates an undesirable level of area-wide congestion. (Source: Texas A&M Transportation Institute)

[2] Means are conditional on share of unionized workers in strong bargaining rights states.

[3] CNG is compressed natural gas, a cleaner type of fuel than gasoline, diesel, or propane.

Table B.2: The Effect of Mayoral Party Affiliation and Transit Agency Service Offering

	Annual Vehicle Miles (th.)		Annual Passenger Miles (th.)		Occupancy Rate	
	(1)	(2)	(3)	(4)	(5)	(6)
MV Window	[-1, 1]	[-0.6,0.6]	[-1, 1]	[-0.6,0.6]	[-1, 1]	[-0.6,0.6]
Mayor Effect	-6.372***	-2.442	-8.942***	-4.465	0.142	0.202
	(2.081)	(2.44)	(3.312)	(4.138)	(0.35)	(0.476)
Observations	1444	762	1444	762	1444	762
Adjusted R2	0.769	0.752	0.659	0.647	0.66	0.627

Notes: Dependent variable is denoted at top of each column. Occupancy Rate is PMT/ VMT. Mayor effect is the total effect of Democrat mayor (λ_1) and Democrat mayor interacted with strong bargaining rights state indicator (λ_2) from the following specification: $y_{it} = \sigma D(S)_{it} + \lambda_1 M_{it} + \lambda_2 M_{it}^2 + \lambda_3 D_{it} M_{it} + \lambda_4 D_{it} M_{it}^2 + \nu_{it}$ or ($|M| \leq 0.6$) where y is a covariate, M is the winning margin, and $D = 1$ if the winning mayor is a Democrat. All regressions are OLS. Sample includes dominant transit agencies only. z_{it} is the vector of covariates listed in Table 2.3 as well as year and state fixed effects. Standard errors clustered at UZA level. $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

Table B.3: Effect of Privatization on Frequency of Fatal and Non-Fatal Incidents

	(1)	(2)	(3)	(4)
Estimator	NB	NB	FRD	CCT
Dependent Var	No. of Yearly Incidents		No. of Yearly Incidents/VMT	
Privatization share	-0.329**	-0.507	-0.027	5.299
	(0.167)	(0.379)	(0.027)	(8.579)
Year & State FE	Y	Y	Y	
UZA & Agency Controls	Y	Y	Y	
Sample	Full	RD	RD	RD
Observations	2597	1030	1030	661
AIC	13954.9	6799.6	—	—

Notes:The dependent variable for (1-2) is the number of annual non-fatal and fatal incidents reported by the transit agency. Regressions (1-2) are negative binomial, and control for log of annual VMT. The dependent variable for (3-4) is the number of annual non-fatal and fatal incidents per VMT. (3) includes controls for quadratic polynomial of winning margin and their interaction with Democrat mayor indicator. The first stage for (3) is reported in Table 2.4, column (3). Excluded instruments in (3) are Democrat (d); and Democrat x StrongBarg. (4) employs local polynomial RD estimation with robust confidence intervals developed by Calonico et al. (2014); (4) is estimated using a triangular kernel and mean-squared error optimal CCT bandwidth selector of 0.240. Agency and UZA controls listed in Table 2.3. AIC measures goodness of fit for nonlinear models; smaller AIC is a better fit. AIC calculated as -2(log likelihood) + 2(dof +1). Standard errors clustered at the UZA level in parentheses.$^{*}p < 0.10$,** $p < 0.05$,*** $p < 0.01$.

Table B.4: Effect of Privatization on Total VMT, PMT and Rider Occupancy

	(1)	(2)	(3)	(4)	(5)	(6)
Estimator	OLS	FRD	OLS	FRD	OLS	FRD
Dependent Var	Log (Vehicle Miles Traveled)		Log (Passenger Miles Traveled)		Log (Occupancy)	
Privatization share	0.144***	0.28	0.014	-0.826	-0.13	-1.233
	(0.035)	(0.42)	(0.089)	(0.777)	(0.084)	(0.81)
Year & State FE	Y	Y	Y	Y	Y	Y
UZA & Agency Controls	Y	Y	Y	Y	Y	Y
Observations*	3672	1430	3672	1430	3672	1430
F-Stat for Excl. IV	–	3.1	–	3.74	–	3.1

Notes: Dependent variable is log annual VMT in (1) and (2), log annual PMT in (3) and (4), and annual occupancy rate in (5) and (6). Occupancy is annual passenger miles divided by annual VMT. All regressions include controls for log(fleet size) and log(average fleet age) as of 1997 to control for initial service demand conditions. The fuzzy RD specifications further control for quadratic polynomials of Democrat winning margin and their interactions with the Democratic mayor dummy, and use both the Democratic mayor assignment variable as well as its interaction with strong state bargaining rights indicator in the first stage. The fuzzy RD estimates are based on dominant agencies only. UZA & Agency Controls listed in Table 2.3. Standard errors are clustered at the UZA level. $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$. *Sample sizes are based on transit agencies that report passenger miles traveled. Of our full sample of 3,706 agency-year observations, 34 observations do not report passenger miles traveled, 14 of which have mayoral election data.

Table B.5: First Stage Regressions Using Contract Cycle as IV

	(1) OLS	(2) OLS
Labor contract cycle	0.016***	0.042***
	(0.005)	(0.005)
Hourly wage (low skill)	-0.009	-0.067
	(0.032)	(0.084)
Road congestion index	0.014	0.103
	(0.111)	(0.134)
Avg. home price	-0.001	-0.004
	(0.004)	(0.006)
Union share*StrongBarg (d)[1]	-0.1	0.217
	(0.178)	(0.357)
Ln(population density)	1.051***	1.225***
	(0.401)	(0.163)
Ln(number of buses)	0.173***	0.277***
	(0.062)	(0.04)
Ln(average bus age)	-0.036**	-0.085***
	(0.014)	(0.0)2
Share of CNG buses	-0.011	0.049
	(0.043)	(0.05)
Share of hybrid buses	-0.255***	-0.237
	(0.082)	(0.333)
Year FE	Y	Y
Agency FE		Y
Observations	130	140
Adjusted R2	0.623	0.726

Notes: The dependent variable is privatization share. The data include 10 transit agencies for which we have union contract cycle data. Column (1) is estimated after first-differencing (FD) while column (2) uses transit agency fixed effects (FE). Contract cycle variable is defined as the number of labor contract cycles since 1998. (d) denotes a binary variable. $*p < 0.10, **p < 0.05, ***p < 0.01$.

[1] Omitted category is weak bargaining rights state.

Table B.6: Labor Cost Regression Using Contract Cycles as IV

	(1) OLS w/ FD	(2) OLS w/ FE	(3) GMM w/ FD	(4) GMM w/ FE
Privatization share	-0.646***	-0.762***	-0.389	-0.869***
	(0.129)	(0.072)	(0.598)	(0.141)
Hourly wage (low skill)	-0.103	0.02	-0.099	0.008
	(0.093)	(0.116)	(0.098)	(0.115)
Road congestion index	0.449*	0.429***	0.458*	0.418***
	(0.267)	(0.139)	(0.273)	(0.136)
House Price	0.014**	0.003	0.014**	0.004
	(0.007)	(0.006)	(0.007)	(0.006)
Union share*StrongBarg (d)[1]	0.065	-0.158	0.092	-0.12
	(0.292)	(0.434)	(0.296)	(0.428)
Ln(population density)	0.228	0.587***	-0.033	0.641***
	(0.296)	(0.145)	(0.698)	(0.151)
Ln(population)	0.190***	0.160**	0.143	0.203**
	(0.057)	(0.068)	(0.099)	(0.081)
Ln(number of buses)	0.021	0.050***	0.03	0.045**
	(0.031)	(0.019)	(0.041)	(0.019)
Ln(average bus age)	0.009	0.059	0.014	0.056
	(0.084)	(0.052)	(0.085)	(0.052)
Share of CNG buses	-0.893***	-0.645***	-0.840***	-0.733***
	(0.305)	(0.207)	(0.302)	(0.216)
Share of hybrid buses	-0.646***	-0.762***	-0.389	-0.869***
	(0.129)	(0.072)	(0.598)	(0.141)
Year FE	Y	Y	Y	Y
Agency FE		Y		Y
F-Stat for Excl. IV	–	–	7.22	59.79
Observations	130	140	130	140
Adjusted R2	0.144	0.933	0.137	0.932

Notes: The dependent variable is log(labor cost per VMT). The data include 10 transit agencies for which we have union contract cycle data. Columns (1) and (3) are estimated after first-differencing (FD) and columns (2) and (4) include transit agency fixed effects (FE). The IV for privatization is the number of labor contract cycles since 1998. (d) denotes a binary variable. $*p < 0.10, **p < 0.05, ***p < 0.01$.

[1] Omitted category is weak bargaining rights state.

Table B.7: Effect of Privatization on Bus Costs per Mile Using Subway as Control

Dependent Var	(1) log(Bus cost per VMT)	(2) log(Bus cost/Subway cost)
Privatization share	-1.008***	-1.002***
	(0.115)	(0.226)
Log(total cost per mile for subway)	0.248***	
	(0.044)	
Hourly wage (low skill)	0.183	0.117
	(0.185)	(0.236)
Ln(Nat. Gas price)	0.017	-0.016
	(0.014)	(0.027)
Ln(Diesel price)	-0.326	-1.529**
	(0.282)	(0.558)
Road congestion index	0.093	0.049
	(0.202)	(0.296)
House Price	0.012	0.052
	(0.015)	(0.030)
Union share*StrongBarg (d)[1]	-0.118	0.925
	(0.579)	(0.900)
Ln(population density)	0.244	-0.462
	(0.302)	(0.365)
Independent agency (d)[2]	-0.223	-0.803***
	(0.218)	(0.179)
City agency (d)[2]	-0.753***	-0.745**
	(0.136)	(0.238)
Ln(number of buses)	0.024	0.087
	(0.048)	(0.073)
Log(average bus age)	0.001	-0.176
	(0.065)	(0.099)
Share of CNG buses	-0.254***	-0.236*
	(0.070)	(0.122)
Share of hybrid buses	-0.191	-0.55
	(0.179)	(0.395)
Year & State FE	Y	Y
Observations	140	140
Adjusted R2	0.977	0.953

Notes: The dependent variable in (1) is log(total bus operating cost per VMT). The dependent variable in (2) is log(total bus operating cost per VMT / total subway operating cost per subway mile traveled). All regressions are OLS. Total operating costs include labor and material costs, fuel costs, administration, and maintenance. (d) denotes a binary variable. The data include 10 transit agencies that run both a bus system and a subway system. Standard errors are clustered at the transit agency level. $*p < 0.10, **p < 0.05, ***p < 0.01$.

[1] Omitted category is weak bargaining rights state.

[2] Omitted agency-type category is "Other", which includes state DOT or subsidiary agencies.

APPENDIX C

APPENDIX FOR CHAPTER 3

C.1 Figures & Tables

Figure C.1: Beijing Subway System Expansion

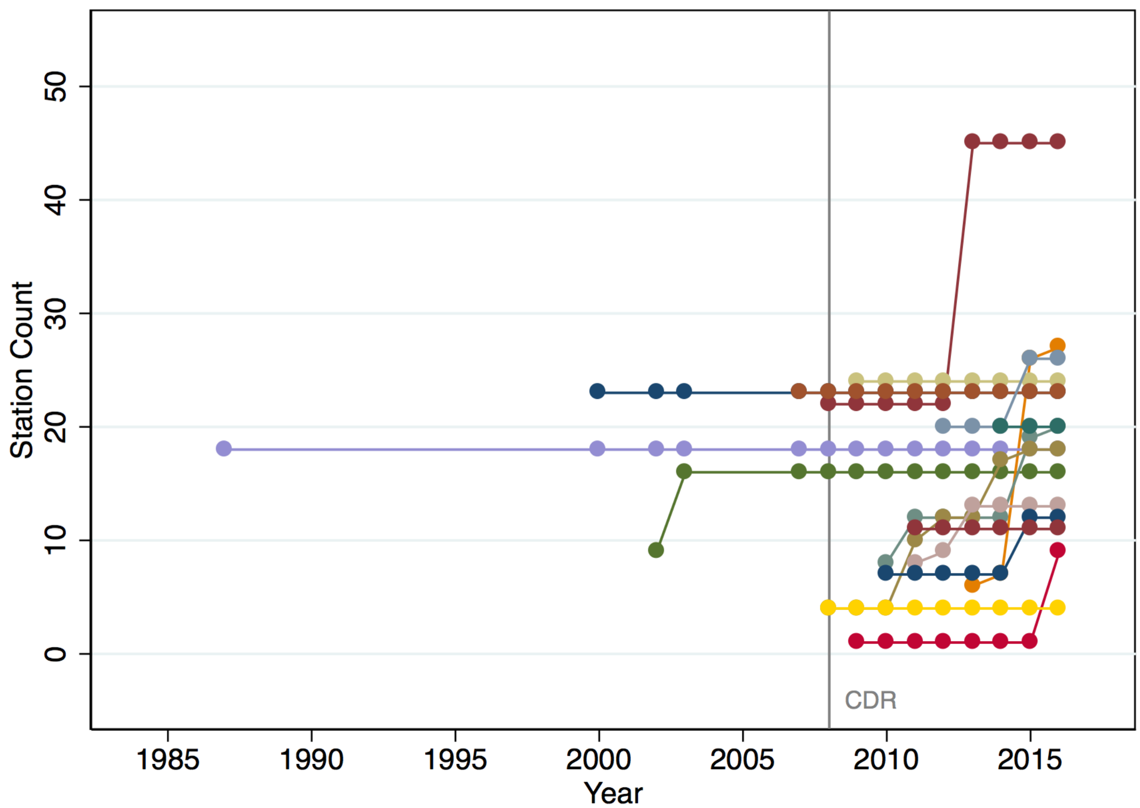

Figure C.2: Income-Distance Gradient with various definitions of CBD

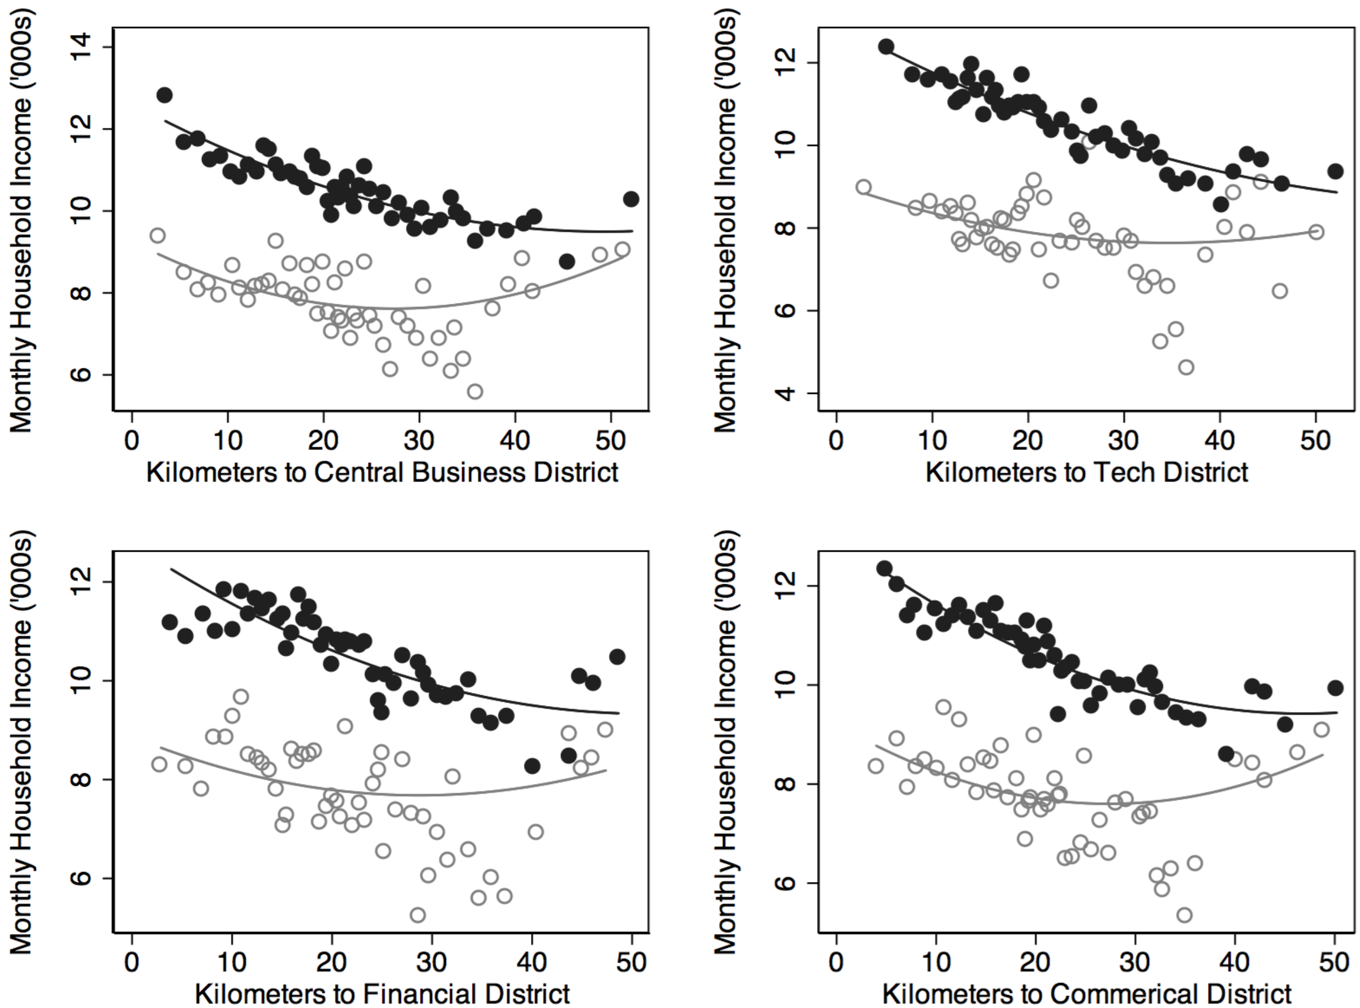

Note: Figure plots mean household income by binned distance to city center. Includes 50 bins each with 330 and 1400 obs per bin in pre and post periods, respectively. Means residualized by distance to the nearest subway.

Table C.1: Effects of CDR policy on City Center-Price gradient - Varying Definition of City Center

	Geographic Center	Technology District	Software District	Financial District	Beijing "CBD"	Embassy District	Business Park	Shopping District
CBD Proximity (km) x CDR	0.006	0.007**	0.009***	0.003	0.010***	0.006	0.010***	0.010**
	(0.005)	(0.003)	(0.003)	(0.005)	(0.003)	(0.004)	(0.003)	(0.005)
CBD Proximity (km)	-0.020	-0.015	-0.013	-0.025	-0.003	-0.005	-0.003	-0.018
	(0.018)	(0.013)	(0.012)	(0.018)	(0.010)	(0.012)	(0.009)	(0.017)
Avg Price Premium / Km	$60.85	$71.82	$89.07	$34.04	$96.67	$58.32	$95.73	$99.45
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Jiedao FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	83711	83711	83711	83711	83711	83711	83711	83711
Adjusted R^2	0.611	0.612	0.613	0.612	0.614	0.611	0.615	0.612

Note: Dependent variable is ln(total price per square meter in 2007 real Yuan). Standard errors clustered at building complex level. Sample spans 12 mos. before and after CDR. Average price premium evaluated at the mean unit size (100 sqm for owner-occupied) within 4 and 6 km of the nearest business district. Controls include distance to nearest subway station, year, month, and district fixed effects; controls for complex age, age2, size, floor-area ratio, green space, no. total floors; controls for housing unit size, decoration level, floor level, facing direction, and no. bedrooms.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.2: Comparison of CDR Effects with Xu et al. (2015)

	CDR is October 11				CDR is July 20			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ln(Distance to Subway)	-0.081**	0.071	-0.022	-0.085**	-0.096**	0.010	-0.057***	-0.124**
	(0.038)	(0.049)	(0.023)	(0.039)	(0.038)	(0.057)	(0.019)	(0.050)
Ln(Distance to Subway) x CDR	-0.069	-0.107***	-0.068***		-0.085	-0.077**	-0.049***	
	(0.049)	(0.030)	(0.023)		(0.055)	(0.031)	(0.013)	
Location Controls	Y			Y	Y			Y
Full Controls			Y				Y	
Jiedao FE		Y	Y			Y	Y	
6 mos. pre CDR				Y				Y
Observations	18989	18981	18981	5917	14061	14051	14051	6016
Adjusted R^2	0.223	0.630	0.720	0.166	0.189	0.654	0.731	0.101

Note: Dependent variable is ln(total price per square meter in 2007 real Yuan). Standard errors clustered at building complex level. Excludes housing units near newly built subway stations. Location Controls include distance to city center, and an indicator for locating within a key school district. Controls include year, month, and district fixed effects; and controls for complex age, age2, size, floor-area ratio, green space, no. total floors; controls for housing unit size, decoration level, floor level, facing direction, and no. bedrooms. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.3: CDR Policy and Income sorting near the Central Business District - Varying Definition of City Center

	Geographic Center	Technology District	Software District	Financial District	Beijing "CBD"	Embassy District	Business Park	Shopping District
Ln(Household Income)× CDR	-0.010**	0.002	-0.009**	0.003	-0.020***	-0.015***	-0.021***	-0.011***
	(0.004)	(0.006)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.004)
Ln(Household Income)	-0.001	-0.011**	-0.002	-0.009***	0.009**	0.002	0.008*	0.000
	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Zip FE	Y	Y	Y	Y	Y	Y	Y	Y
Subway Line FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Observations	18135	18135	18135	18135	18135	18135	18135	18135
Adjusted R^2	0.979	0.957	0.959	0.980	0.967	0.972	0.963	0.978

Note: Dependent variable is ln(Distance to CBD (km)). Income is household monthly income ('000 yuan). CDR equals 1 after July 20 2008. Standard errors clustered by household. Sample spans July 20, 2006-July 20, 2010. Demographic controls include husband and wife age, employment rank, education, employer type, and tenure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure C.3: Price Proximity Gradient, excluding areas near newly-built stations

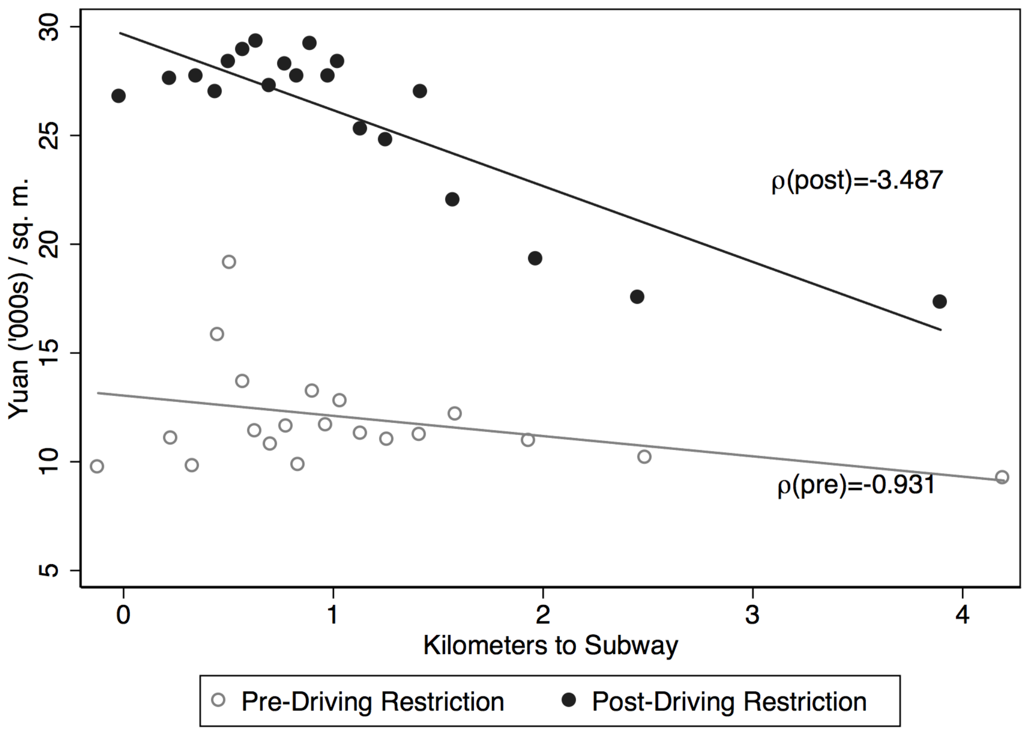

Note: Figure plots mean house price for each of 20 distance bins. Each dot represents 12,100 and 76,500 obs per bin in pre and post, respectively. Means are residualized by distance to the central business district. ρ(pre) and ρ(post) are regression coefficients. Includes years 2005-2016. Source: Real estate transaction dataset.

Table C.4: Effects of CDR policy on Subway-Price gradient, excluding areas near newly-built stations

	(1)	(2)	(3)	(4)	(5)
Subway Proximity (km) x CDR	0.037	0.012	0.013***	0.014***	0.015***
	(0.029)	(0.008)	(0.004)	(0.004)	(0.004)
Subway Proximity (km)	0.097***	0.058***	0.049***	0.048***	0.046***
	(0.015)	(0.011)	(0.017)	(0.011)	(0.011)
Avg Price Premium / Km	$373.45	$117.89	$127.32	$136.63	$150.62
Controls		Y	Y	Y	Y
Jiedao FE			Y	Y	Y
Year-Month FE				Y	
DistrictxYear-Quarter FE					Y
Observations	12486	12486	12479	12479	12479
Adjusted R^2	0.194	0.617	0.727	0.738	0.746

Note: Dependent variable is ln(total price per square meter in 2007 real Yuan). Sample includes housing units in building complexes that do not change in their proximity to subway stations from 2005 through 2016. Standard errors clustered at jiedao level. Sample spans 12 mos. before and after CDR. Average price premium evaluated at the mean unit size (100 sqm for owner-occupied) within 1 and 1.5 km of a subway station. Controls include year, month, and district fixed effects; controls for complex age, age2, size, floor-area ratio, green space, no. total floors; controls for housing unit size, decoration level, floor level, facing direction, and no. bedrooms. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

BIBLIOGRAPHY

Abramowitz, A. I. (2010). *The disappearing center: Engaged citizens, polarization, and American democracy*. Yale University Press.

Albouy, D., Christensen, P., and Sarmiento-Barbieri, I. (2018). Unlocking amenities: Estimating public-good complementarity. Technical report, National Bureau of Economic Research.

Alesina, A., Baqir, R., and Easterly, W. (2000). Redistributive public employment. *Journal of Urban Economics*, 48(2):219–241.

Alm, J., Buschman, R. D., and Sjoquist, D. L. (2011). Rethinking local government reliance on the property tax. *Regional Science and Urban Economics*, 41(4):320–331.

Alonso, W. et al. (1964). Location and land use. toward a general theory of land rent. *Location and land use. Toward a general theory of land rent.*

Andreen, W. L. (2013). Success and backlash: The remarkable (continuing) story of the clean water act. *Geo. Wash. J. Energy & Envtl. L.*, 4:25.

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

APTA (2015). 2015 public transportation fact book appendix. Historical Tables 1& 7.

Ash, E., Morelli, M., and Van Weelden, R. (2015). Forthcoming."elections and divisiveness: Theory and evidence.". *Journal of Politics*.

Baicker, K. (2001). Government decision-making and the incidence of federal mandates. *Journal of public economics*, 82(2):147–194.

Baicker, K. and Gordon, N. (2006). The effect of state education finance reform on total local resources. *Journal of Public Economics*, 90(8-9):1519–1535.

Baltimore Sun Editorial Board (2019). How to help with escalating water bills. *The Baltimore Sun*.

Banzhaf, H. S. (2018). Difference-in-differences hedonics. Working Paper.

Banzhaf, H. S. and Walsh, R. P. (2008). Do people vote with their feet? an empirical test of tiebout. *American Economic Review*, 98(3):843–63.

Baum-Snow, N. and Kahn, M. E. (2000). The effects of new public projects to expand urban rail transit. *Journal of Public Economics*, 77(2):241–263.

Baumol, W. J. (1967). Macroeconomics of unbalanced growth: the anatomy of urban crisis. *The American economic review*, pages 415–426.

Bel, G. and Rosell, J. (2016). Public and private production in a mixed delivery system: Regulation, competition and costs. *Journal of Policy Analysis and Management*, 35(3):533–558.

Berechman, J. (2013). *Public transit economics and deregulation policy*. Elsevier.

Black, A. (1991). Privatization of urban transit: A different perspective. *Transportation Research Record*, 1297:69–75.

Bleakley, H. and Lin, J. (2012). Portage and path dependence. *The Quarterly Journal of Economics*, 127(2):587–644.

Bockstael, N. E., Hanemann, W. M., and Kling, C. L. (1987). Estimating the value of water quality improvements in a recreational demand framework. *Water Resources Research*, 23(5):951–960.

Boustan, L. P. and Margo, R. A. (2009). Race, segregation, and postal employment: New evidence on spatial mismatch. *Journal of Urban Economics*, 65(1):1–10.

Boyle, K. J., Poor, P. J., and Taylor, L. O. (1999). Estimating the demand for protecting freshwater lakes from eutrophication. *American journal of agricultural economics*, 81(5):1118–1122.

Brown, D. (2018). Environmental Engineer and Manager, City of Philadelphia Water Department.

Brueckner, J. K. and Neumark, D. (2014). Beaches, sunshine, and public sector pay. theory and evidence on amenities and rent extraction by government workers. *American Economic Journal: Economic Policy*, 6(2):198–230.

Brueckner, J. K. and Rosenthal, S. S. (2009). Gentrification and neighborhood housing cycles: will america's future downtowns be rich? *The Review of Economics and Statistics*, 91(4):725–743.

Brueckner, J. K., Thisse, J.-F., and Zenou, Y. (1999). Why is central paris rich and downtown detroit poor?: An amenity-based theory. *European economic review*, 43(1):91–107.

Cain, L. (2005). Sanitation in chicago: A strategy for a lakefront metropolis. Technical report, The Encyclopedia of Chicago, Chicago, IL.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Card, D., Lee, D. S., Pei, Z., and Weber, A. (2015). Inference on causal effects in a generalized regression kink design. *Econometrica*, 83(6):2453–2483.

Carlino, G. A. and Saiz, A. (2008). Beautiful city: Leisure amenities and urban growth. *Working Paper No. 08-22*.

Carrillo, P. E., Malik, A. S., and Yoo, Y. (2016). Driving restrictions that work? quito's pico y placa program. *Canadian Journal of Economics/Revue canadienne d'économique*, 49(4):1536–1568.

Cellini, S. R., Ferreira, F., and Rothstein, J. (2010). The value of school facility investments: Evidence from a dynamic regression discontinuity design. *The Quarterly Journal of Economics*, 125(1):215–261.

Chay, K. Y. and Greenstone, M. (2005). Does air quality matter? evidence from the housing market. *Journal of Political Economy*, 113(21):376–424.

Chen, Y., Jin, G. Z., Kumar, N., and Shi, G. (2013). The promise of beijing: Evaluating the impact of the 2008 olympic games on air quality. *Journal of Environmental Economics and Management*, 66(3):424–443.

Chong, A., La Porta, R., Lopez-de Silanes, F., and Shleifer, A. (2014). Letter grading government efficiency. *Journal of the European Economic Association*, 12(2):277–298.

Conlan, T. J. (1994). *Federally induced costs affecting state and local governments*. US Advisory Commission on Intergovernmental Relations.

Cooper, R. and John, A. (1988). Coordinating coordination failures in keynesian models. *The Quarterly Journal of Economics*, 103(3):441–463.

Copeland, C. (1999). Clean water act: a summary of the law. Congressional Research Service, Library of Congress Washington, DC.

Copeland, C. (2012). Clean water act and pollutant total maximum daily loads (tmdls). Congressional Research Service, Library of Congress Washington, DC.

Copeland, C. (2015). Funding for epa water infrastructure: A fact sheet. Congressional Research Service, Library of Congress Washington, DC.

Crémer, J. and Palfrey, T. R. (2000). Federal mandates by popular demand. *Journal of Political Economy*, 108(5):905–927.

Cromwell, E., Ihlanfeldt, K., et al. (2015). Local government responses to exogenous shocks in revenue sources: Evidence from florida. *National Tax Journal*, 68(2):339–376.

Davis, L. W. (2008). The effect of driving restrictions on air quality in mexico city. *Journal of Political Economy*, 116(1):38–81.

Deaton, A. (1974). A reconsideration of the empirical implications of additive preferences. *The Economic Journal*, 84(334):338–348.

Deming, D. J., Cohodes, S., Jennings, J., and Jencks, C. (2016). School accountability, postsecondary attainment, and earnings. *Review of Economics and Statistics*, 98(5):848–862.

Deprez, E. (2013). N.y. mta labor talks drag as union fights part-time bus drivers.

Diamond, R. (2017). Housing supply elasticity and rent extraction by state and local governments. *American Economic Journal: Economic Policy*, 9(1):74–111.

Dilger, R. J. (2013). Unfunded mandates reform act: History, impact, and issues. Congressional Research Service, Library of Congress Washington, DC.

DiNardo, J. and Lee, D. S. (2004). Economic impacts of new unionization on private sector employers: 1984–2001. *The Quarterly Journal of Economics*, 119(4):1383–1441.

Ding, C., Song, Y., and Knaap, G. (2005). Growth scenarios for beijing 2020. Lincoln Institute of Land Policy Working Paper.

DiSalvo, D. (2010). The trouble with public sector unions. *National Affairs*, 5(Fall):3–19.

DiSalvo, D. (2015). *Government against itself: Public union power and its consequences*. Oxford University Press, USA.

Duflo, E. and Pande, R. (2007). Dams. *The Quarterly Journal of Economics*, 122(2):601–646.

Earnhart, D. (2004). Regulatory factors shaping environmental performance at publicly-owned treatment plants. *Journal of Environmental Economics and Management*, 48(1):655–681.

Economic Development Research Group, Inc. (2016). Failure to act: Closing the infrastructure investment gap for america's economic future. Technical report, American Society of Civil Engineers, Washington, DC.

EPA (1973). Estimate of municipal wastewater treatment facility requirements. Technical report, Clean Watershed Needs Survey.

EPA (1976). An analysis of construction cost experience for wastewater treatment plants. Technical Report 430/9-76-002, US Environmental Protection Agency Office of Water Program Operations, Washington, DC.

EPA (1988). The municipal sector study: Impacts of environmental regulations on municipalities. Technical Report MCDEPA 230-09/88-038, US Environmental Protection Agency, Washington, DC.

EPA (1994). Current 301(h) waiver recipients and 301(h) applications pending final decision. Accessed: 2019-04-18.

Eskeland, G. S. and Feyzioglu, T. (1997). Rationing can backfire: the "day without a car" in mexico city. *The World Bank Economic Review*, 11(3):383–408.

Fairfax, V. and Hamilton, V. (2000). Progress in water quality: an evaluation of the national investment in municipal wastewater treatment. *Washington, DC: US EPA*.

Feler, L. and Senses, M. Z. (2017). Trade shocks and the provision of local public goods. *American Economic Journal: Economic Policy*, 9(4):101–43.

Ferreira, F. and Gyourko, J. (2009). Do political parties matter? evidence from us cities. *The Quarterly journal of economics*, 124(1):399–422.

Fields, G. and Emshwiller, J. R. (2011). Federal offenses: A sewage blunder earns engineer a criminal record. *Wall Street Journal*.

Fisher, R. C. (2015). *State and local public finance*. Routledge.

Foster, H. S. and Beattie, B. R. (1979). Urban residential demand for water in the united states. *Land Economics*, 55(1):43–58.

Freeman, R. B. and Valletta, R. (1988). The effects of public sector labor laws on labor market institutions and outcomes. In *When public sector workers unionize*, pages 81–106. University of Chicago Press.

FTA (2016). National transit database. Accessed: 2016-06-06.

Gagnepain, P. and Ivaldi, M. (2002). Incentive regulatory policies: the case of public transit systems in france. *RAND Journal of Economics*, 33:605–629.

Gamper-Rabindran, S., Mastromonaco, R., and Timmins, C. (2011). Valuing the benefits of superfund site remediation: Three approaches to measuring localized externalities. *NBER Working Paper*, (16655).

GAO (1980). Epa should help small communities cope with federal pollution control requirements, report to congress by the comptroller general of the united states. Technical report, US General Accounting Office.

Gentzkow, M., Shapiro, J., Taddy, M., et al. (2016). Measuring polarization in high-dimensional data: Method and application to congressional speech. Technical report.

Gerber, E. R. and Hopkins, D. J. (2011). When mayors matter: estimating the impact of mayoral partisanship on city policy. *American Journal of Political Science*, 55(2):326–339.

Giuliano, G. and Lave, C. (1989). The high cost of a bargain: Winning the right to use part-time transit drivers. *Transportation Research Part A: General*, 23(2):151–159.

Glaeser, E. L., Kahn, M. E., and Rappaport, J. (2008). Why do the poor live in cities? the role of public transportation. *Journal of urban Economics*, 63(1):1–24.

Glaister, S. (1974). Generalised consumer surplus and public transport pricing. *The Economic Journal*, 84(336):849–867.

Greene, W. H. (2003). *Econometric analysis*. Pearson Education India.

Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the of manufactures. *Journal of Political Economy*, 110(6):1175–1219.

Greenstone, M. and Gallagher, J. (2008). Does hazardous waste matter? evidence from the housing market and the superfund program. *The Quarterly Journal of Economics*, 123(3):951–1003.

Guerrieri, V., Hartley, D., and Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, 100:45–60.

Guo, T., Englehardt, J., and Wu, T. (2014). Review of cost versus scale: water and wastewater treatment and reuse processes. *Water Science and Technology*, 69(2):223–234.

Hahn, J., Todd, P., and Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209.

Hanford, P. L. and Sokolow, A. D. (1987). Mandates as both hardship and benefit: The clean water program in small communities. *Publius: The Journal of Federalism*, 17(4):131–146.

Hart, O., Shleifer, A., and Vishny, R. W. (1997). The proper scope of government: theory and an application to prisons. *The Quarterly Journal of Economics*, 112(4):1127–1161.

Hines, J. R. and Thaler, R. H. (1995). Anomalies: The flypaper effect. *The Journal of Economic Perspectives*, 9(4):217–226.

Hudson, S., Hull, P., and Liebersohn, C. (2015). Interpreting instrumented difference-in-differences. Technical report, Working Paper (available upon request).

Hudson River Watershed Alliance (2015). Watershed rules and regulations for protection and drinking water in new york. New York State Department of Environmental Conservation.

Imazeki, J. and Reschovsky, A. (2004). Is no child left behind an un (or under) funded federal mandate? evidence from texas. *National Tax Journal*, pages 571–588.

Institute, B. T. (2010). Beijing transport annual report. `http://www.bjtrc.org.cn/JGJS.aspx?id=5.2&Menu=GZCG`. Accessed: 2019-03-2.

Iseki, H., Ford, A., and Factor, R. J. (2005). Contracting practice in fixed-route transit service: Case studies in california. *Transportation research record*, 1927(1):82–91.

Jain, P. (2018). Delhi pollution: Vehicles, not stubble burning, main cause, says teri study. *India Today*.

Kahn, M. E. (2000). Smog reduction's impact on california county growth. *Journal of Regional science*, 40(3):565–582.

Kahn, M. E. (2001). The beneficiaries of clean air act regulation. *Regulation Magazine*, 24:34.

Kasperowicz, P. (2018). GOP to attack unfunded federal mandates next week. *Washington Examiner*, July 5.

Keiser, D. A. and Shapiro, J. (2018). Consequences of the clean water act and demand for water quality. *The Quarterly Journal of Economics*.

Koch, E. I. (2011). *Mayor: an autobiography*. Simon and Schuster.

Kuminoff, N. V. and Pope, J. C. (2014). Do "capitalization effects" for public goods reveal the public's willingness to pay? *International Economic Review*, 55(4):1227–1250.

Lake, E. E. H., Oster, W. M., et al. (1979). *Who pays for clean water? the distribution of Water pollution control costsan urban systems research report*. Number 333.9114 L3.

Lave, C. (1991). Measuring the decline in transit productivity in the us. *Transportation Planning and Technology*, 15(2-4):115–124.

Lechner, M. et al. (2011). The estimation of causal effects by difference-in-difference methods. *Foundations and Trends in Econometrics*, 4(3):165–224.

Lee, D. S. (2008). Randomized experiments from non-random selection in us house elections. *Journal of Econometrics*, 142(2):675–697.

Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of economic literature*, 48(2):281–355.

Leggett, C. G. and Bockstael, N. E. (2000). Evidence of the effects of water quality on residential land prices. *Journal of Environmental Economics and Management*, 39(2):121–144.

LeRoy, S. F. and Sonstelie, J. (1983). Paradise lost and regained: Transportation innovation, income, and residential location. *Journal of Urban Economics*, 13(1):67–89.

Levin, J. and Tadelis, S. (2010). Contracting for government services: Theory and evidence from us cities. *The Journal of Industrial Economics*, 58(3):507–541.

Li, S., Kahn, M. E., and Nickelsburg, J. (2015). Public transit bus procurement: The role of energy prices, regulation and federal subsidies. *Journal of Urban Economics*, 87:57–71.

Lin, Y. (2016). The long shadow of industrial pollution: Environmental amenities and the distribution of skills.

Lipton, D. (2004). The value of improved water quality to chesapeake bay boaters. *Marine Resource Economics*, 19(2):265–270.

Lopez-de Silanes, Florencio, A. S. and Vishnyn, R. W. (1997). Privatization in the united states. *The RAND Journal of Economics*, 28(3):447–471.

Lu, X. (2016). Effectiveness of government enforcement in driving restrictions: a case in beijing, china. *Environmental Economics and Policy Studies*, 18(1):63–92.

Lu, X. (2018). Housing markets and automobile policy. Manuscript.

Lutz, B., Molloy, R., and Shan, H. (2011). The housing crisis and state and local government tax revenue: Five channels. *Regional Science and Urban Economics*, 41(4):306–319.

Lutz, B. F. (2008). The connection between house price appreciation and property tax revenues. *National Tax Journal*, pages 555–572.

McCarty, N., Poole, K. T., and Rosenthal, H. (2016). *Polarized America: The dance of ideology and unequal riches*. mit Press.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of econometrics*, 142(2):698–714.

McQuillin, E. (1912). *A Treatise on the law of municipal corporations*, volume 4. Callaghan & Company.

Melnik, W. (2017). Municipal government reaction to mass layoffs in ohio. Working Paper.

Melosi, M. V. (2000). *Sanitary City*. the University of Pittsburgh Press.

Metcalf, L. and Eddy, H. P. (1922). *Sewerage and Sewage Disposal: A Textbook*. McGraw Hill.

Miceli, T. J. and Segerson, K. (1999). Threshold rules for funding environmental mandates: accountability and the unfunded mandate reform act. *Land Economics*, pages 375–389.

Miller, G. (2010). The arbitration game. *Governing*, 8.

Mills, E. S. (1967). An aggregative model of resource allocation in a metropolitan area. *The American Economic Review*, 57(2):197–210.

Minnesota Pollution Control Agency (2009). Low dissolved oxygen in water: Causes, impact on aquatic life-an overview. Technical report, St. Paul, MN.

Missouri (1901). *Missouri v. Illinois and Sanitary District of Chicago, 180 U.S. 208*.

Mohring, H. (1972). Optimization and scale economies in urban bus transportation. *The American Economic Review*, 62(4):591–604.

Muth, R. F. (1967). The distribution of population within urban areas. In *Determinants of Investment Behavior*, pages 271–299. NBER.

Newman, B. (1976). Mud and debt: Village is all fed up with expert planners and federal strings. *The Wall Street Journal*.

NLC (2017). Energy, environment and natural resources 2018. Technical report, National League of Cities.

Oates, W. E. (1969). The effects of property taxes and local public spending on property values: An empirical study of tax capitalization and the tiebout hypothesis. *Journal of political economy*, 77(6):957–971.

Okun, D. A. (1996). From cholera to cancer to cryptosporidiosis. *Journal of Environmental engineering*, 122(6):453–458.

Olmstead, S. M. (2010). The economics of water quality. *Review of Environmental Economics and Policy*, 4(1):44–62.

Olmstead, S. M. and Kuwayama, Y. (2015). Water quality and economics: willingness to pay, efficiency, cost-effectiveness, and new research frontiers. *Handbook on the Economics of Natural Resources*, page 474.

Oum, T. H., Waters, W. G., and Yong, J.-S. (1992). Concepts of price elasticities of transport demand and recent empirical estimates: an interpretative survey. *Journal of Transport Economics and policy*, pages 139–154.

Poor, P. J., Pessagno, K. L., and Paul, R. W. (2007). Exploring the hedonic value of ambient water quality: A local watershed-based study. *Ecological Economics*, 60(4):797–806.

PriceWaterhouse (1994). Impact of unfunded federal mandates on us cities. *Second Printing*, page 4.

Rappaport, J. and Sachs, J. D. (2003). The united states as a coastal nation. *Journal of Economic growth*, 8(1):5–46.

Reback, R., Rockoff, J., and Schwartz, H. L. (2014). Under pressure: Job security, resource allocation, and productivity in schools under no child left behind. *American Economic Journal: Economic Policy*, 6(3):207–41.

Rechtschaffen, C. (2003). Enforcing the clean water act in the twenty-first century: Harnessing the power of the public spotlight. *Ala. L. Rev.*, 55:775.

Richmond, J. (2001). *The private provision of public transport*.

Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy*, 82(1):34–55.

Samuelson, P. A. (1954). The pure theory of public expenditure. *The review of economics and statistics*, pages 387–389.

Seward, J. (2017). Senate passes comprehensive mandate relief legislation. The New York State Senate.

Shleifer, A. (1998). State versus private ownership. *Journal of economic perspectives*, 12(4):133–150.

Sieg, H., Smith, V. K., Banzhaf, H. S., and Walsh, R. (2004). Estimating the general equilibrium benefits of large changes in spatially delineated public goods. *International Economic Review*, 45(4):1047–1077.

Skidmore, M. and Scorsone, E. (2011). Causes and consequences of fiscal stress in michigan cities. *Regional Science and Urban Economics*, 41(4):360–371.

Small, K. A. and Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica (pre-1986)*, 49(1):105.

Stock, J. and Yogo, M. (2005). *Asymptotic distributions of instrumental variables statistics with many instruments*, volume 6. Chapter.

Stoddard, A., Harcum, J. B., Simpson, J. T., Pagenkopf, J. R., and Bastian, R. K. (2003). *Municipal wastewater treatment: evaluating improvements in national water quality*. John Wiley & Sons.

Sun, W., Zheng, S., Geltner, D. M., and Wang, R. (2017). The housing market effects of local home purchase restrictions: evidence from beijing. *The Journal of Real Estate Finance and Economics*, 55(3):288–312.

Tang, M. and Coulson, N. E. (2017). The impact of china's housing provident fund on homeownership, housing consumption and housing investment. *Regional Science and Urban Economics*, 63:25–37.

Tarr, J. (2016). Infrastructure: Water and sewers. In *Oxford Research Encyclopedia of American History*.

TCRP (2004). Transit pricing and fares. Accessed: 2016-06-06.

Thomas, H. (1998). The effect of state policies on the location of manufacturing: Evidence from state borders. *Journal of Political Economy*, 106(4):667–705.

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of political economy*, 64(5):416–424.

Train (1975). *Train versus City of New York, 420 U.S. 35*.

US Census Bureau (2015). Historical finances of individual governments, 1967-2012. Technical report, US Department of Commerce, Washington, DC.

Viard, V. B. and Fu, S. (2015). The effect of beijing's driving restrictions on pollution and economic activity. *Journal of Public Economics*, 125:98–115.

Walsh, P. J., Milon, J. W., and Scrogin, D. O. (2011). The spatial extent of water quality benefits in urban housing markets. *Land Economics*, 87(4):628–644.

Walters, G. (2017). Personal interview.

Waxman, A. (2017). The long road to work: The divergent effects of transportation policies by worker skill in a locational sorting model. *working paper*.

Weiland, P. S. (1998). Environmental regulations and local government institutional capacity. *Public Administration Quarterly*, pages 176–203.

Westerling, K. e. (2011). Owner and operator of wastewater treatment facilities sentenced for violating the clean water act. *Water Online*.

Winston, C. (2000). Government failure in urban transportation. *Fiscal Studies*, 21(4):403–425.

Winston, C. (2010). *Last exit: Privatization and deregulation of the US transportation system*. Brookings Institution Press.

Wolff, H. (2014). Keep your clunker in the suburb: Low-emission zones and adoption of green vehicles. *The Economic Journal*, 124(578):F481–F512.

World Health Organization (2019). Public health, environmental and social determinants of health. `https://www.who.int/phe/health_topics/outdoorair/databases/cities/en/`. Accessed: 2019-03-31.

Xu, Y., Zhang, Q., and Zheng, S. (2015). The rising demand for subway after private driving restriction: Evidence from beijing's housing market. *Regional Science and Urban Economics*, 54:28–37.

Yang, J., Liu, Y., Qin, P., and Liu, A. a. (2014). A review of beijing's vehicle registration lottery: Short-term effects on vehicle growth and fuel consumption. *Energy Policy*, 75:157–166.

Zhang, W., Lawell, C.-Y. C. L., and Umanskaya, V. I. (2017). The effects of license plate-based driving restrictions on air quality: Theory and empirical evidence. *Journal of Environmental Economics and Management*, 82:181–220.