Essays on Real Estate and Urban Economics

A dissertation submitted to the
Graduate School of the University of Cincinnati
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Ph.D.)

in the Department of Economics
of the Carl H. Lindner College of Business

by

Saannidhya Rawat

Dissertation Committee Members:
David Brasington (Chair)
Gary Painter
Eunjee Kwon
Olivier Parent

2026

Abstract

Cities and neighborhoods depend on sustained investment in roads, housing, and local public goods, but the incentives to maintain and improve the built environment differ across public institutions, ownership structures, and governance regimes. This dissertation studies how institutions and incentives governing local investment shape roads, housing reinvestment, and urban public goods, and how those investment decisions affect property values, housing conditions, and neighborhood trajectories. The first chapter examines public disinvestment in local infrastructure by studying close elections over road maintenance tax renewals in Ohio. Combining election returns, local government finances, housing transactions, and AI-based measures of road quality from satellite imagery, it uses a regression discontinuity design to identify the effects of failed renewals. It finds that when a locality cuts renewal road levies, road maintenance funds fall by 11%, road quality declines by 23%, and house prices fall by about 9%, showing that even maintenance of existing infrastructure is capitalized into local real estate values. These price effects emerge with a delay, consistent with gradual deterioration, and are strongest in urban areas, emphasizing that local fiscal institutions affect housing markets through the long-run quality of neighborhood infrastructure. The second chapter turns from public to private investment in the housing stock. Using parcel-level ownership records and building-permit data for single-family homes in Minneapolis and Charlotte, it measures reinvestment directly rather than inferring it from resale prices. It finds that owner-occupiers in Minneapolis file 26% more permits and undertake 55% more permitted work than otherwise comparable rentals, while larger landlords file about 40% fewer permits and invest about 45% less than smaller landlords. Evidence from Charlotte points in the same direction. Together, the first two chapters show that both public finance and private ownership determine whether the U.S. built environment is maintained or allowed to depreciate.

The third chapter studies formal urban governance and public investment in rapidly urbanizing settlements in India. Using a fuzzy multi-threshold regression discontinuity design around Census Town eligibility, it exploits quasi-random variation in the probability that a settlement receives statutory urban status. Meeting the Census Town thresholds increases statutory recognition by 7.1 percentage points, and the resulting transition to urban local body status expands schools, hospitals, and financial access while also reducing sports facilities, consistent with land reallocation in denser and formalizing settlements. The results indicate that settlements near the urban threshold gain not only administrative recognition but also materially different access to local services once they enter the formal urban governance system. Taken together, the three essays show that the quality and trajectory of places depend on the rules and incentives that govern who invests in them: local tax institutions shape the maintenance of public infrastructure, ownership structure shapes reinvestment in the housing stock, and municipal status shapes the provision of urban public goods. Across settings, the dissertation highlights that underinvestment need not arise from physical scarcity alone; it often follows from the institutional rules that define who bears costs, who captures benefits, and who has authority to act. The dissertation therefore contributes to both real estate and urban economics by linking property values, housing maintenance, and local public goods within a common framework centered on institutions, incentives, and investment in the built environment. More broadly, the findings imply that policies governing local public finance, the ownership of residential property, and the timing of municipalization can have durable effects on the physical condition of neighborhoods and on the opportunities available to households living in them.

Acknowledgements

I would like to thank David Brasington, my dissertation committee chair, for his unwavering support throughout my doctoral studies. Thank you for all the lunches, meetings, and patient guidance through countless data problems. I should also thank Riley, Daniel, and Casey for letting me borrow their dad from time to time.

I am equally thankful to the rest of my committee. Gary Painter generously shared his guidance, knowledge, and hard-earned experience from the real estate side of the world. Eunjee Kwon believed in me from day one. Olivier Parent sent many opportunities my way and has been a steady source of support for both me and the Economics Department.

I am also grateful to other faculty members and post-docs whose teaching, conversations, and examples shaped my time at Cincinnati: Jeff Mills, Hernan Moscoso Boedo, Debashis Pal, Michael Jones, Na Young Lee, Rene Saran, Iryna Topolyan, Erwin Erhardt, Lenisa Chang, Asawari Deshmukh, Dongchen Zhao, Clemens Pilgram, and Maeve Maloney.

My thanks also go to the PhD students across Lindner for making the long years of graduate school far less lonely. Thank you to Vikram Krishnaveti, Ian Bryant, Anthony Chance, Nathan Hudson, Sara Katanchian, Max Richards, and Sachin Sisodiya. My third floor warriors, Yujin Lee, Chad Dulle, Albert Choi, Emma Neybert, Sodiq Babatunde, Ben Fagan, Sharmeen Merchant, Alberto Barchetti, Hyerin Han, Amin Aminimehr, Adam Arveson, Adriana Gonzalez Sanchez, Jiawei Huang, Amy Koomson, Jack Luu, Angel Vila, Ruthairut Wootisarn, Chia-Chun Yang, and Seungjoo Yoon, for their camaraderie and solidarity.

Most importantly, I thank my brother, who lived with me and saw me struggle through this PhD every single day, and my partner, Shreeya, who has been my rock and source of strength. Finally, I would like to acknowledge my parents and my grandparents for all of their sacrifices. Words are not enough to thank them, so I will not try.

Contents

List of Tables

List of Figures

Chapter 1

The Effect of Local Road Maintenance Tax Cuts on House Values

1.1 Introduction

Roads are an important form of infrastructure investment that affect households' amenity values, firms' production functions, and people's commuting costs. A great deal of existing research has focused on the effects of new roads[1], especially in developing nations, providing valuable policy insights and spurring development initiatives like China's Belt and Road Initiative (Huang, 2016) and India's Pradhan Mantri Gram Sadak Yojana (Asher and Novosad, 2020). However, the literature is constrained by two fundamental challenges. The first is endogenous selection as infrastructure investments are not exogenous events but are deliberate policy choices. New road placement is selected by policymakers and is endogenous to anticipated economic growth, while repavement schedules and decisions are systematically correlated with confounding local factors such as political cycles or unobserved community characteristics that independently drive economic outcomes (Duranton and Turner, 2011; Gonzalez-Navarro et al., 2023). This identification problem is exacerbated by the second challenge: measurement. Most studies rely on road quality metrics that are inherently subjective, labor-intensive, and generated at low frequency (Attoh-Okine and Adarkwa, 2013). Such data are not only prone to measurement error that produces biased estimates, but their lack of scalability limits the geographic and temporal scope of empirical analysis (List, 2022).

In this paper, we avoid endogeneity issues by studying property taxes that fund the maintenance of existing roads and by utilizing a quasi-experimental setting that arises from analyzing close elections. We establish a novel dataset that allows us to study how changes in local taxes affect local infrastructure maintenance and neighborhood house prices. We match voting data on local road maintenance taxes from the Ohio Secretary of State with financial records of cities from Ohio Auditor of State and property sale prices from CoreLogic.

[1] See Ghani, Goswami and Kerr (2016), Beenstock, Feldman and Felsenstein (2016), Levkovich, Rouwendal and Van Marwijk (2016), Hoogendoorn et al. (2019), Theisen and Emblem (2021), Gibbons et al. (2019), Fretz, Parchet and Robert-Nicoud (2022). Some papers have also observed adverse effects of new roads such as state-led repression (González et al., 2025) and environmental degradation (Asher, Garg and Novosad, 2020).

Our data also includes remote sensing data of local roads in Ohio from Google Earth. Using these rich data, we first estimate the expected loss in the road maintenance budget for a local government after a tax cut from a failed renewal of road maintenance levy. We fine-tune an Artificial Intelligence (AI) vision model using more than 53,000 satellite images from Brewer et al. (2021) to predict changes in road quality after a road maintenance tax cut. Finally, we use the Regression Discontinuity (RD) estimator to estimate the effect of cutting local road maintenance taxes on housing prices.

Our empirical findings reveal that when a city cuts its renewal road maintenance taxes, it faces an 11% loss in maintenance funds, its road quality declines by 23% and its house prices decrease by around $15,350 (9%). After establishing the main results, we explore heterogeneity in the effects. We find that the effects on house prices are driven by urban areas rather than rural areas, with urban areas experiencing a 7.6% decline in house prices while rural areas experience inconsistent changes in house prices. Moreover, we find an effect on the intensive margin with evidence for dosage-response based on the size of the levy. We also observe that higher-priced homes experience larger reductions in value relative to lower-priced homes. This suggests that wealthier households are more sensitive to road quality, as a poor road in front of a high-priced home is more likely to be noticed than a poor road in front of a low-priced home. We rationalize our main findings using a Dynamic General Equilibrium (DGE) framework and show how public disinvestment negatively affects property values, highlighting the trade-off between short-term tax savings and long-term private capital losses. Following Hendren and Sprung-Keyser (2020), we compute a conservative estimate of the Marginal Value of Public Funds (MVPF) using our empirical estimates and find a 1.4 net welfare loss to households for every dollar of revenue loss to the government. Our focus on road maintenance instead of new development, our approach of leveraging close elections as a quasi-experiment to study the effects of local road maintenance tax cuts, and our use of an AI vision model to measure changes in road quality distinguish our study from the existing literature.

These results are consistent with other work on the impact of roads on house prices. For instance, Gonzalez-Navarro and Quintana-Domeque (2016) finds that paving roads in Acayucan, Mexico, increases property values by 28%, while Theisen and Emblem (2021) finds 13% higher housing prices in towns nearest to a new highway in Norway. Unlike Beenstock, Feldman and Felsenstein (2016) and Diao, Leonard and Sing (2017), we do not find any anticipation effects but we do find statistically significant effects starting in the fourth year after the vote and continuing at least through the ninth year. This delayed effect highlights how the consequences of road tax cuts are not immediate but gradually accumulate as it takes time for roads to deteriorate and for the decline in road quality to be capitalized into house prices.

Roadmap. The rest of the paper is organized as follows. Section 1.2 reviews the related literature and positions our contribution. Section 1.3 provides background information, establishes the road quality metrics and presents the variables used in the study. Section 1.4 outlines the empirical strategy. Section 1.5 presents the results of the study and shares the relevant robustness checks. Section 1.6 examines the mechanisms underlying our findings[2]. Section 1.7 implements our MVPF calculation to assess the impact of this tax cut policy. Section 1.8 concludes.

1.2 Literature Review

Roads and House Prices. Gonzalez-Navarro and Quintana-Domeque (2016) studies the effects of paving roads in Acayucan, Mexico. The government identified 56 neighborhoods that needed paved roads. Of these, 28 were randomly selected for paving. About 1,000 people were surveyed in these neighborhoods before paving (2006) and after (2009), although

[2] In this section, we highlight the main findings from the DGE framework and relegate the model details to Appendix A1.4.

11 of the 28 treatment groups were still in the process of being paved at the time of the post-treatment survey. It finds a 17% increase in property values as measured by professional appraisals, a 28% increase in homeowner-estimated property values, a 36% increase in rents, and a 72% increase in vacant land values, along with effects on a few non-housing outcomes. Our study evaluates cuts in road maintenance rather than a switch from unpaved to paved roads, and our Ohio geography contrasts with that of a developing nation. Our housing values are not based on professional appraisals or homeowner valuation but on actual sales transactions; and the fact that we observe sales transactions every year lets us analyze pre- and post-treatment trends in house prices over a long timeframe rather than a one-time change in house values.

We find almost no research on road maintenance. The exceptions include theoretical work by Rioja (2003), which develops a dynamic general equilibrium model to study the optimal amount of road maintenance, which is found to depend on the size of new infrastructure investments and the productivity of infrastructure. Rioja (2003) solves and parameterizes the model, finding that increasing funding away from new infrastructure and toward the maintenance of existing infrastructure decreases the depreciation rate of existing infrastructure, increases the stock of existing infrastructure, and increases economic output and consumption levels.

Another study, Chaurey and Le (2022), assesses the impact of infrastructure maintenance in India. As their research notes, the Indian national government identified 17 major states as the poorest and made an index of "backwardness" for districts in these states. The 115 most backward districts were awarded 450 million rupees and were given discretion on how to spend these funds to "improve and maintain" or "make complementary investments" to existing infrastructure. Money could be used to widen roads, add lanes, add electricity transmission and distribution infrastructure, build new road links to markets, and build bridges, for example. The treatment is therefore a mixture of road and non-road spending and of maintenance and new construction, especially, presumably, in districts that did not already have paved roads and electricity, where Chaurey and Le (2022) finds the largest effects. Unlike Chaurey and Le (2022), we focus on the effects of infrastructure spending on real estate values rather than employment outcomes.

The only other study we find on maintenance and house prices is the classic capitalization study by Edel and Sclar (1974). It studies the effect of local public spending on median house values in the Boston area using decennial Census years from 1930 to 1970. Controlling for latitude and longitude, the tax rate, population density, tenure status, and school expenditures, its ordinary least squares regressions do not find a statistically significant link between road maintenance spending and house prices.

Identification. Asher and Novosad (2020) studies the impact of new roads on villages in India. Like our study, its identification strategy is regression discontinuity, although ours is sharp rather than the fuzzy form. It argues the main obstacle to identification in prior studies is that the placement of new roads is usually correlated with economic (or political) characteristics rather than exogenous. Its findings suggest this is a serious problem with the literature because, unlike prior studies, it finds no strong link between economic growth and new road placement, suggesting that the estimates of previous studies that find a link are driven by road placement in villages that are already growing. Asher and Novosad (2020) touts its use of village-level rather than regional-level data. We, too, look at economic outcomes at the level of village, city, and township, the most local levels of government. A surprising finding of Asher and Novosad (2020) is that investment in transportation infrastructure does not affect village incomes, assets, or agricultural output. Its measure of assets is a village-level average of a series of binary indicators of ownership of a variety of assets, along with separate regressions for the presence of a 'solid house', refrigerator, and phone; whereas we study the effect of local road tax cuts on housing sale prices. Of course, our use of a developed geography contrasts with rural villages in India. Our efforts to achieve identification focus on the maintenance of existing roads, which avoids the endogeneity of the placement of new roads. We also find similarities with Boudot-Reddy and Butler (2024), which employs a similar identification strategy but studies new roads instead of the maintenance of existing roads and finds that villagers reward the political party responsible for spearheading the road expansion.

Another study with an identification strategy similar to ours is Cellini, Ferreira and Rothstein (2010), which studies the effect of new capital projects for schools, funded via new local bond issues and raised by referendums. Both our study and Cellini, Ferreira and Rothstein (2010) employ a dynamic regression discontinuity design and analyze changes in regional property values. Moreover, both papers rely on broadly similar identification strategies and assignment mechanisms: in each case, votes in favor of or against a local referendum serves as the source of exogenous variation. Despite these similarities, a few key distinctions stand out. First, Cellini, Ferreira and Rothstein (2010) looks at the impact of bond elections for additional funding on new capital infrastructure projects, while we focus on the impact of cutting renewal referendums for existing road maintenance. Second, Cellini, Ferreira and Rothstein (2010) implements a fuzzy regression discontinuity approach, whereas ours is sharp. Lastly, we focus on the Intent-to-Treat (ITT) estimator, which analyzes the effect of an election on outcomes without controlling for the results of subsequent elections[3] (more details in Section 1.4). As Cellini, Ferreira and Rothstein (2010) shows, when elections are independent, the ITT estimator is equal to the Treatment on the Treated (TOT) estimator.

Road Quality. Measuring road quality remains empirically challenging. Currier, Glaeser and Kreindler (2023) exploits vertical-acceleration signals from millions of Uber trips to construct a Road Roughness Index (RRI) and shows that it moves predictably with resurfacing events. Although RRI is highly correlated with engineering benchmarks such as the International Roughness Index (IRI) and the Pavement Condition Index (PCI), each of these

[3]Recent papers such as Hsu and Shen (2024) and Biasi, Lafortune and Schönholzer (2025) have focused on improving upon methods from Cellini, Ferreira and Rothstein (2010). However, these papers focus on the effects of new capital projects rather than maintaining existing capital infrastructure. Given the exogeneity of the timing of renewal elections and our focus on existing maintenance, the ITT estimator is more appropriate for our setting.

legacy measures has drawbacks that limit their use for large-scale economic analysis: IRI data are collected only where a costly profilometer van has recently driven, and PCI ratings depend on labour-intensive, inspector-specific visual surveys that are updated irregularly. Moreover, the Uber-based approach cannot cover streets outside the platform's service footprint—precisely where many suburban and rural housing markets are located. In contrast, high-resolution satellite imagery is available nationwide and available at low marginal cost. By fine-tuning a vision model on a labelled road-surface dataset, we generate a consistent, wall-to-wall measure of road quality that spans urban, suburban and rural networks alike, and captures the visual cues that prospective home-buyers directly observe.

We find another study, Wong et al. (2017), which examines the effect of local political reform in rural Chinese villages on local infrastructure projects and finds that these reforms result in younger village leaders and higher quality roads. It measures road quality by creating a village-level road quality rating, which is based on the observations of field surveyors that use their "road quality evaluation scheme". They emphasize minimization of subjectivity in the evaluation process, as the surveyors are trained to use a standardized set of criteria to assess road quality, such as number of bends per 100 meters and road width. Unlike Wong et al. (2017), which relies on surveyor-based evaluations, our road quality measure is derived from high-resolution satellite imagery processed through our fine-tuned vision model, ensuring consistency and scalability across diverse geographic regions.

Transportation and House Prices. We highlight a substantial literature studying the effect of transportation infrastructure on house prices. Hoogendoorn et al. (2019) studies the effect of the opening of a tunnel on house prices in the Netherlands, noting that prior research on transportation infrastructure in developed regions often suffers from reverse causality. It argues that the opening of the Westerscheldtunnel is a fairly exogenous event, with natural borders that prevent contamination of results by the surrounding environment. Hoogendoorn et al. (2019) finds that half the capitalized value of the tunnel occurs more than a year before the tunnel opens and argues that the exogeneity of the tunnel's opening, along with hedonic controls, time trends, and postcode fixed effects, identifies its estimates. Our data also pertains to a developed nation. One novelty of our study is how ordinary the events are that we study. While the opening of a new tunnel is significant, it is rare. Votes to renew infrastructure spending are common events in many local governments in the United States, and the amount of road maintenance spending is regularly determined by governments worldwide, either through voting or directly by bureaucrats. It is therefore important to study the effects of road maintenance spending on house prices.

Li et al. (2016) studies the overall effect on apartment prices of new subway lines in Beijing, but the estimates may represent the net effect of competing factors. Gibbons and Machin (2005), studying the construction of new rail stations for the London underground and light rail services, notes that the effect on house prices captures the net effect of better access, increased crime, and increased noise pollution. Levkovich, Rouwendal and Van Marwijk (2016) looks at the effect of highway development on house prices in the Netherlands. It separates out accessibility effects from noise pollution and increased traffic effects by looking at different neighborhoods near the highway development. Its repeat sales difference in differences model finds increased house prices from anticipation effects (Kohlhase, 1991). Beenstock, Feldman and Felsenstein (2016) also finds anticipation effects for house prices for the development of a highway across Israel.

Contribution. Our paper contributes to the literature on three main fronts. First, we contribute to the literature on the effects of public infrastructure on house prices. Our quasi-experimental design complements recent evidence from developing nations, such as Mexico (Gonzalez-Navarro and Quintana-Domeque, 2016) and Indonesia (Gertler et al., 2024). By exploiting close elections for renewal referendums earmarked to fund road maintenance, we mitigate the endogeneity concerns common in the literature and estimate the capitalization of road maintenance into property values. Second, we compute a novel measure of road quality using a fine-tuned vision model trained on satellite images. Current literature has relied on measures based on Pavement Condition Index (PCI) or International Roughness

Index (IRI) to measure road quality, but these are not available for many localities. Similarly, Currier, Glaeser and Kreindler (2023) measures road roughness using vertical acceleration data from a smartphone app, but this also requires large-scale data available mainly for large metropolitan areas and is not available for smaller towns. Our approach allows us to measure road quality at the local level, bypassing the limitations of existing measures. Third, by assessing the long-term capitalization of road maintenance, we uncover a dynamic trade-off between immediate tax savings and deferred infrastructure quality, highlighting the role of local government investment in public goods for sustaining property values (Oates, 1972; Rioja, 2003). While *immediate* tax cuts may initially increase house prices by lowering ownership costs and making properties more attractive to buyers, our results suggest that these savings are eventually eclipsed by the *future* decline in local public amenities.

1.3 Background & Data

1.3.1 How are roads funded in Ohio?

Roads in Ohio are funded through a combination of federal, state, and local sources. A significant portion of road funding provided by federal and state governments comes from gas taxes, which are currently set at $0.18 per gallon for federal tax and $0.38 per gallon for Ohio state tax. Additional sources for these two levels of government include vehicle registration fees, license plate fees, tolls, and driver's license fees. Funding for local governments largely comes from property taxes. Below, we provide an overview of road funding in Ohio from national, state, and local sources.

Federal Funding. U.S. Department of Transportation is responsible for federal road infrastructure funding and allocates funds through the Federal Highway Administration (FHWA). The FHWA provides funding for the construction, maintenance, and operation of highways, bridges, and tunnels through the Highway Trust Fund (HTF), which is funded by the federal gas tax. The HTF is divided into two accounts: the Highway Account and the Mass Transit Account. The Highway Account is used to fund highway construction and maintenance, while the Mass Transit Account is used to fund public transportation projects. Under the Bipartisan Infrastructure Law (BIL) (U.S. Department of Transportation, 2022), the HTF funds key programs like the National Highway Performance Program (NHPP), which targets the National Highway System (NHS), and the Surface Transportation Block Grant (STBG) program, which provides flexibility for a broader range of federal-aid roads. However, federal funding for local road infrastructure is largely restricted to major roads, leaving most neighborhood streets dependent on local revenue. As of 2022, about $4/5^{\text{th}}$ of all road funding came from state and local governments (Peter G. Peterson Foundation, 2024).

State Funding. The Ohio Department of Transportation (ODOT) is responsible for maintaining the Interstate system, the U.S. and State Routes. In Ohio, gas taxes, licensing fees and user fees account for 69% of the state's road funding (Boesen, 2021), which can be used for both maintenance and new construction. However, about 70% of funds goes to highway construction[4], 2% is given to local governments as grants and only 4% of state funds are directed towards roadway maintenance (Ohio Department of Transportation, 2023). The remaining funds are part of payroll and operating expenses and other miscellaneous expenses. Hence, most of the road maintenance funding for local, neighborhood-specific roads in Ohio comes from local governments.

Local Funding. Local governments in Ohio fund roads mainly through property taxes, although the extent of this funding varies across different localities. These municipalities have the authority to levy property taxes specifically for road maintenance, providing a crucial source of funding for the upkeep of local infrastructure. For instance, as per our correspondence with Beavercreek Township, 61% of its road funds come from property taxes, and only 8% from gas taxes. Moreover, 77% of funding for roads in Beavercreek Township is

[4]According to (Ohio Department of Transportation, 2023), most of this category is used for preserving existing infrastructure rather than new capacity.

provided by the local government and 11% of the overall local government budget is allocated to roadway maintenance (Public Service Department of Beavercreek Township, 2025). We can see that local roads are primarily funded by local governments in Ohio and road maintenance funding is a significant part of the local government's budget.

1.3.2 Local Taxation in Ohio

Background. Ohio consists of 88 counties, each covering about 464 miles^2 (1,200 kilometers^2). Each county was historically divided into about 15 equally-sized townships, which do not cross county lines. Citizens can petition to incorporate as a village, which has a different type of government structure than a township and the ability to levy both income and property taxes, whereas townships may only levy property taxes. When a village exceeds 5,000 population in Ohio, it is reclassified as a city. Villages and cities may cross township and county lines, dissolve, or annex parts of contiguous townships. Villages, cities and townships, which we call "cities" for brevity, are the most local governmental unit in Ohio. Each local government covers about 18.2 miles^2 (47.1 km^2) on average. The Ohio Revised Code lets local governments collect a small amount of tax without a vote. Beyond this limited amount local governments put tax levies on the ballot to ask for additional money from voters. Our data on renewal road tax levies include 3,184 referendums and cover 617 local governments in Ohio.

The type of tax levy we study is for the renewal of road tax levies. Most of the renewal taxes we consider have stated purposes of "road and bridges repair", "road repair", "street fund", and "street improvements", although there are less common stated purposes like "repair and maintenance of streets and sewers system" and "resurfacing and rehabilitation of city streets." The construction of new roads and bridges, in contrast, would be funded with a tax levy for additional money, not a renewal tax; and it would likely be funded by a bond levy lasting 20 or 30 years. We eliminate from our dataset stated purposes that might suggest new road construction like a 30-year 1.9-mill new tax in Moscow Village for "permanent improvements" and 0.5-mill new tax for 20 years in Shawnee Hills Village for "general construction and road and bridges repair." Our dataset includes tax levies such as a 2-mill, 5-year renewal in Adams Township (Champaign County) in 1995; a 3-mill, 5-year renewal in Lore City Village in 2016; and a 2.5-mill, 5-year renewal in Pataskala City in 2007.

Levies that had originally passed typically expire, and the most common duration to collect a levy is five years, representing about 90% of the road tax levies in our sample. If a tax levy is renewed, taxes and funding continue. If 50% or fewer votes approve the levy, it fails. When a tax levy for additional funding fails, there is no increase in funding, but existing funding from other tax levies continues as normal. When a renewal tax levy fails, funding from that tax levy stops. 99% of the road tax levies in our sample are property taxes and 1% are income taxes.

Renewal levies. Most RD studies that use voting data to look at the impact of funding changes examine new tax levies for additional funding. Cellini, Ferreira and Rothstein (2010) observes that votes for additional tax money may not be statistically independent; a vote may be proposed until it passes. We minimize this source of endogeneity of new votes by only considering renewal votes (Brasington, 2017). While a government may choose when to put a vote for additional funding on the ballot and keep proposing the new tax until it is passed, when a vote passes, it has an expiration date. So if a road tax levy for additional funding passes in 2007 to last five years, in 2012 voters will have the chance to renew or reject the tax. The timing of the vote in 2012 is not endogenous, having been set in 2007. If voters renew it in 2012, it will be up for renewal again in 2017. The exogeneity of the timing of a renewal tax levy contrasts with the timing of new taxes for additional money as explained by Cellini, Ferreira and Rothstein (2010) because local governments can endogenously choose the timing of a new tax levy to coincide with positive economic conditions or positive sentiment toward government action, making an election more likely to pass.

Spending impact of failing to renew a levy. When a renewal tax levy fails, the local government loses the funding from that tax levy. The local government may still have tax levies for other purposes in effect, but the road tax levy that failed to renew is no longer collected. We determine the dollar amount in consideration when a household is making a decision on whether to vote for or against a renewal road tax levy, and we call it Average Road Tax per household, which is the tax amount that a household will pay if the levy is successfully renewed. Average Road Tax per household is computed following Equation 1:

$$\text{Average Road Tax per household}_{it} = \text{millage rate}_{it} \times \text{Average Assessed value}_{it} \tag{1}$$

where i is the city, t is the year, and millage-rate [5] is the predetermined millage amount set for the road tax levy. The Average Assessed value is 35% of the average appraisal value[6], which for our study equals the average sale price of houses in the city for that year. From this, we can also compute the average road tax per city by multiplying the average road tax per household with the number of households in the city.

Although we acknowledge a large variation in the appraisal value across cities and in millage rates across referendums, we find that the average road tax per household is $76, and the average road tax per city is $167,011 as shown in Table 1. We do not observe any significant difference in the average road tax per household between cities that renew and cut road tax levies, which suggests the levies up for renewal are not systematically different between areas that renew and areas that fail to renew them. Whenever a renewal tax levy fails, the local government loses the funding from that tax levy, which is the average road tax per city. This loss in funding directly impacts local government's spending budget needed to

[5]Millage-rate is property tax rate expressed in mills (tax per $1,000 of assessed value).

[6]Assessment ratio of 35%, as set by Ohio Department of Taxation (2020).

maintain local roads in Ohio and accounts for 11% of a local government's budget for road maintenance [7].

Table 1. Spending impact of failing to renew a Road Tax Levy

	Aggregate	**Renewed**	**Cut**
Panel A: road tax per household			
Mean	76	75	79
	(55)	(53)	(62)
Panel B: road tax per city			
Mean	167,011	167,648	163,547
	(340,268)	(340,628)	(338,671)

Notes: This table presents descriptive statistics for two measures of the money collected via road tax levies. **Panel A** reports the mean and standard deviation (SD) of the road tax per household per year. **Panel B** reports the mean and SD of the total road tax collected per city. "Aggregate" denotes the full sample, while "Renewed" and "Cut" refer to levies that were renewed and failed to renew respectively. All monetary values are in constant 2010 U.S. dollars and rounded to the nearest integer.

1.3.3 Running Variable

The running variable plays a critical role in RD, which in this study represents the proportion of votes against the renewal of a road tax levy. A vote share of 50% or more against the renewal road tax levy means the levy fails and the tax will no longer be collected, resulting in

[7] For townships within our effective bandwidth, we use the average expenditure by the public works department as our base, which is \$1,528,404 as per their audited financial reports. All local roads in such municipalities are funded via this department. Average loss in maintenance spending for areas that cut their renewal levy is \$163,547, thus giving us $11\% = \frac{163{,}547}{1{,}528{,}404} \times 100$

a stoppage of road funding via that particular tax levy. There are 3,184 referendum results in our sample, 83% of which renew the tax, and 17% of which cut taxes and road maintenance. The Great Financial Crisis falls in the middle of our dataset, so readers might wonder if voting behavior was affected, but we find vote share the same to two decimal points during and outside the years 2008-2009. Our key identification assumption is that the election results are not predetermined and vote share is not precisely manipulated to fall just above or below the cutoff. This assumption allows us to exploit the randomization around the cutoff and provides the variation needed to identify the causal effect of cutting road tax. We test this assumption using a density test detailed below and covariate balance tests (see Appendix A1.3).

Density test. A classic RD assumption states that agents cannot precisely manipulate the running variable to fall just above or below the cutoff. In our context, it means that the election results are not determined prior to when the ballot takes place. In other words, no individuals, organizations, higher levels of government, foreign governments, or the firm that programs the voting machines are dictating the precise vote share for the renewal road tax levy referendums raised by a city. The standard way to test this assumption is to perform a density test like that of Cattaneo, Jansson and Ma (2020), which is based on the idea that manipulation of elections might cause a clustering of votes just to one side of the cutoff, with a pronounced drop-off on the other side of the cutoff. The p-value of this density test is 0.98. A histogram of vote share is shown in Figure 1 that graphically illustrates the lack of abrupt change in density.

Although Table 3 shows covariate balance between sets of cities that pass and fail to renew road tax levies, covariate values could still jump from one side of the cutoff to another. A drop in education levels, for example, could cause a drop in house prices that might coincide with a change in treatment, so that what might look like a treatment effect from cutting taxes and spending might in fact be caused by lower education levels. Graphs that suggest covariate smoothness around the cutoff are found in Appendix A1.3. A formal way to assess covariate

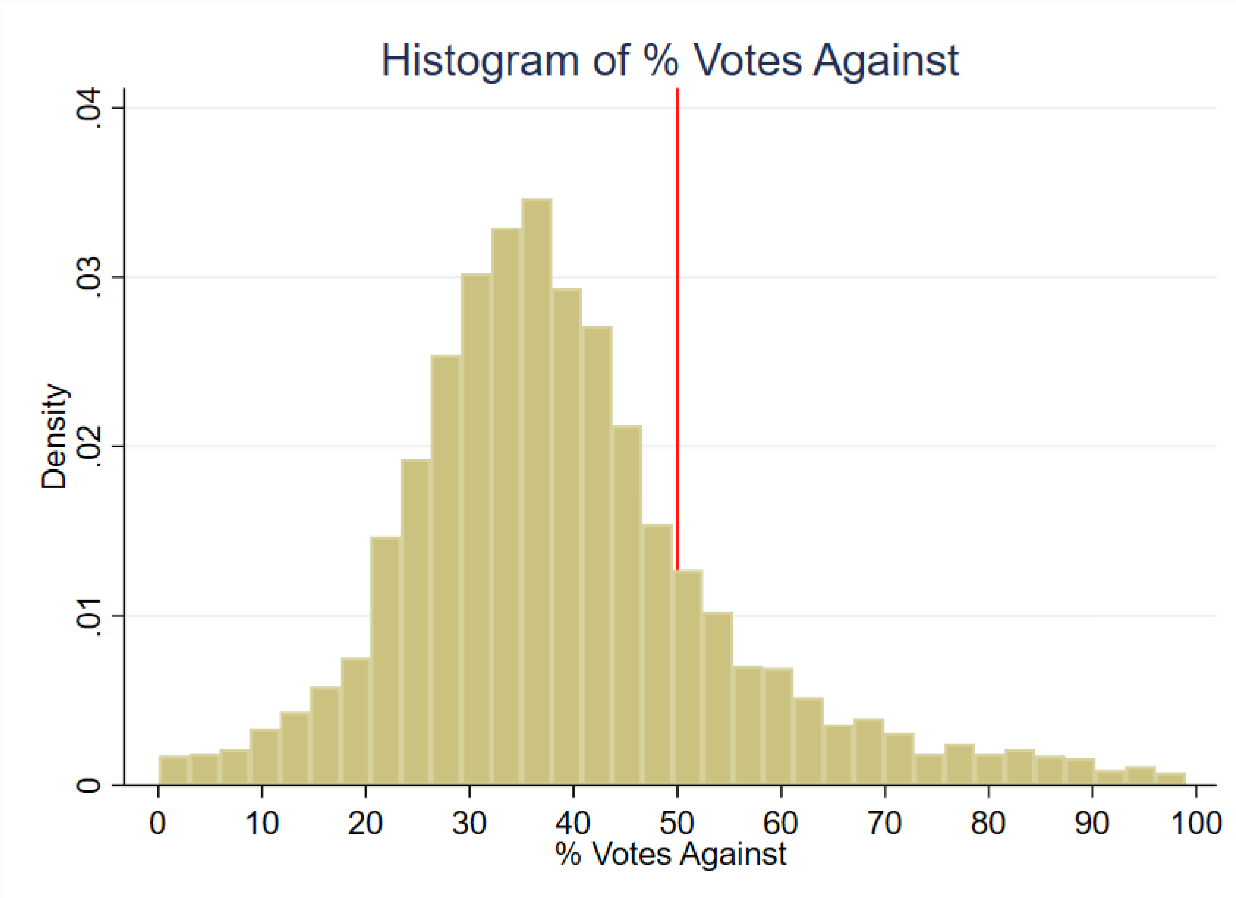

Figure 1. Histogram of Running Variable

smoothness is to use each covariate as an outcome variable in a regression of the running variable and the treatment effect dummy. When we do so, the p-value of the treatment effect dummy varies from 0.234 to 0.981, as shown in Table 24, indicating no statistically significant jump in covariate values.

1.3.4 Tax Cuts & Evidence of Road Quality

Roads in Ohio last 15-20 years and deteriorate faster with poor drainage, utility cuts to the road, and snowplow damage (City of Hudson, 2020). To understand the effect of cutting local road taxes on road quality, we fine-tune an AI vision model on satellite imagery data from Brewer et al. (2021). Following the seminal work on text-based Transformers in Natural

Language Processing (NLP) by Vaswani (2017), vision transformers (ViTs) were introduced by Google Brain's team and are part of a recent class of deep learning models that have shown to outperform classic Convolutional Neural Networks (CNNs) on image classification tasks (Dosovitskiy, 2020). Nevertheless, a more recent class of vision models learns from ViTs and uses CNNs to achieve high accuracy on image classification tasks (Woo et al., 2023). We use one such model to classify road quality from satellite images[8]. Below, we outline the steps we take to assess road quality.

Satellite Images for fine-tuning. Fine-tuning a vision model involves taking a pretrained model and adapting it for a specific image-classification task. In order to ensure appropriate fine-tuning and enable our vision model to accurately predict road quality, we need large-scale road-image data. We use the road-image dataset of Brewer et al. (2021) which consists of 53,677 labeled satellite images of roads in different conditions, and a classification representing quality of each road: 0 (poor), 1 (decent) and 2 (high). We divide the data into training, validation and test datasets, and use 70% of the data for training, 15% for validation and 15% for testing. For a detailed analysis of the road-image dataset, see the Appendix, Section A1.5. In section 1.5.1, we present the results from our fine-tuned model, when applied to roads in areas with close elections within the effective bandwidth.

Ohio Satellite Images. While the fine-tuning of our vision image-classification model uses images from Brewer et al. (2021), our testing sample uses hand-collected satellite images from Google Earth Pro for roads in different cities in Ohio. Even though Google Earth Pro covers all of Ohio and provides satellite images of roads for cities we study, these images may differ in resolution and quality depending on the location. The images are also not always available for all time periods, and the time period of the images may not be consistent across different locations[9]. For our analysis, it is imperative to ensure we have images before and after the referendum for both the treatment and control groups. However, we also

[8]see the Appendix, Section A1.5 for details.

[9]For example, Jersey Township in Licking County, Ohio had a road tax levy **renewal** in 2019, but the most temporally adjacent satellite images available are from 2017 and 2023. Marion Township in Hocking County had a road tax levy **cut** in 2018, and the proximate satellite images available are from 2014 and 2021.

found that some areas did not have images available for both periods, or the images were of poor quality. We restrict our sample to include only those areas where we could obtain high-resolution images for both periods[10]. We collect road images for cities with votes within the average effective RD bandwidth provided in Table 5 and ensure that we have pre- and post-referendum images for both- the treatment and control groups[11]. Figure 2 presents two examples of road satellite images for Ohio, one from a road classified as high and the other as poor quality.

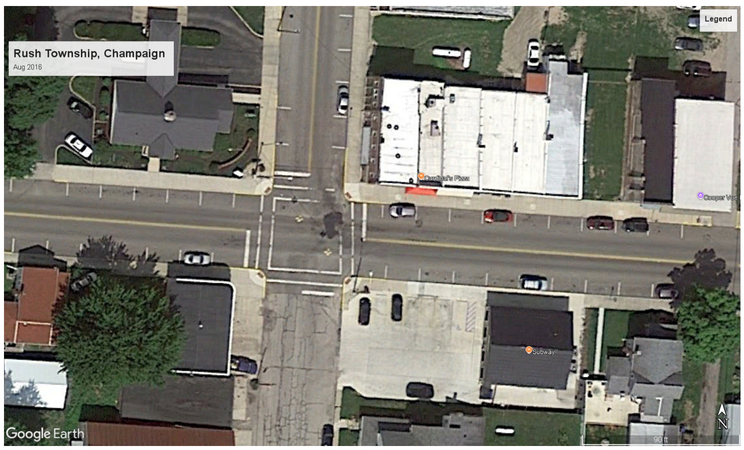

(A) A **Poor** Quality Road

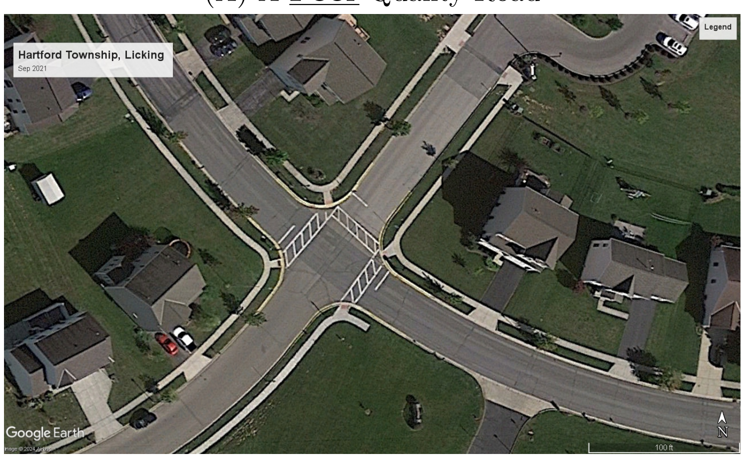

(B) A **High** Quality Road

Figure 2. Road Quality Satellite images for Ohio

[10]On average, pre-referendum images are from 2.3 years before the election, and post-referendum images are from 3.3 years after the election.

[11]Other papers such as Gonzalez-Navarro and Quintana-Domeque (2016) also have one pre- and post-intervention year, collected via surveys, rather than a comprehensive panel.

Road Quality Rating. Our fine-tuned vision model classifies each satellite image into one of three categories: 0 (poor), 1 (average), or 2 (high). This rating is based on visual cues such as pavement quality, visible cracks, patchwork, and overall surface condition, as learned by the vision model from the labeled training data. For each image, the model outputs both the predicted class and a confidence score, which reflects the model's certainty in its prediction. These ratings provide a standardized, objective measure of road quality across locations and time periods, enabling us to compare changes in road conditions before and after tax levy referendums for both treated and control groups.

Road Quality Score. Even though our original model output is a road quality rating[12], we also convert the discrete model output to a continuous measure of road quality between 1 and 100. This conversion uses two variables: the predicted road quality rating and confidence score[13], and allows us to map our road quality rating to a number similar to Pavement Condition Rating (PCR) used by ODOT, and call this Road Quality Score (RQS). The formula for this conversion is as follows:

$$\mathrm{RQS} = 1 + [(\hat{r} + \tau \cdot \hat{c}) \cdot w \cdot \mathbb{I}\{\hat{r} \neq 0\} + (1 - \tau) \cdot \hat{c} \cdot w \cdot \mathbb{I}\{\hat{r} = 0\}] \tag{2}$$

where $\hat{r}$ is the predicted road quality rating, $\hat{c}$ is the confidence score, w is a relative weight for allotted size of each road quality rating group, 0, 1, and 2. τ is a parameter for confidence weights[14]. The RQS ranges from 1 to 100, where 1 indicates poor quality and 100 indicates high quality. We use this RQS in our analysis to assess the impact of cutting local road taxes on road quality.

Next, we share a summary of means and standard deviations of the predicted road quality rating and road quality score for the treatment and control groups, before and after the referendum.

[12] This is a direct result of our training data from Brewer et al. (2021) which uses the discrete 0 (low), 1 (medium), 2 (high) road-quality metric.

[13] During the prediction process, we asked the fine-tuned vision model its confidence in the prediction i.e. the probability of the prediction being correct, a number between 0 and 1.

[14] We set this equal to 0.4.

Table 2. Predicted Road Quality by Treatment Status and Period

	Treated (Failed Levy)	**Control (Renewed Levy)**
Panel A: Road Quality Rating (0, 1, 2)		
Pre-Election	1.62	1.56
	(0.490)	(0.499)
Post-Election	1.47	1.56
	(0.505)	(0.499)
Panel B: Road Quality Score (1–100)		
Pre-Election	74	71
	(17.3)	(17.9)
Post-Election	68	70
	(18.0)	(17.2)

Notes: The control group comprises townships within the effective RD bandwidth that successfully renewed their road tax levies; the treated group comprises those that failed to renew. Ratings are derived from our fine-tuned vision model. Scores follow Equation 2. Standard deviations are in parentheses.

Panel A of Table 2 summarizes the Road Quality Rating (RQR), showing that treated areas i.e. jurisdictions that cut their renewal road tax levies, started with a slightly higher average RQR (1.62) than control areas (1.56) before the referendum, but experienced a notable decline from 1.62 to 1.47 after the referendum, while control areas saw no change and remained at 1.56.

Panel B of Table 2 presents the continuous Road Quality Score (RQS), where treated areas also began with a higher average RQS (74) compared to control areas (71) before the referendum. After the referendum, RQS dropped notably from 74 to 68 in treated cities, but only slightly from 71 to 70 in control cities. This pattern suggests that failing to renew a road tax levy leads to a decline in both RQR and RQS, while renewing the levy is associated with fairly stable road quality.

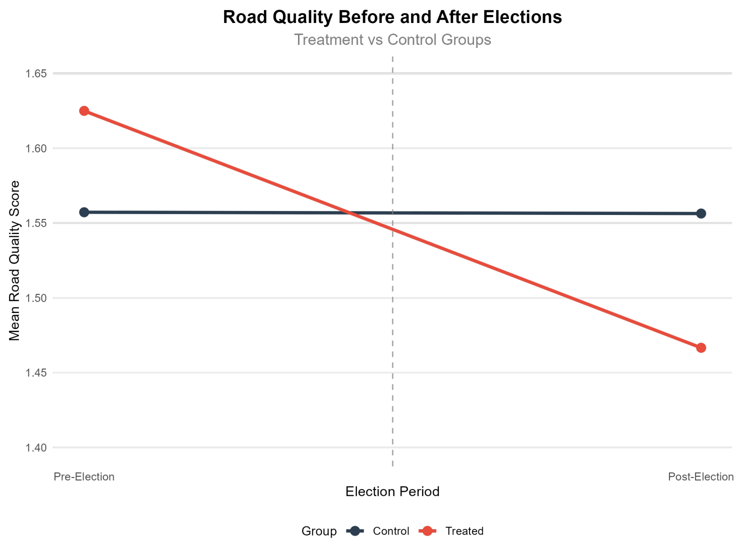

(A) Before and after referendum

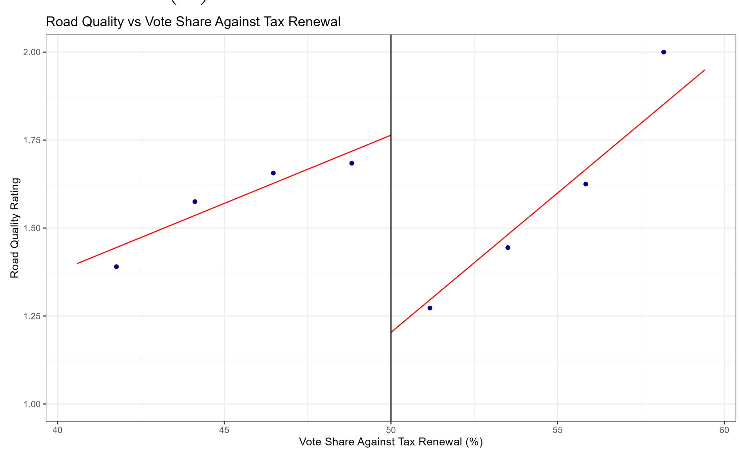

(B) Road Quality vs Vote Share Against Renewal (%)

Figure 3. Difference in Road Quality Rating by Treatment Status and Vote Share

Figure 3 Panel A graphically illustrates the difference in road quality rating before and after the referendum for both treated and control groups. The graph shows a clear decline in road quality rating for the treated group after the referendum, while the control group remains relatively stable. Panel B of Figure 3 shows road quality rating five years after the referendum graphed alongside the percent of votes against the tax levy. The points represent the local average road quality rating for the representative bins of vote shares within the average effective RD bandwidth around the cutoff of 50%. The graph suggests a discontinuity in road quality rating at the cutoff, which is the basis of our identification strategy, suggesting that the failure to renew a road tax levy has a negative impact on road quality. Section 1.5.1 provides a more detailed analysis of the impact of cutting road taxes on road quality, using the predictions from the model as the outcome variable, and employs regression discontinuity to estimate the effect of cutting road maintenance taxes on road quality.

1.3.5 Outcome Variable: Median House Price

Our house price data comes from a CoreLogic® dataset of actual sales transaction prices in Ohio from 1995 through 2021 containing over 7 million observations. The dependent variable, *Median House Price*, reflects the median sale price of houses within a specific city and year. For example, for the houses sold in Delaware Township during the year 2002, the median sale price was $205,041. We take precaution to only include arm's-length transactions, and we restrict our attention to single-family residential structures for comparability. The overall sample mean for the 10-year period from the time of votes considered in this study is $166,082 in constant 2010 dollars with a standard deviation of $372,135 which suggests the presence of outliers. Although our use of median sale price addresses outliers, one of our robustness checks drops 1% tails and re-estimates the treatment effects (Figure 9).

Figure 4 shows house prices from five years after the vote graphed alongside the percent of votes against the tax levy[15]. The points represent the mean house price for the 10 representative bins of vote shares for the average effective bandwidth around the cutoff of 50%. The graph suggests a discontinuity in house prices at the cutoff, which is the basis of our identification strategy, suggesting that the failure to renew a road tax levy has a negative impact on house prices.

[15] after deflating to 2010 U.S. dollars

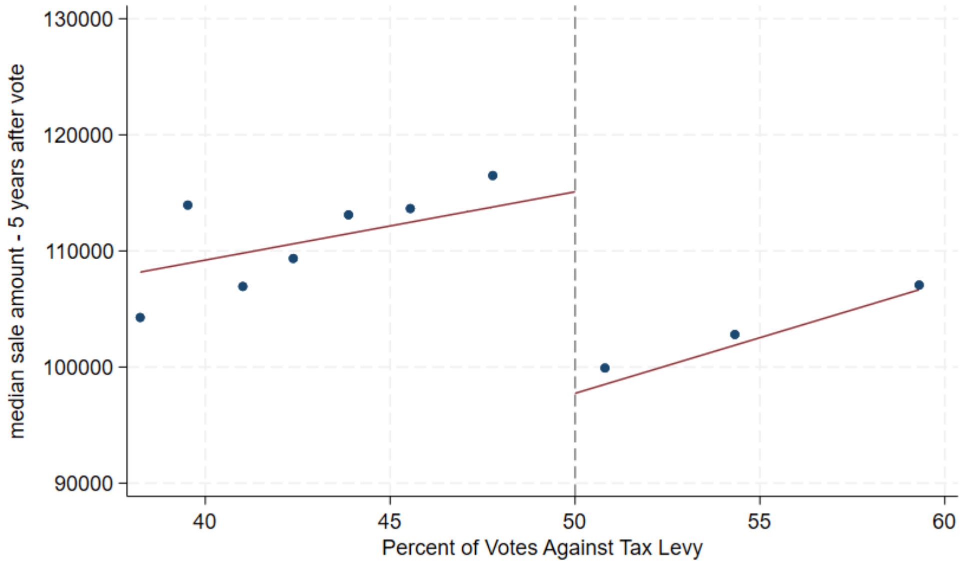

Figure 4. Median Sale Price of Houses: 5 years after vote

1.3.6 Covariates

Table 3 shows covariate means for both the global sample of all votes in the data set as well as the local sample within a representative effective bandwidth of the decision cutoff. The effective bandwidth displayed in Table 3 is the mean bandwidth for all the housing outcome regressions. The first columns for the global sample show similar values of characteristics between cities that renew and cut road taxes and spending. The second set of columns for the effective bandwidth show even closer similarity between cities that renew and cut road taxes and spending.

The data, captured at the time of the vote and as observed in Table 3, shows minimal differences within the effective bandwidth: mean population differs by only 301, and median family income varies by $436, measured in 2010 U.S. dollars. Other variables, including poverty rates, married household percentages, educational attainment, age distribution, and racial composition, show differences of two percentage points or less, bolstering the comparability of the treatment and control groups. This similarity suggests that any observed differences in outcomes can more confidently be attributed to the treatment effect rather than to pre-existing differences[16].

1.4 Empirical Strategy

In this section, we describe our empirical strategy to estimate the causal effect of cutting road maintenance spending on house prices. One key feature of our quasi-experimental design is the exogeneity of the timing of the election. The timing is determined by the natural expiration of a road maintenance tax levy, which is typically 5 years, and is not impacted by factors such as the prevailing economic conditions or whether a road tax levy was passed or failed in earlier years.

1.4.1 Regression Discontinuity in Panel Data Setting

Suppose that local government in area i and year t conducts a referendum to renew an existing road tax levy. Let v_{it} be the vote share against the renewal tax levy and v^* be the threshold determining the result of the referendum (levy fails to renew if $v_{it} > v^*$). Let $F_{it} = 1(v_{it} > v^*)$ be an indicator to represent if the renewal road tax levy fails, and y_{it} be the outcome variable median housing sale price. We can write Equation 3 as follows:

$$y_{it} = \alpha + \theta F_{it} + \epsilon_{it} \tag{3}$$

where α is the intercept, θ is the parameter of interest representing the causal effect of cutting a renewal road tax levy and ϵ_{it} is the error term representing all other determinants of the

[16] We also find similar covariate means for cities that have no road tax levies, which we explain in Section 1.5.4

Variable	Global			Effective	
	Full Sample	Renewed	Cut	Renewed (Control)	Cut (Treatment)
Population	5,072	4,733	5,139	4,885	5,186
	(7,936)	(7,291)	(8,058)	(7,036)	(8,229)
Poverty Rate	0.11	0.11	0.11	0.10	0.11
	(0.08)	(0.07)	(0.08)	(0.07)	(0.08)
% with Kids	0.39	0.40	0.39	0.39	0.39
	(0.08)	(0.08)	(0.08)	(0.07)	(0.08)
% Households with Children under 18	0.09	0.09	0.09	0.09	0.08
	(0.06)	(0.05)	(0.06)	(0.06)	(0.05)
% Less than High School Education	0.16	0.18	0.15	0.18	0.16
	(0.11)	(0.12)	(0.11)	(0.12)	(0.10)
% Some College Education	0.25	0.24	0.25	0.24	0.25
	(0.06)	(0.06)	(0.06)	(0.07)	(0.06)
% Renters	0.20	0.20	0.20	0.19	0.20
	(0.11)	(0.10)	(0.11)	(0.09)	(0.11)
Unemployment Rate	0.05	0.05	0.05	0.05	0.06
	(0.04)	(0.03)	(0.04)	(0.03)	(0.04)
% White	0.96	0.97	0.96	0.97	0.97
	(0.07)	(0.07)	(0.07)	(0.07)	(0.08)
% Black	0.02	0.02	0.02	0.02	0.02
	(0.07)	(0.06)	(0.07)	(0.07)	(0.07)
% Married	0.59	0.60	0.59	0.61	0.60
	(0.09)	(0.08)	(0.09)	(0.08)	(0.09)
% Separated	0.01	0.01	0.01	0.01	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Income Heterogeneity Index	0.10	0.09	0.10	0.09	0.09
	(0.08)	(0.07)	(0.08)	(0.07)	(0.06)
Median Family Income	61,018	58,761	61,467	59,934	60,370
	(17,649)	(13,915)	(18,270)	(13,655)	(15,713)
% Under 5 Years Old	0.06	0.06	0.06	0.06	0.06
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
% Aged 5 to 17	0.20	0.21	0.20	0.20	0.20
	(0.05)	(0.04)	(0.05)	(0.04)	(0.05)
% Aged 18 to 64	0.60	0.60	0.60	0.60	0.60
	(0.05)	(0.05)	(0.05)	(0.04)	(0.05)
% Racial Minority	0.04	0.03	0.04	0.03	0.03
	(0.07)	(0.07)	(0.07)	(0.08)	(0.07)
Number of Observations	3,184	2,656	528	653	269

Table 3. Variable Means by Road Tax Levy Renewal Status

outcome. Around a narrow enough window around the threshold v^*, we can estimate the causal effect of cutting a renewal road tax levy on the outcome variable y_{it} by comparing the outcome variable for cities that narrowly pass the referendum to those that narrowly fail it.

1.4.2 Intent-to-Treat (ITT) Estimator

We follow a model of RD design similar to Cellini, Ferreira and Rothstein (2010) and estimate the **Intent-to-Treat** or ITT estimator. We prefer using the ITT estimator instead of the alternative **Treatment on the Treated** (TOT) estimator because the ITT estimator is more suited to our setting given the independence of the renewal elections. As described in Cellini, Ferreira and Rothstein (2010), when the elections are independent, the ITT estimator equals the TOT estimator.

We operationalize our ITT estimator using Equation 4:

$$Y_{i,t+\tau} = \alpha_\tau + \kappa_t + F_{it}\theta_\tau^{ITT} + P_g(v_{it}, \gamma_\tau) + Z_{it}\beta_\tau + \epsilon_{i,t+\tau} \tag{4}$$

Equation 4 shows a city i that holds an election in year t and we study this city's outcome τ years later. $Y_{i,t+\tau}$ represents the outcome variable for city i at year $t+\tau$. We define treatment as failure of a city, village or township to renew its road maintenance tax levy, which is represented by the indicator F_{it}. θ_τ^{ITT} is the causal effect of failing to renew road tax on the outcome. $P_g(v_{it}, \gamma_\tau)$ is a polynomial function of the running variable v_{it}, which is the percent of votes against the renewal tax levy. α_τ and κ_t represent timing and year-specific fixed effects. Z_{it} is a vector of control variables that include city-level demographics, economic conditions, and other relevant covariates. $\epsilon_{i,t+\tau}$ is the error term.

We use the bandwidth selection method of Calonico et al. (2019) to find the mean optimal bandwidth h and then conduct a local polynomial regression after choosing a weighting scheme k. The bandwidth h determines the size of the neighborhood around the cutoff v^*, defined as

$(v^* - h, v^* + h)$. Only observations within this neighborhood are used to compute the bias-corrected treatment effect estimate $\hat{\tau}$. For a sufficiently small neighborhood, the continuity assumption central to the RD estimator is considered valid. We also cluster the standard errors by city to account for any serial correlation between years within each city. The weighting scheme k determines the weights of the observations within the neighborhood $(v^* - h, v^* + h)$ and is crucial in estimating θ_τ. Common weighting schemes include uniform, triangular, and Epanechnikov. We use the default Mean Squared Error Regression Discontinuity (MSERD) method to compute the effective bandwidth (h) and bias bandwidth (b) for the outcome variable. This method identifies the bandwidth that minimizes the trade-off between bias and variance of the treatment effect estimate. All observations are used to estimate h and b, but only those within the effective bandwidth h are used to identify our treatment effect estimates θ_τ for different τ years.

1.5 Results

1.5.1 Road Quality Decline

First, we present the results of our analysis on the impact of cutting road maintenance taxes on road quality. As discussed in Section 1.3.4, we use a fine-tuned AI vision model to assess road quality based on satellite imagery data. The model gives us a road quality rating of 0 (poor), 1 (medium), or 2 (high), which we also convert to a continuous variable called Road Quality Score (RQS) ranging from 1 to 100.

Table 4 reports RD estimates of the discontinuity in post-election road quality at the 50% threshold. Column (1) uses the categorical Road Quality Rating (RQR, 0–2), and Column (2) uses the continuous Road Quality Score (RQS, 1–100), both derived from the fine-tuned vision model described in Section 1.3.4. As shown in Table 4, we find that cutting road

Table 4. Change in Road Quality after an Election

	(1)	(2)
RD Estimate	-0.477**	-15.634**
	(0.222)	(7.366)
Covariates	✓	✓
Clustered SEs	✓	✓
Effective Sample	178	178

Notes: Column (1) shows the RD estimate of the effect of cutting road maintenance taxes on the Road Quality Rating (RQR, 0–2), and Column (2) shows this effect on the continuous Road Quality Score (RQS, 1–100). Standard errors are in parentheses. The table reports covariate-adjusted sharp RD estimates from a polynomial regression ($p = 1$, $q = 2$) at the 50% vote-share cutoff. The bandwidth is mean effective RD bandwidth from Table 5. Inference is based on bias-corrected point estimates with robust standard errors. Standard errors are clustered by ten-digit FIPS code. Significance levels: *p $<$ 0.10, **p $<$ 0.05, ***p $<$ 0.01.

maintenance taxes does lead to a decline in road quality of -0.477 for the RQR metric. This suggests that relative to similar cities that renew their road maintenance taxes, cities that cut their road maintenance taxes experience a decline of almost half a category in road quality after the election. The estimate for the RQS metric is -15.634, indicating a decline of about 15.6 points on a 100-point scale. Using the pre-election means among treated cities of 1.625 (RQR) and 74 (RQS), these estimates of difference in road quality correspond to approximately 29.4% and 21.2% reduction in road quality. The limited sample size is an artifact of our focus on close elections, limited availability of high-resolution satellite imagery for some cities and years, and the cleaning process to ensure that the road segments analyzed are representative of public roads. Although these results may be sensitive to data availability, sample size and the performance of the vision model, they provide suggestive evidence that cutting road maintenance taxes leads to a significant decline in road quality. The next section presents our main results on the impact of cutting road maintenance taxes on house prices.

1.5.2 House Prices Decline

Table 5 below shows the ITT estimates of failing a road tax levy on housing sale prices. Each treatment effect estimate represents the discount in median sale price for cities that cut road taxes relative to otherwise similar cities that renew the taxes. We show dynamic effects starting three years before the vote and extending to ten years after the vote. Treatment effect estimates for years 4 through 9 after the vote are statistically significant in Table 5. The estimate for year $t + 10$ has a smaller estimate and is only significant at the 10% level, suggesting that the effect of tax and service cuts on house prices may peter out ten years after the vote. Overall, we find an average reduction of $15,349 in median house price over the 10-year period for houses in cities that vote to cut road tax levies, representing 9% of overall house value[17]. The 9% reduction we estimate may be compared to Gonzalez-Navarro and Quintana-Domeque (2016) which observes a change in property values of 17-28% after paving previously unpaved roads and the 13% increase in house prices Theisen and Emblem (2021) finds in towns nearest to a newly-constructed highway[18].

[17] The average treatment effect estimate of $15,349 was divided by the mean house sale price in the dataset of $166,000 to get 9%.

[18] Note that these studies focus on the effect of developing paved roads whereas we focus on deterioration of existing paved roads.

Table 5. Effect on median house prices of failing to renew a road tax levy

Panel A: Years $t-3$ through $t+3$							
Year	$t-3$	$t-2$	$t-1$	t	$t+1$	$t+2$	$t+3$
Treatment effect	5,307	166	–273	–4,261	–3,908	–11,001	–14,733
Standard error	(7,341)	(6,943)	(7,391)	(7,955)	(8,719)	(9,405)	(7,989)
Effective bandwidth (h)	10.30	10.06	8.76	7.97	9.78	9.63	11.93
Bias bandwidth (b)	20.22	17.48	14.39	13.03	17.50	18.22	22.69
Effective observations	724	725	650	587	774	778	1,003
Total observations	2,552	2,631	2,696	2,784	2,764	2,699	2,640
Panel B: Years $t+4$ through $t+10$							
Year	$t+4$	$t+5$	$t+6$	$t+7$	$t+8$	$t+9$	$t+10$
Treatment effect	–21,701	–21,706	–17,365	–15,975	–21,984	–19,857	–16,090
Standard error	(7,747)	(8,751)	(8,355)	(7,248)	(9,074)	(7,751)	(9,027)
Effective bandwidth (h)	8.52	11.20	9.79	13.16	8.25	7.28	6.24
Bias bandwidth (b)	16.25	20.16	17.32	23.42	15.28	14.09	16.27
Effective observations	688	922	764	1,061	591	505	402
Total observations	2,614	2,531	2,438	2,324	2,199	2,115	2,017

Notes: Outcome is median house price in constant 2010 U.S. dollars. Unit of observation is the city–year, so a treatment effect of –\$21,701 in year $t+4$ means that four years after the vote, cities that fail to renew road taxes have median sale prices \$21,701 lower than similar cities that renew. Covariates include the demographic and socioeconomic controls listed in Table 3.

Figure 5 provides an event-study plot that summarizes the treatment effects for each time period. In the graph, we include placebo years up to 3 years before the treatment to

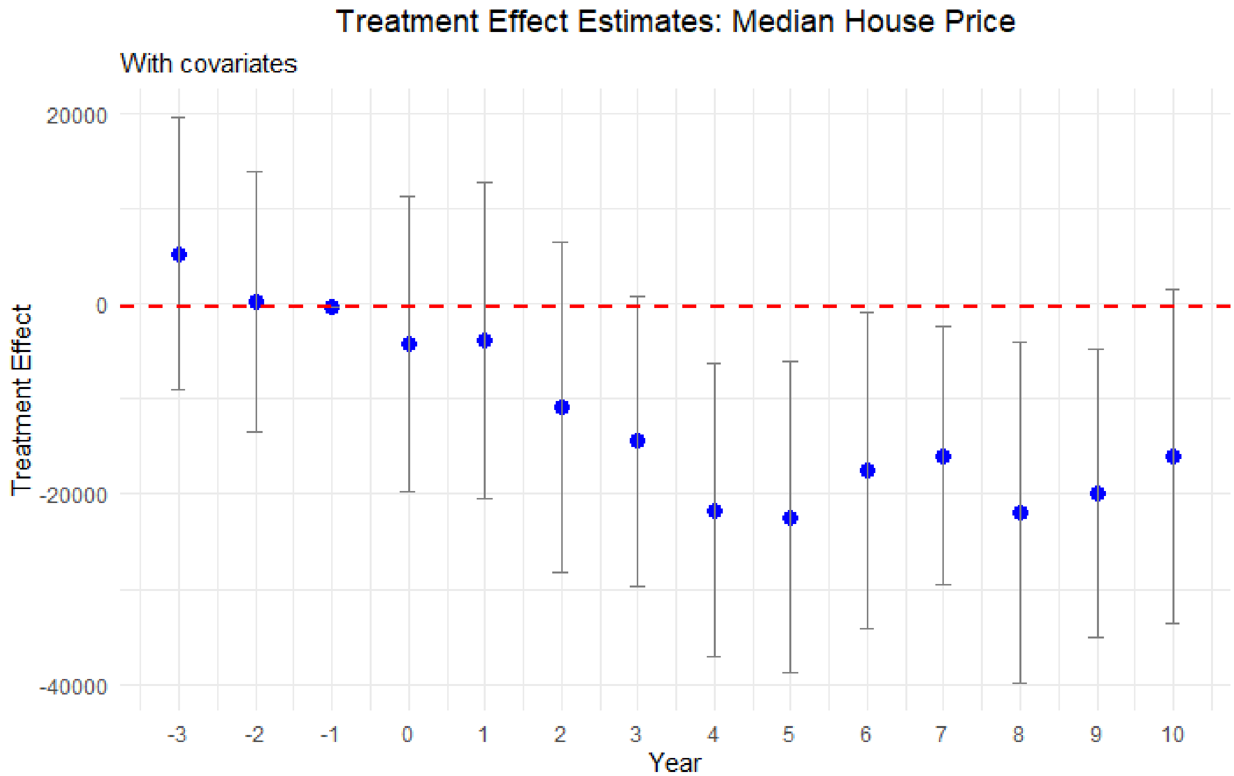

Figure 5. Effect plot for Median Housing Price

show that the median housing prices are statistically identical for cities above and below the threshold prior to the treatment. Each dot represents the treatment effect point estimate for that year and the bar around it represents the 95% confidence interval for that estimate. For year 0, which is the year of the vote, we see a slight decrease in the estimate. However, this effect is not statistically significant, as evidenced by the confidence interval containing the null effect. Up to year 3, we observe that the treatment effect estimates are fairly close to zero, and the confidence interval overlaps with zero. As stated previously and shown in Table 5, we start to see a sizable increase in treatment effect from year 4 onwards and continue to observe it through year 9 after the vote.

1.5.3 Heterogeneity Analysis

We show the results of our heterogeneity analysis, where we explore the differential impact of cutting road maintenance spending on house prices in urban and rural neighborhoods, for tax levies of different sizes, and for different housing price quantiles.

Urban vs Rural neighborhoods: Analyzing regional heterogeneity is important because the impact of public goods like road maintenance on house prices can differ substantially between areas due to factors such as differences in elasticity of housing supply. In this section, we check how treatment effects differ between urban and rural areas[19].

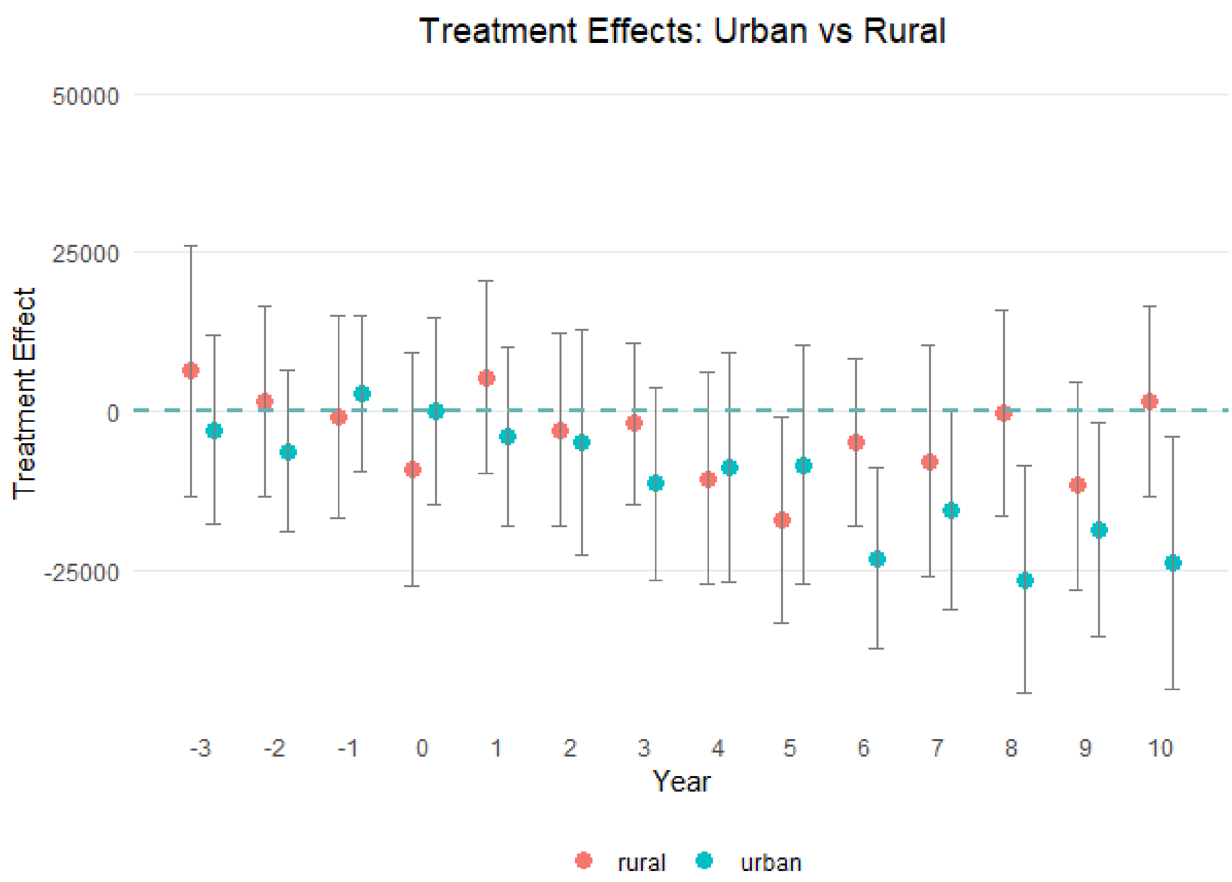

Figure 6. Median Housing Price in Urban and Rural Areas

[19]Urban areas are identified using data from U.S Census Bureau. We strictly consider urban areas with atleast 0.3 allocation factor in 2010 and do not consider urbanized clusters, making our classification of urban areas more restrictive.

As shown by Figure 6, we do not find any significant differences in housing prices in rural areas after a renewal tax levy fails to pass. On the other hand, we do find a statistically significant decline in housing prices in urban areas starting six years after voting. The standard errors are somewhat smaller for the rural estimates due to a larger number of observations. The stronger effects in urban than rural areas may stem from a more inelastic housing supply (Brasington, 2002). Overall, we find that housing prices decrease by $13,302 on average over the decade[20] after cutting road maintenance tax levies in urban areas.

Tax magnitude: We check for dose-response by splitting the sample by the size of the tax levy. First, we consider splitting the sample based on the median size of the tax levies. The median size of the tax levies in our sample is 1.9 mills, however we find that the tax levies above and below the median size do not show any significant differences in treatment effects. Next, we focus on the top quartile of the tax levies, which in our sample is the set of tax levies ranging from 2.1 to 8 mills. Figure 7 below presents the results for this subsample. The treatment effect estimates for these larger tax cuts are both statistically significant and of greater magnitude, indicating a more pronounced decline in house prices when the reduction in road tax levies is more substantial.

Despite the reduction in sample size leading to wider standard errors in Figure 7, the overall pattern in decline of house price remains, with mean treatment effect for this subsample being approximately $30,000 which is about double the effect observed in the full sample. This finding supports the presence of an intensive margin effect: the larger the tax cut, the greater the negative impact on housing prices. Such a pattern is consistent with a dose-response relationship, where the magnitude of the fiscal shock translates directly into the size of the decline in property values.

[20]This is equal to $7.6\% = \frac{13{,}302}{175{,}217} \times 100$, where the denominator is average sale price of houses in urban areas in our dataset.

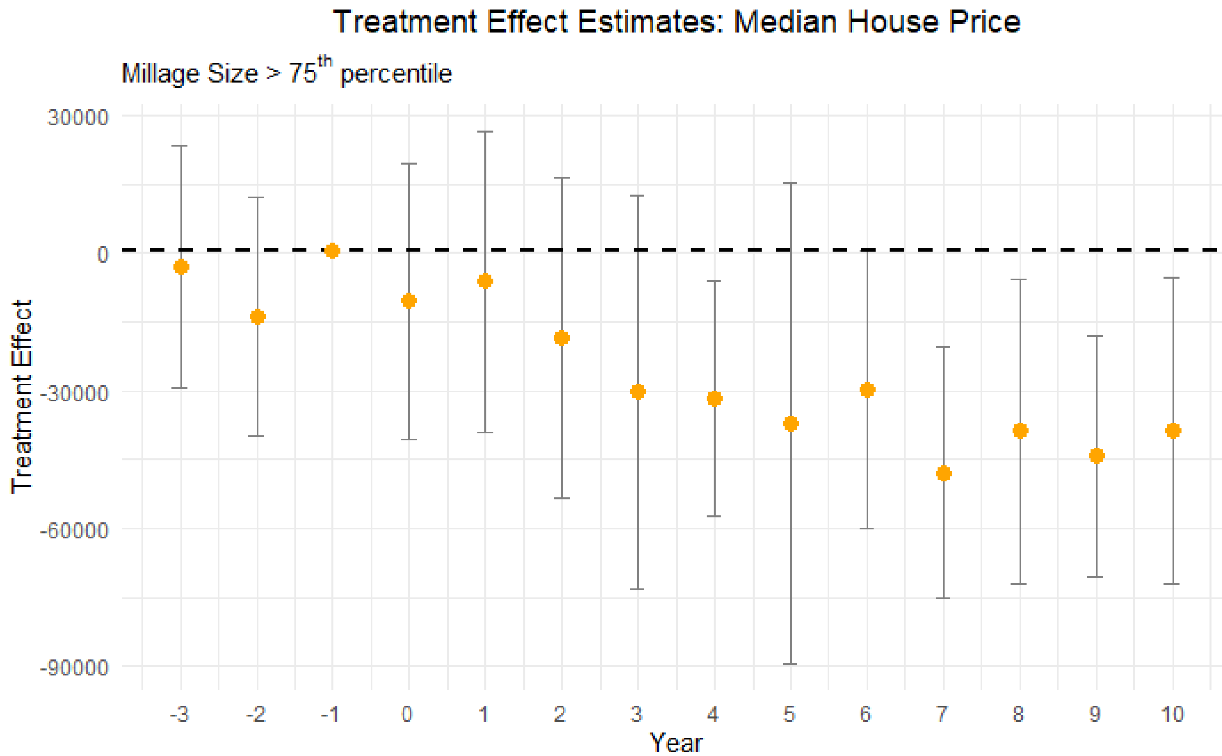

Figure 7. Effect of Large Road Maintenance Tax Cuts (>2 mills) on Median House Price

RD Quantile Estimation: We analyze our results by estimating quantile-level treatment effects, as suggested by Frandsen, Frölich and Melly (2012), to study how the treatment's impact varies at different quantiles of the outcome variable.

Table 6 shows the treatment effect heterogeneity of cutting road spending on high and low quantiles of median house prices. The top percentiles consistently exhibit a statistically significant decline in house sale prices, beginning in year 6 after the reduction in road spending. In contrast, the lower percentiles do not demonstrate a consistent treatment effect. This suggests a differential impact, where higher-valued properties are more sensitive to road disrepair than lower-valued houses. Figure 8 contrasts the treatment effects of the 20th and 80th percentiles of house sale prices in an effect plot to highlight this differential impact of reduction in road maintenance spending.

Table 6. Quantile-level Treatment Effects of Cutting Road Spending on Median House Prices

Pct.	$t+4$	$t+5$	$t+6$	$t+7$	$t+8$	$t+9$	$t+10$
10%	-6,433	-22,570	-9,602	-12,984	-11,217	-6,569	-1,326
	(9,364)	(9,065)	(9,205)	(8,420)	(9,136)	(10,809)	(8,793)
20%	-5,400	-15,070	4,014	-14,682	-15,040	-3,228	624
	(9,983)	(9,886)	(7,443)	(8,502)	(8,160)	(10,435)	(8,509)
70%	-21,760	-11,171	-38,082	-36,685	-21,356	-25,605	-18,600
	(12,333)	(11,806)	(12,835)	(12,163)	(12,218)	(13,984)	(9,872)
80%	-28,478	-16,379	-38,460	-37,470	-28,950	-27,800	-18,658
	(13,343)	(11,404)	(18,623)	(12,169)	(12,507)	(12,421)	(11,808)
90%	-51,470	-34,604	-38,510	-27,039	-29,010	-49,093	-36,662
	(18,409)	(15,837)	(22,194)	(16,308)	(16,640)	(14,498)	(19,110)

Notes: The outcome is median house price in constant 2010 U.S. dollars. The unit of observation is the city-year, so a treatment effect of −\$28,478 means that at the 80th percentile of house prices four years after the vote, cities that fail to renew road taxes and the associated spending have houses that sell for \$28,478 less than cities that vote to renew them. Covariates include demographic and socioeconomic characteristics as described in Table 3.

1.5.4 Robustness Tests

We conduct several robustness tests to ensure the validity of our results. We test the sensitivity of our treatment effect estimates by removing contaminated observations and winsorizing the outcome variable to reduce the influence of outliers. Lastly, we consider other tax levies to alleviate endogeneity concerns.

Removing contaminated observations: In this test, we focus on independence and exogeneity. Estimates may be biased if tax levies for additional money pass after renewal tax levy decisions. To address this concern, we exclude observations from our analysis if a tax levy for additional funding is introduced and passed within a ten-year period following a renewal tax levy vote, as these votes may contaminate the treatment effect of the renewal tax levy vote. For example, consider a scenario where a city votes to fail its renewal tax levy in the year 2000 i.e. funding for road maintenance is cut. If that city subsequently introduces and passes a tax levy for additional road spending in 2004, we exclude all votes for that city

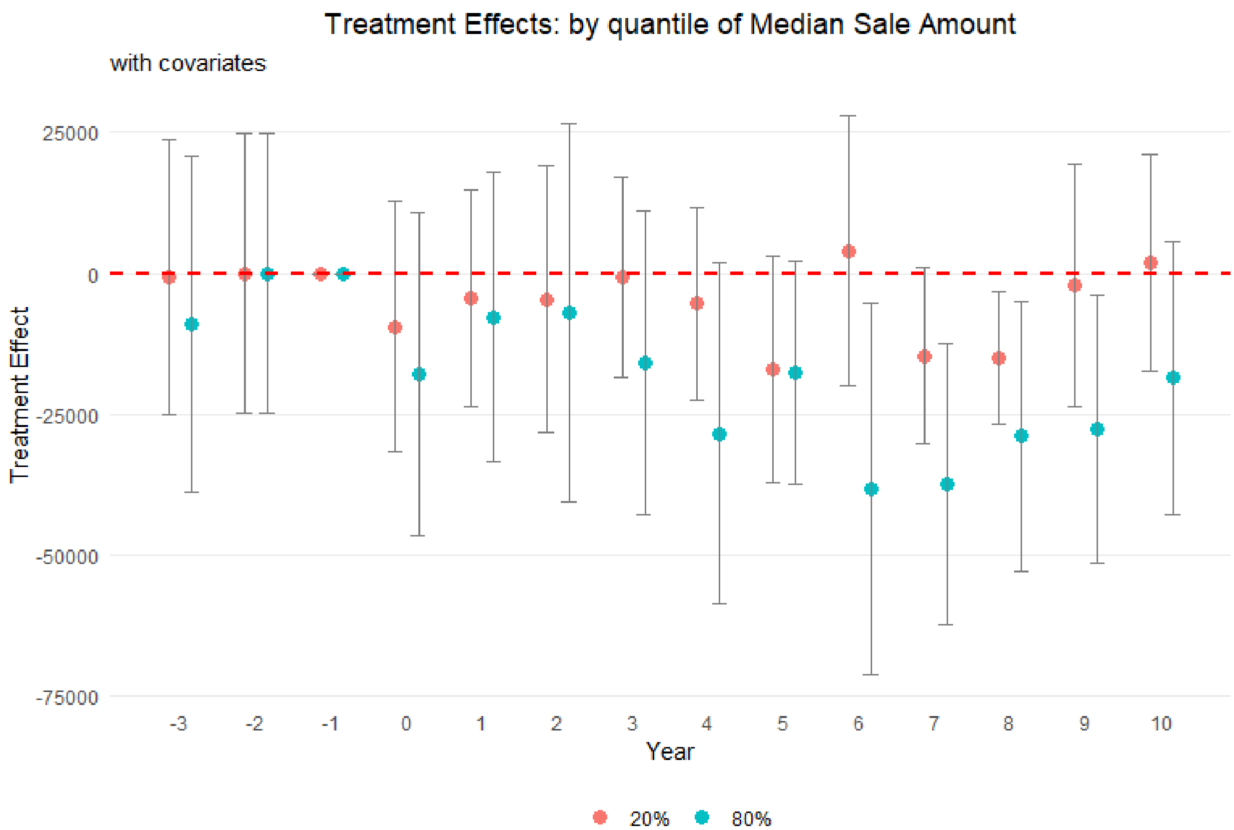

Figure 8. Treatment Effect of Cutting Road Maintenance Taxes for 20^{th} and 80^{th} Percentiles of Median House Price

from 2005 through 2010. This exclusion ensures that the effect on house prices from the 2000 vote are captured for uncontaminated years but not for years after 2004 when the effect of additional road taxes may counteract the drop in tax money from the elections in the year 2000.

Table 7. Effect on median house prices of failing to renew a road tax levy – uncontaminated observations only

Panel A: Years $t-3$ through $t+3$							
Year relative to vote	$t-3$	$t-2$	$t-1$	t	$t+1$	$t+2$	$t+3$
Treatment effect	5,295	–369	1,207	–2,482	–3,107	–12,236	–15,026
Standard error	(7,255)	(7,359)	(7,776)	(8,482)	(9,270)	(10,123)	(8,746)
Effective bandwidth (h)	10.25	9.70	8.84	7.48	9.19	8.20	10.09
Bias bandwidth (b)	20.08	16.41	14.45	12.03	16.04	16.33	20.21
Effective observations	719	682	630	510	680	604	763
Total observations	2,552	2,574	2,593	2,640	2,627	2,536	2,449
Panel B: Years $t+4$ through $t+10$							
Year relative to vote	$t+4$	$t+5$	$t+6$	$t+7$	$t+8$	$t+9$	$t+10$
Treatment effect	–21,082	–16,827	–17,565	–18,179	–20,427	–29,970	–23,593
Standard error	(9,145)	(11,228)	(10,128)	(9,202)	(9,829)	(9,059)	(10,838)
Effective bandwidth (h)	6.85	5.64	8.48	9.34	7.41	5.51	6.20
Bias bandwidth (b)	13.19	14.82	14.53	18.06	14.23	11.74	16.54
Effective observations	473	367	573	609	442	303	322
Total observations	2,389	2,274	2,145	2,016	1,890	1,787	1,666

Notes: Sample excludes any potentially contaminated city–year observations. Outcome is median house price in constant 2010 U.S. dollars. Each coefficient is the difference between treated and control cities in the indicated year; standard errors are heteroskedasticity-robust. All regressions control for the demographic and socioeconomic covariates listed in Table 3.

Upon implementing this data filtration, we observe that the treatment effect of the renewal levies on housing prices, measured from $t+1$ to $t+10$, remains consistent with our initial findings. This consistency in treatment effect, despite the exclusion of potentially confounding data, lends credence to our results. The standard errors increase slightly due to the reduction in sample size caused by the aforementioned data filtration process.

Placebo cutoffs: In our primary analysis, the pivotal threshold for the vote share running variable is 50%, indicating whether a renewal levy passes or fails. Although we find significant treatment effects using this 50% threshold, it could be random jumps in the data rather than cutting road taxes and funding that are responsible for the significant estimates. To this end we conduct a series of placebo tests using alternative cutoffs: 30%, 40%, 60%, and 70%. Table 8 below summarizes the results from the placebo cutoffs analysis.

Table 8. Robust Treatment Effect Estimate for Placebo Cutoffs

Years after vote	30%	40%	60%	70%
$t+4$	2,578	9,149	9,419	-12,987
	(8,209)	(7,284)	(11,462)	(14,365)
$t+5$	-6,381	-29,077	6,383	41,683
	(9,086)	(20,680)	(10,786)	(17,836)
$t+6$	7,681	5,573	-1,095	-14,226
	(9,616)	(8,120)	(8,612)	(15,733)
$t+7$	1,162	3,982	-12,050	22,261
	(10,468)	(8,191)	(9,396)	(19,765)
$t+8$	4,334	12,881	3,593	31,696
	(9,670)	(8,625)	(10,061)	(6,902)
$t+9$	851	7,381	-6,935	42,790
	(5,921)	(6,333)	(7,114)	(15,663)
$t+10$	10,032	-35,038	324	-8,566
	(10,599)	(35,569)	(8,220)	(17,281)

Notes: Robust treatment effect estimates at placebo cutoffs using the estimator from Calonico et al. (2017). The unit of observation is the city-year. Standard errors are reported in parentheses below each estimate.

Table 8 does not show consistently significant treatment effects for any of the placebo cutoffs for our parameter of interest. This absence of significance at thresholds other than the true 50% reinforces the idea that the effects we observe at the 50% mark are not a mere coincidence or a result of random variation in the data, but are indeed attributable to the dynamics surrounding the passing or failing of renewal tax levies.

Winsorization: The debate over whether to include or exclude outliers continues, with some research suggesting that trimming outliers does not improve mean squared error (Bollinger and Chandra, 2005). We now drop the 1% tails to help curtail the influence of outliers. The overall sample mean after dropping 1% tails is $150,505 in constant 2010 dollars with a standard deviation of $115,646. After performing this winsorization step, we re-estimate the treatment effect of failing to renew a road tax levy on housing outcome variables. The results from this estimation process are summarized in Figure 9 below:

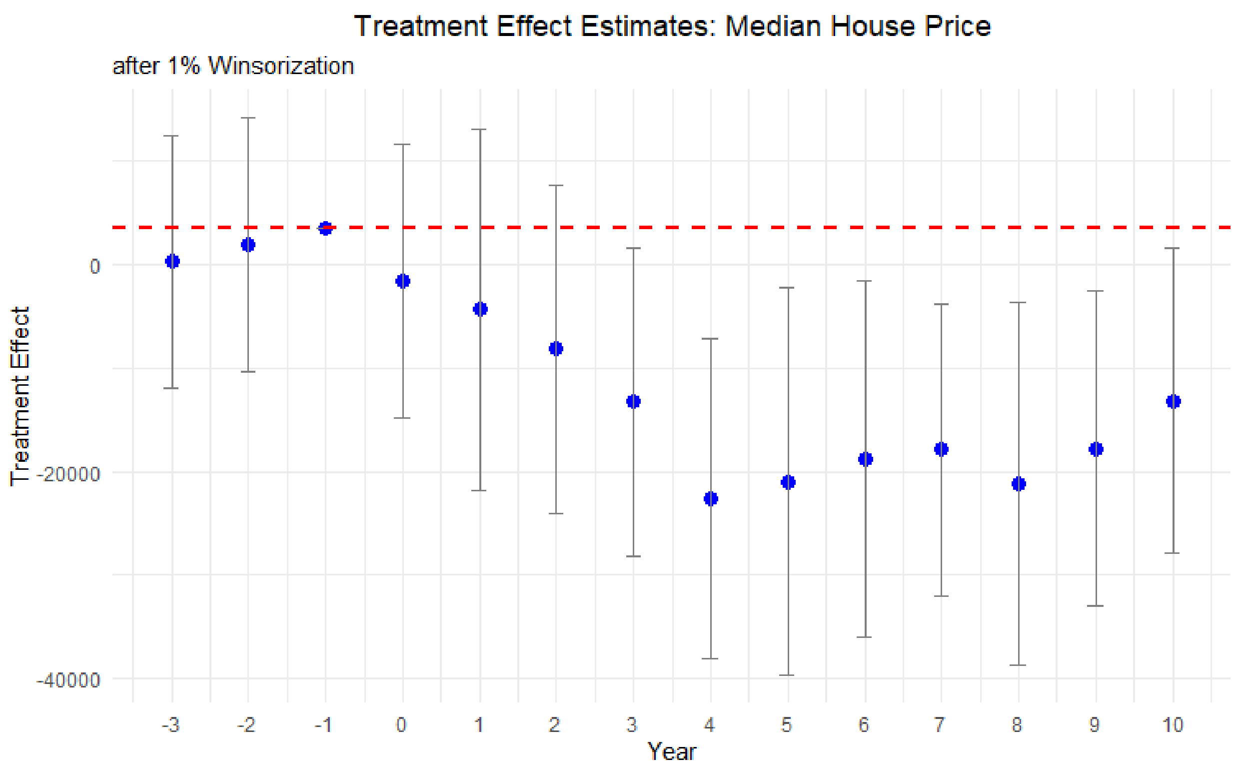

Figure 9. Effect of Cutting Road Maintenance Taxes on Median Housing Price after 1% Winsorization

The treatment effect estimates with winsorization mimic those from our baseline regression results qualitatively and quantitatively.

Other Tax Levies: A potential threat to identification is that a cut in road tax levy might be correlated with cuts in other local tax levies, or may simply reveal an underlying, time-varying taste for lower taxes or smaller government. If the same electorate simultaneously rejects (or approves) levies for police, fire, recreation, or schools, our estimated road-maintenance effect could in fact be picking up broader changes in local public-service bundles that are themselves capitalised into housing prices. To address this concern, we examine whether the election results of a renewal road levy is systematically correlated with the results of other levy referenda held on the same ballot or within the same fiscal year i.e. whenever a road tax levy is cut, are there elections and subsequent tax cuts for other tax levies as well?

Table 9. Association of Road Tax Levy Referenda Results with Other Types of Levies

	Police	Fire	Recreational	School
Estimate	0.248	0.053	0.015	-0.024
	(0.153)	(0.043)	(0.317)	(0.092)

Notes: This table presents coefficients from regressions associating road tax levy referendum outcomes with outcomes for other types of levies (police, fire, recreational, school). Standard errors are reported in parentheses. All regression control for year and neighborhood fixed effects, as well as neighborhood characteristics. The school levy analysis is conducted at the county level, while all other analyses are at the county subdivisions level. Statistical significance levels are indicated as follows: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 9 reports coefficients from separate regressions in which each column regresses the pass/fail outcome of another levy type on the pass/fail outcome of road tax levy[21], controlling for year fixed effects, jurisdiction fixed effects, and demographic covariates. For school levies we cluster to the county level as an approximation, while all other specifications are at the township/city ("county subdivision") level. The point estimates are small and statistically indistinguishable from zero: a failed road-maintenance renewal is neither more nor less likely to coincide with the failure of levies earmarked for police, fire, recreation, or schools. For example, the largest coefficient, 0.248 for police levies, has a standard error of 0.153 and is not significant even at the 10 percent level. These results suggest that voters treat the road-maintenance question separately from other local tax questions, and that decisions to cut or renew road funding are not simply proxies for a general anti-tax or austerity sentiment. Consequently, our main RD estimates are unlikely to be confounded by contemporaneous changes in other public services.

[21] This is represented by a dummy variable indicating a cut in road tax levies, with 1 as denoted as a cut and 0 denoted as a renewal.

A more comprehensive approach to address this concern is to check for balance in other tax levies around the cutoff. As the theory of regression discontinuity suggests, cities close to the 50% vote cutoff are as good as randomized along the cutoff. If true, there should be no difference in funding changes between the pass-levy and fail-levy groups. To confirm this, we collect all 50,000 local tax levies for all municipal purposes from 1991 to 2021, including current expense tax levies, parks and recreation, and police funding, for example. We identify the tax levies for the cities within the effective bandwidth and compare the proportion of cities that renew and fail to renew road tax levies that have other tax levies on the ballot for different purposes at the same time as the road tax levy. The pass-levy group has a non-road tax levy on the ballot 48.1% of the time, and similarly the fail-levy group has a non-road tax levy on the ballot 47.96% of the time. The two groups are balanced, as suggested by RD theory. We next focus on the set of cities within our effective bandwidth with non-road tax levies on the ballot. The cities that barely vote to renew road taxes experience a change in other taxes, either an increase in new spending or a cut in existing spending, 12% of the time. The corresponding number for the set of cities that barely votes to cut road taxes is 15%. Again, randomization near the cutoff suggests that both observable and unobservable characteristics are balanced between pass and fail road tax levy groups. Thus, the difference in house prices can be attributed to tax cuts in road maintenance funding and decreases in road quality.

1.6 Mechanisms

A key question surrounding our findings is *why* cutting road taxes would lead to a sustained decline in local house prices. In this section, we argue that these results are consistent with classical insights from the literature on local public finance, especially the work of Oates (1972) and Edel and Sclar (1974). We develop a theoretical framework using Dynamic General

Equilibrium (DGE) model to explain the mechanisms behind our empirical findings. The model captures the dynamic interplay between local government budget constraints, road quality, and housing prices. Specifically, cutting road-maintenance tax levies reduces local government funds for maintaining infrastructure, which in turn constrains the government's ability to provide road-upkeep services. These lower levels of maintenance funding result in a deterioration of road quality, which then directly and indirectly lowers housing values. For brevity, we summarize the key insights here and relegate the full model description to Appendix A1.4. Below, we outline the three-step channel through which road tax cuts lead to lower house prices.

1. Road Tax Cuts and Reduced Road Maintenance Funds. Local governments in Ohio, like those in many other U.S. states, rely on property taxes and dedicated road levies to fund local public goods. The Oates (1972) decentralization theorem highlights that local authorities can generally provide public services in a manner better aligned with resident preferences than would a higher-level government, provided they have adequate revenue sources. When a renewal levy fails at the ballot box, one of these critical revenue sources vanishes. Consistent with the premise that local road quality is a local public good, losing a dedicated revenue stream substantially weakens a city's ability to maintain or upgrade roads. Empirically, we infer the size of the drop in maintenance funds when a maintenance levy fails to renew in Table 1.

2. Decline in Road Quality. Once the road tax levy is not renewed and the budget is cut, the local government has fewer resources to maintain roads, but does it reflect in the quality of roads? We study this question and explain our method to measure road quality in Section 1.3.4. Results from our AI model-based road-quality measures in Section 1.5.1 show a marked decrease in road quality following the cut of renewal road tax levies. The resulting deterioration in road quality is primarily seen through remote sensing data, but may be experienced via ride discomfort, declining aesthetics, and reduced usability of neighborhood roads. This is precisely what a local public goods framework would predict: fewer financial resources for upkeep reduces the quality of a local public good. Residents observe these changes slowly over time, often after four to five years, when road conditions become visibly poor.

3. Capitalization into House Prices. Falling road quality imposes a disamenity on local residents via reduced local infrastructure that is capitalized into house values (Oates, 1969). A couple of mechanisms help explain the ultimate discount in property values:

- **Appearance and Neighborhood Appeal.** Rough roads, potholes, and poorly patched surfaces reduce aesthetic appeal of the houses in a neighborhood. Potential buyers, seeing visible signs of neglect, bid less, thereby reducing the sale price of houses.
- **Travel Costs.** Lower quality roads and chronic maintenance issues translate into slower commute times for residents, particularly in urban areas. This is a direct disutility for homeowners, as time lost in traffic takes time away from leisure or work.

Short-Run versus Long-Run Trade-Off. Importantly, a trade-off emerges over time. In the short run, residents in cities that voted against a road tax levy experience immediate relief in terms of lower property tax bills[22]. Their out-of-pocket expenses decline, which can be a tangible financial benefit, particularly in communities where property tax burdens were already viewed as high. However, as road quality starts to visibly deteriorate from around the fourth year onward (see Table 5), the same residents find themselves negatively affected by a decline in house prices. This lag reflects the time it takes for infrastructure disrepair to become apparent, be it through more frequent potholes or visibly eroding surfaces. By the time these problems are evident, the city's real-estate market has had sufficient time to register the disamenity, capitalizing it as a discount in property values. This dynamic captures a fundamental insight of local public finance: the preferences of current taxpayers can diverge from the longer-term public interest (Buchanan and Tullock, 1962; Alesina and Tabellini, 1990).

[22]Table 1 reflects the relief in property tax payments for households when a road tax levy is not renewed.

1.7 Welfare Analysis

Following Hendren and Sprung-Keyser (2020) and Finkelstein and Hendren (2020), we conduct welfare analysis to evaluate the overall impact of the road maintenance tax cuts and compute a conservative estimate of the Marginal Value of Public Funds (MVPF). We define MVPF as the ratio of household's willingness to pay (WTP) for the tax cut to the government's net cost of providing the tax cut. Specifically, we compute MVPF as follows:

$$\text{MVPF} = \frac{\text{WTP}}{\text{Net Cost to Government}} = \frac{\text{Tax Savings - House Value Decline}}{\text{Direct Long Run Cost} + \text{Fiscal Externality}} \tag{5}$$

Where,

- **Direct Long Run Cost** is the tax revenue loss per household to the government due to the road maintenance tax cut over the life of the tax cut.
- **Fiscal Externality** is the additional loss in property tax revenue to the government due to the decline in tax base following the decline in house prices.
- **Tax Savings** is the overall value of tax savings from the road maintenance tax cut.
- **House Value Decline** is the average of the yearly capitalized loss in house values over the period of analysis.

Using this framework, we compute an MVPF of -\$1.4, indicating that for every dollar of revenue loss the government, the households incur a net loss of \$1.4, after accounting for both direct costs and fiscal externalities. This negative MVPF suggests that the road maintenance tax cut leads to a net welfare loss for society, making it a welfare destroying policy. To put this in perspective, although Hendren and Sprung-Keyser (2020) do not report a sub-category for public infrastructure and focus on government spending programs instead of tax cuts, their MVPF estimates for policies targeting adults range from 0.5 to 2.

1.8 Conclusion

Most studies examine the impact of new infrastructure on housing values, particularly in developing nations (Asher and Novosad, 2020; Huang, 2016; Li et al., 2016). Recent work emphasizes the challenge of endogenous placement where new projects are correlated with unobserved growth potential (Biasi, Lafortune and Schönholzer, 2025; Donaldson and Hornbeck, 2016). We circumvent this identification problem by studying the maintenance of existing infrastructure in a developed economy. This focus is particularly salient for policymakers in the U.S. and other advanced economies, where the primary infrastructure challenge is not new construction but the preservation of a massive stock of existing assets.

By constructing a novel panel of over 3,000 renewal referendums and matching election results to local financial records and property sales, we design a quasi-experiment to estimate the causal effect of road maintenance tax cuts. We also introduce a scalable method for measuring infrastructure condition by fine-tuning an AI vision model on satellite imagery. Our results reveal that local governments failing to renew road levies face an 11% loss in maintenance funds. This fiscal shock translates into a 23% decline in road quality and, subsequently, a 9% decrease in median house prices, with effects emerging four years after the tax cut.

Our heterogeneity analysis indicates that these effects are driven by urban areas, which see a 7.6% decline in house prices, while rural areas show inconsistent effects. We also find that higher-priced homes are more sensitive to road quality shocks, and that larger tax cuts lead to larger price declines consistent with a dose-response relationship. Using a Dynamic General Equilibrium (DGE) framework, we rationalize these findings as a trade-off where the immediate benefit of lower taxes is eventually outweighed by the capitalization of public capital depreciation. Finally, our welfare analysis yields an MVPF of -1.4, implying a net welfare loss to the households due to the road maintenance tax cuts.

From a policy perspective, our findings underscore the hidden costs of deferred maintenance. While tax cuts offer short-term relief, they can trigger a vicious cycle of infrastructure decay and asset devaluation that leaves homeowners worse off. Future research could extend our RD framework to other public goods funded by local levies, or other geographies. Advances in AI vision architectures open new avenues for researchers that could potentially allow real-time evaluation of road quality, and enable timely assessments of infrastructure policies and their welfare consequences.

References

Alesina, Alberto, and Guido Tabellini. 1990. "A Positive Theory of Fiscal Deficits and Government Debt." *Review of Economic Studies*, 57(3): 403–414.

Asher, Sam, and Paul Novosad. 2020. "Rural roads and local economic development." *American Economic Review*, 110(3): 797–823.

Asher, Sam, Teevrat Garg, and Paul Novosad. 2020. "The ecological impact of transportation infrastructure." *The Economic Journal*, 130(629): 1173–1199.

Attoh-Okine, Nii, and Offei Adarkwa. 2013. "Pavement condition surveys–overview of current practices." *Delaware Center for Transportation, University of Delaware: Newark, DE, USA*.

Beenstock, Michael, Daniel Feldman, and Daniel Felsenstein. 2016. "Hedonic pricing when housing is endogenous: The value of access to the trans-Israel highway." *Journal of Regional Science*, 56(1): 134–155.

Biasi, Barbara, Julien Lafortune, and David Schönholzer. 2025. "What Works and for Whom? Effectiveness and Efficiency of School Capital Investments Across the U.S." *The Quarterly Journal of Economics*, qjaf013.

Boesen, Ulrik. 2021. "How Are Your State's Roads Funded?" Accessed: 2025-01-23.

Bollinger, Christopher R, and Amitabh Chandra. 2005. "Iatrogenic specification error: A cautionary tale of cleaning data." *Journal of Labor Economics*, 23(2): 235–257.

Boudot-Reddy, Camille, and André Butler. 2024. "Paving the road to re-election." *Journal of Public Economics*, 239: 105228.

Brasington, David M. 2002. "Edge versus center: Finding common ground in the capitalization debate." *Journal of Urban Economics*, 52(3): 524–541.

Brasington, David M. 2017. "School spending and new construction." *Regional Science and Urban Economics*, 63: 76–84.

Brewer, Ethan, Jason Lin, Peter Kemper, John Hennin, and Dan Runfola. 2021. "Predicting road quality using high resolution satellite imagery: A transfer learning approach." *Plos one*, 16(7): e0253370.

Buchanan, James M., and Gordon Tullock. 1962. *The Calculus of Consent: Logical Foundations of Constitutional Democracy.* Ann Arbor:University of Michigan Press.

Calonico, Sebastian, Matias D Cattaneo, Max H Farrell, and Rocío Titiunik. 2017. "rdrobust: Software for regression-discontinuity designs." *The Stata Journal*, 17(2): 372–404.

Calonico, Sebastian, Matias D Cattaneo, Max H Farrell, and Rocío Titiunik. 2019. "Regression discontinuity designs using covariates." *Review of Economics and Statistics*, 101(3): 442–451.

Cattaneo, Matias D, Michael Jansson, and Xinwei Ma. 2020. "Simple local polynomial density estimators." *Journal of the American Statistical Association*, 115(531): 1449–1455.

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

Chaurey, Ritam, and Duong Trung Le. 2022. "Infrastructure maintenance and rural economic activity: Evidence from India." *Journal of Public Economics*, 214: 104725.

City of Hudson. 2020. "Asphalt Road Resurfacing Frequently Asked Questions (FAQ's)." Accessed: 2025-05-19.

Currier, Lauren, Edward L Glaeser, and Gabriel E Kreindler. 2023. "Infrastructure Inequality: Who Pays the Cost of Road Roughness?" National Bureau of Economic Research Working Paper w31981.

Diao, Mi, Devon Leonard, and Tien Foo Sing. 2017. "Spatial-difference-in-differences models for impact of new mass rapid transit line on private housing values." *Regional Science and Urban Economics*, 67: 64–77.

Donaldson, Dave, and Richard Hornbeck. 2016. "Railroads and American economic growth: A "market access" approach." *The Quarterly Journal of Economics*, 131(2): 799–858.

Dosovitskiy, Alexey. 2020. "An image is worth 16x16 words: Transformers for image recognition at scale." *arXiv preprint arXiv:2010.11929*.

Duranton, Gilles, and Matthew A Turner. 2011. "The fundamental law of road congestion: Evidence from US cities." *American Economic Review*, 101(6): 2616–2652.

Edel, Matthew, and Elliott Sclar. 1974. "Taxes, spending, and property values: Supply adjustment in a Tiebout-Oates model." *Journal of Political Economy*, 82(5): 941–954.

Finkelstein, Amy, and Nathaniel Hendren. 2020. "Welfare analysis meets causal inference." *Journal of Economic Perspectives*, 34(4): 146–167.

Frandsen, Brigham R., Markus Frölich, and Blaise Melly. 2012. "Quantile treatment effects in the regression discontinuity design." *Journal of Econometrics*, 168(2): 382–395.

Fretz, Stephan, Raphaël Parchet, and Frédéric Robert-Nicoud. 2022. "Highways, market access and spatial sorting." *The Economic Journal*, 132(643): 1011–1036.

Gertler, Paul J, Marco Gonzalez-Navarro, Tadeja Gračner, and Alexander D Rothenberg. 2024. "Road maintenance and local economic development: Evidence from Indonesia's highways." *Journal of Urban Economics*, 143: 103687.

Ghani, Ejaz, Arti Grover Goswami, and William R Kerr. 2016. "Highway to success: The impact of the Golden Quadrilateral project for the location and performance of Indian manufacturing." *The Economic Journal*, 126(591): 317–357.

Gibbons, Stephen, and Stephen Machin. 2005. "Valuing rail access using transport innovations." *Journal of Urban Economics*, 57(1): 148–169.

Gibbons, Stephen, Teemu Lyytikäinen, Henry G Overman, and Rosa Sanchis-Guarner. 2019. "New road infrastructure: The effects on firms." *Journal of Urban Economics*, 110: 35–50.

González, Felipe, Josepa Miquel-Florensa, Mounu Prem, and Stéphane Straub. 2025. "The dark side of infrastructure: Roads, repression and land in authoritarian Paraguay." *The Economic Journal*, 135(666): 653–669.

Gonzalez-Navarro, Marco, and Climent Quintana-Domeque. 2016. "Paving streets for the poor: Experimental analysis of infrastructure effects." *Review of Economics and Statistics*, 98(2): 254–267.

Gonzalez-Navarro, Marcos, R David Zarate, Remi Jedwab, and Nick Tsivanidis. 2023. "Land transport infrastructure." *VoxDevLit*, 9(1): 3.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A unified welfare analysis of government policies." *The Quarterly journal of economics*, 135(3): 1209–1318.

Hoogendoorn, Sander, Jochem van Gemeren, Paul Verstraten, and Kees Folmer. 2019. "House prices and accessibility: Evidence from a quasi-experiment in transport infrastructure." *Journal of Economic Geography*, 19(1): 57–87.

Hsu, Yu-Chin, and Shu Shen. 2024. "Dynamic regression discontinuity under treatment effect heterogeneity." *Quantitative Economics*, 15(4): 1035–1064.

Huang, Yiping. 2016. "Understanding China's Belt & Road initiative: motivation, framework and assessment." *China Economic Review*, 40: 314–321.

Kohlhase, Janet E. 1991. "The impact of toxic waste sites on housing values." *Journal of Urban Economics*, 30(1): 1–26.

Levkovich, Oleg, Jan Rouwendal, and Rolf Van Marwijk. 2016. "The effects of highway development on housing prices." *Transportation*, 43(2): 379–405.

Li, Shanjun, Jun Yang, Ping Qin, and Shouyang Chonabayashi. 2016. "Wheels of fortune: Subway expansion and property values in Beijing." *Journal of Regional Science*, 56(5): 792–813.

List, John A. 2022. *The voltage effect: How to make good ideas great and great ideas scale.* Crown Currency.

Oates, Wallace 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.

Oates, Wallace E. 1972. *Fiscal Federalism.* New York:Harcourt Brace Jovanovich.

Ohio Department of Taxation. 2020. "Real Property Tax – General." Accessed: 2025-01-30.

Ohio Department of Transportation. 2023. "2023 Annual Report." Accessed: 2024-01-23.

Peter G. Peterson Foundation. 2024. "The Highway Trust Fund Explained." Accessed: 2025-01-23.

Public Service Department of Beavercreek Township. 2025. "Road funding & maintenance in Beavercreek Township, Ohio." *Personal communication*, Email to Saani Rawat, January 21, 2025.

Rioja, Felix K. 2003. "Filling potholes: Macroeconomic effects of Maintenance versus New Investments in public infrastructure." *Journal of Public Economics*, 87(9-10): 2281–2304.

Theisen, Theis, and Anne Wenche Emblem. 2021. "The road to higher prices: Will improved road standards lead to higher housing prices?" *The Journal of Real Estate Finance and Economics*, 62: 258–282.

U.S. Department of Transportation. 2022. "The Bipartisan Infrastructure Law Will Deliver for Ohio." Accessed: 2025-01-23.

Vaswani, A. 2017. "Attention is all you need." *Advances in Neural Information Processing Systems*.

Wong, Ho Lun, Yu Wang, Renfu Luo, Linxiu Zhang, and Scott Rozelle. 2017. "Local governance and the quality of local infrastructure: Evidence from village road projects in rural China." *Journal of Public Economics*, 152: 119–132.

Woo, Sanghyun, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. 2023. "Convnext v2: Co-designing and scaling convnets with masked autoencoders." 16133–16142.

Chapter 2

Single-family Homes and Reinvestment: Variation by Ownership Type

2.1 Introduction

Literature from the 1980s and 1990s (Galster, 1983; Shilling, Sirmans and Dombrow, 1991; Gatzlaff, Green and Ling, 1998) generated what has become a stylized fact in the mind of researchers and practitioners; namely, owners of single-family homes that also occupy them reinvest more in single-family homes than owners that rent out single-family homes. Despite being viewed as a settled question, this question has garnered renewed relevance given recent shifts in the U.S. single-family housing market. Over the past decade, large landlords have substantially expanded their holdings of single-family rentals (SFRs), prompting widespread policy interest in whether the size and ownership structure of the holdings of SFRs will affect the reinvestment of the single-family housing stock.

The increased attention in the literature on holdings of single-family homes by institutional or large investors has focused on a number of questions, including how such investors manage their portfolio and what impacts their ownership has both on households and the communities in which those homes are located (Gorback, Qian and Zhu, 2025; Polimeni and An, 2024; Giacoletti et al., 2025; Coven, 2023). For example, Billings and Soliman (2023) observe that homes in neighborhoods with more investor purchases of single-family homes pull fewer building permits themselves; however, they do not measure associations between ownership of a given property and reinvestment into that particular property. Finally, An et al. (2024) observes that properties owned by larger portfolio owners more frequently experience code violations.

The older literature on reinvestment and maintenance of single-family homes also suffers from a lack of direct test of reinvestment in single-family homes with different tenure types. Other than Galster (1983), which draws on a survey conducted in Wooster, Ohio, research like Shilling, Sirmans and Dombrow (1991) and Gatzlaff, Green and Ling (1998) estimate models that focus differences in the resale values of otherwise similar single-family homes that are renter and owner-occupied. Both attribute these differences in price appreciation to differences in reinvestment, without measuring reinvestment of properties directly. Collectively, both the older and more recent literature study what is effectively a long-term consequence of differential reinvestment and depreciation.

This study proposes a direct test of whether home reinvestment differs across ownership and resident types using tax parcel data and building permits from two American Cities: Minneapolis, Minnesota, and Charlotte, North Carolina. In addition to looking at differences between owners and renters, we investigate differences between different types of landlords based on the size of their holdings. To our knowledge, this paper thus constitutes the first study to measure both reinvestment activity and tenure status in two major metropolitan areas without relying on proxies for either variable of interest. This study is timely given the policy backdrop of increased prevalence of single-family houses being renter-occupied and the growth of institutional ownership across the U.S (Naamane, 2024). This increase in concentration of ownership has changed property transaction prices and thus who can access neighborhoods (Coven, 2023; Polimeni and An, 2024), and led to decreases in local property tax receipts (Austin, 2022).

Our findings demonstrate substantial differences in reinvestment by tenure status and by landlord size. In Minneapolis where we have the best data, owner-occupied single-family homes file 26% more building permits and invest 55% more in permitted work than otherwise comparable single-family rental units. Moreover, we detect a strong inverse relationship between the size of a landlord's portfolio and the degree of reinvestment; specifically, larger landlords file 40% fewer permits and invest 45% less than smaller landlords. Notably, we find smaller statistically significant difference in the volume of plumbing permits across these groups relative to building permits, suggesting that reactive work (e.g., urgent repairs) is less likely to be cut than proactive work (e.g., roof inspections), and hence, showing that large landlords tend to 'procrastinate' on home reinvestment. Even with much less coverage of permits in Charlotte, the results affirm the inverse relationship between the size of landlord holdings and the level of permit and reinvestment activity in single-family homes. These results point toward the possibility that ongoing shifts in single-family rental ownership - particularly the growth in SFHs owned by large landlords, could reduce reinvestment in the existing housing stock, with implications for property conditions and neighborhood trajectories.

2.2 Related Literature

Abundant literature exists on how reinvesting in existing homes staves off depreciation. Using a variety of different data sources and empirical approaches, Chinloy (1980), Harding, Rosenthal and Sirmans (2007) and Wilhelmsson (2008) find that resale value of homes reflect the amount of reinvestment performed on them, and that reinvestment can counteract depreciation of the housing stock. The amount of reinvestment performed on a home depends on a variety of factors including life-cycle considerations for the owner, the property, and the neighborhood (Galster, 1987, 174). Davidoff (2004) finds that all else held equal, elderly homeowners also spend reinvest less than do younger owners of similar homes, while Helms (2003) finds that expectations about a neighborhood's trajectory can influence reinvestment decisions. However, measuring differences in property reinvestment across tenure types is fraught by data limitations. While studies on other reinvestment-related topics employ either local building permit data or the American Housing Survey to measure reinvestment activity, a common limitation is that data on reinvestment performed specifically on rental housing units are seldom available. For example, Helms (2003) note that while their building permit data records detailed information on building maintenance, there is no systematic recording of which units are rental properties. Similarly, the American Housing Survey's survey instrument only asks owner-occupant respondents its questions regarding building maintenance.

A number of studies from the 1990s attempt to back into differences in reinvestment between rental and owner-occupied housing units by comparing price appreciation between sales. Both Shilling, Sirmans and Dombrow (1991) and Gatzlaff, Green and Ling (1998) find differences in how the resale values of otherwise similar single-family homes develop over time between renter and owner-occupied units using transaction data. Both attribute these differences in price appreciation to differences in reinvestment, without actually measuring reinvestment of properties. To this date, the only study on this question that explicitly relates measures of reinvestment to tenure status is Galster (1983). Drawing upon a survey conducted in Wooster, Ohio, Galster demonstrates that owner-occupied SFRs experience more reinvestment than rental SFRs. However, part of this effect may reflect higher incomes of owners rather than purely tenure-status related differences, as higher income homeowners also perform more reinvestment than homeowners with lower incomes.

An entirely separate literature focuses on relationships between ownership types and the condition of the housing stock. Perhaps inspired by Matt Desmond's 2016 book 'Evicted' (Desmond, 2016), a group of studies evaluates the associations between ownership type and dereliction of properties. Decker (2023) describes and documents the practice of "milking" properties i.e. landlords achieving financial returns by allowing properties to dilapidate while collecting rental income from them, finding that this practice is somewhat common in low-end small rental properties with no more than four units. Similarly, Shannon, Skobba and Durham (2023) find for a sample from rural Georgia that properties owned by landlords with five or more properties have a significantly greater risk of dilapidation compared to properties owned by landlords with fewer properties. Rose and Harris (2022) use data from both routine building inspections from Rochester, New York, to show that code violations are the least common in owner-occupied properties and are most common in properties owned by absentee landlords. Further, the authors suggest that this outcome results from systematic differences in reinvestment across types of ownership. Finally, An et al. (2024) observe that properties owned by larger scale owners more frequently experience code violations. Collectively, while not directly measuring the same thing as the 1980 and 1990s literature on reinvestment, these papers study what is effectively a long-term consequence of differential reinvestment and depreciation.

Most recently, following increased attention on holdings of single family homes by institutional or otherwise large investors, a new literature studies how such investors manage their portfolio, and what impacts their ownership has both on households and the communities in which those homes are located (Gorback, Qian and Zhu, 2025; Polimeni and An, 2024; Giacoletti et al., 2025; Coven, 2023). Most notably - using the same data from Charlotte, North Carolina as we employ in this study - Billings and Soliman (2023) observe that homes in neighborhoods with more investor purchases of single family homes pull fewer building permits themselves, however, they do not measure associations between ownership of a given property and reinvestment into that particular property.

Contribution. Taken together, these lines of inquiry underscore a persistent gap in how scholars measure and compare reinvestment across tenure status, and further across different forms of rental ownership. Existing work either uses indirect proxies (e.g., price appreciation) or focuses on homeowner self-reports, leaving large segments of the rental market under-examined. Our paper directly addresses this limitation by leveraging parcel-level data on building permits and parcel ownership records from two sizable metropolitan areas. By matching permit filings to property ownership information, we provide first micro-level assessment of reinvestment across different ownership categories. In doing so, we contribute fresh empirical insight into how changing landlord structures, especially the growth of large-scale owners, might systematically influence reinvestment and, ultimately, the long-term condition of the single-family housing stock.

2.3 Data

2.3.1 Sources

Our paper draws upon data from two study areas: Minneapolis, Minnesota, and Charlotte, North Carolina. For each study area, we draw upon two primary types of data: Permits that record any modifications to properties including reinvestment activity, and tax assessor data that report parcels' characteristics ownership histories. As we note below, the permit and tax assessor data are much more complete in Minneapolis for two reasons. First, permits are required for much smaller projects and many more project types in Minneapolis. Second, the homestead exemption is much more generous in Minneapolis than in most states, which incentivizes local government to ensure that properties are occupied by owners.

Permit Data The construction permits data for Minneapolis, Minnesota are from the City of Minneapolis's Construction Code Services office, available via the city's Open Data web site (OpenData Minneapolis, 2024*b*). Minneapolis requires building permits for more types of work than most other jurisdictions, exempting only minor construction or mechanical work, as well as flooring (City of Minneapolis, 2024).[23] As such, Minneapolis' construction permits are uniquely suited for the purposes of this study, as they create a log of all but the smallest home improvements performed on any structure within the city.

Our permits data for Minneapolis cover the time frame from January 2017 to December 2024, and include information on the applicant, on the type of work performed (both in categorical terms and via short description), on what property the work is to be performed, and the date a permit was issued.[24] For single-family homes, work is broken out into three

[23] Examples for construction work that is exempt from permit requirements include constructing decks, fencing, flooring, countertops, or installation of appliances. See, City of Minneapolis (2024).

[24] While the Minneapolis 2040 Plan and its associated upzoning does fall into the time frame covered by our data, we believe it is unlikely to affect reinvestment decisions for existing properties due to legal uncertainty surrounding the plan's implementation.

categories: Mechanical permits cover work such as heating, ventilation, air conditioning, refrigeration, and gas piping. Plumbing permits cover any installation or replacement of water fixtures, water piping, water heaters, backflows, or gas appliances such as stoves. Finally, building permits cover any other work performed on a house, such as roofing, window replacements, or remodels. Permits for any work beyond mechanical work or plumbing - approximately 40% of permits in our data - report the dollar value of the work performed.

The permits data for Charlotte spans from January 2004 to December 2023 and is provided by Mecklenburg County GIS Services Department (Mecklenburg County GIS, 2025*a*), covering all of the same fields as our Minneapolis data. Unfortunately, the permit data for Charlotte only includes building permits, as opposed to building, plumbing and mechanical permits available for Minneapolis. Moreover, compared to Minneapolis, building permits data for Charlotte cover a narrower range of types of work, as Mecklenburg County only requires building permits for more extensive forms of remodeling. One advantage of Charlotte data is that while Minneapolis' permit data spans eight full years, our Charlotte data covers a much longer time horizon.

Tax Assessor Parcel Data To establish ownership and housing tenure status for each parcel, we rely on annual tax parcels data. For Minneapolis, we take this data from Hennepin County Tax Assessor for each year from 2017 to 2024 (Minnesota Geospatial Commons, 2017, 2018, 2019, 2020, 2021, 2022, 2023, 2024)[25]. Since Hennepin County, Minnesota is only a small portion of the Minneapolis-Saint Paul metropolitan area we also use tax parcels data from the other six counties of that metropolitan area[26] for the years 2005-2024 in addition to that for Hennepin County for identifying large owners within the region.

[25] While our permit data begins in 2016, it does not cover the entire year. For this reason, we do not include permits from 2016 in our analysis.

[26] Specifically, we use tax parcels data for Anoka, Carver, Dakota, Hennepin, Ramsey, Scott, and Washington counties, Minnesota.

To construct a similar dataset for Mecklenburg County, North Carolina, we use tax parcel data covering the period from 2004 to 2023 (Mecklenburg County GIS, 2025*b*)[27]. These data contain information on all parcels within the county, encompassing ownership details, property characteristics, land use, property values and parcel boundaries. In addition to the aforementioned permit and parcel datasets, we collect data on property characteristics such as the size, bedrooms, bathrooms, heated area, parcel area above ground, and construction year of any given property, as well as which neighborhood a given property is located in, identified by different community, neighborhood variables as well as census tracts. For our Minneapolis data, we take this information from the Hennepin County Tax Assessor (OpenData Minneapolis, 2024*a*). For our Charlotte study area, property characteristics are reported by the tax assessor in the tax parcels data. In both study areas, we use spatial matching to identify the census block location for each property[28].

2.3.2 Data Processing

For our analyses, we compile the aforementioned sources into a dataset that links property ownership histories with reinvestment records as measured through parcel permits. In this dataset, each observation represents one parcel in one year. We begin this process by compiling ownership histories for each parcel from the annual tax parcel datasets. In Minneapolis, the same tax parcel data also reports for any given parcel-year whether the owner of that parcel claimed the homestead tax exemption, which we interpret as that parcel being owner-occupied in that year. The homestead exemption in Minneapolis reduces the taxable value by 40% up to $95,000 in value, which is one of the most generous homestead exemptions in the US. Since Mecklenburg County data do not report homestead exemptions, we identify owner-occupants

[27]2024 data were available but had undergone major changes compared to previous decade so we did not consider it in our final analysis for consistency.

[28]All spatial matching is performed in *R* using the *sf* and *tigris* software packages. Since census blocks are perfect subsets of census blockgroups and census tracts, the same matching also tells us which blockgroup and tract a given property is located within.

in our Charlotte study area by flagging parcels for which the property's tax address - i.e. the address to which the property tax bill is mailed - is identical to the property's street address. This approach was previously employed by Ihlanfeldt (2021).

We use Federal Reserve Bank (FRB) of Minneapolis's definition of very large landlords and specified a rental property owner as a 'large landlord' if the number of single-family properties owned by the landlord exceeded a certain threshold (Ky and Starling, 2023). In the primary models we estimate, we use 20 as the threshold to identify large landlords, while using thresholds of 10, 50 and 100 to test the robustness of our findings and assess how reinvestment behavior changes across different scales of landlord ownership[29][30].

2.4 Empirical Strategy

First, we compare the simple means for the number of permits per parcel-year between the different types of housing tenure and ownership, and perform a simple Welch Two Sample Difference of Means t-test. While this does not account for systematic differences in upkeep—such as those based on demographic lines or neighborhood characteristics, as noted by Galster (1987, 174)—it provides a crude glimpse into whether any differences exist between groups.

Knowing that owner-occupiers, small landlords, and large landlords likely differ in both property and location characteristics, we next turn to regression-based models that can control for those differences. Our primary interest lies in modeling two related outcomes:

[29] Note that we exclude all permits for the construction of new properties or for the demolition of existing properties, as well as permits that have been canceled.

[30] We explored other ways to identify SFHs owned by large landlords. For example, Government Accountability Office (GAO) (Naamane, 2024).

1. **The count of permits** for each single-family property in a given year, modeled via a *negative binomial* (NB) regression. To capture over-dispersion in count data, we employ a negative binomial regression for the count of permits:

$$\ln(\mu_{it}) \;=\; \beta_1 \mathbf{1}_{\{\text{Owner-Occupied}_{it}\}} \;+\; \beta_2 \mathbf{1}_{\{\text{Large Landlord}_{it}\}} \;+\; \alpha_{c(i)} \;+\; \gamma_t \;+\; \mathbf{Z}'_{it}\boldsymbol{\beta}_5 \;+\; \varepsilon_{it},$$

   where $\ln(\mu_{it})$ is the expected number of permits for property i in year t. The variables $\mathbf{1}_{\{\text{Owner-Occupied}_{it}\}}$ and $\mathbf{1}_{\{\text{Large Landlord}_{it}\}}$ denote whether the property is owner-occupied or whether it is owned by a large landlord. Depending on the specification, we include *municipality*, *neighborhood*, or *community* fixed effects, denoted collectively here by $\alpha_{c(i)}$, and year fixed effects γ_t. The vector $\mathbf{Z}_{it}$ captures continuous property characteristics (e.g., bedrooms, baths, heated area, lot size, and age). We rely on standard negative binomial likelihood to estimate this model, as it appropriately handles over-dispersion in the permit counts.

2. **The permit value** (in dollars), using an *OLS* model on the log-transformed permit value:

$$\ln\Big(\text{Value}_{it}\Big) \;=\; \beta_1 \mathbf{1}_{\{\text{Owner-Occupied}_{it}\}} \;+\; \beta_2 \mathbf{1}_{\{\text{Large Landlord}_{it}\}} \;+\; \alpha_{c(i)} \;+\; \gamma_t \;+\; \mathbf{Z}'_{it}\boldsymbol{\beta}_5 \;+\; \eta_{it}.$$

In this framework, the key distinction is that some specifications use *municipality* fixed effects while others use *census tract* fixed effects (or different scales of geographic controls). Missing or zero permit values remain a limitation for Minneapolis data: some projects may not have a recorded value, and large-scale or formal projects might be more likely to specify costs. Hence, interpretation of these permit-value results should be mindful of potential non-random missingness. In contrast, the Charlotte data exhibit a more complete and consistently reported record of permit values.

In both the NB count model and the OLS log-value model, we always include year fixed effects and the same set of housing characteristics ($\mathbf{Z}_{it}$). The difference lies in whether we control for *community/neighborhood/census tract* fixed effects, as well as in how we define a 'large landlord' (e.g., 10+, 20+, 50+, or 100+ properties). This allows us to separately identify how rental units differ from owner-occupied ones, and how large-scale landlords differ from smaller investors, while properly controlling for the location- and parcel-specific factors that could influence reinvestment.

2.5 Results

2.5.1 Single-Family-Home Characteristics

Table 10 reports mean property characteristics by ownership type for single-family homes in Minneapolis (Panel A) and Charlotte (Panel B). In Minneapolis, owner-occupied homes tend to have slightly larger lot sizes than those owned by small landlords, while properties owned by large landlords (20+ properties) stand out for having smaller lot sizes, fewer bathrooms, and somewhat older buildings. These differences are not dramatic but suggest that large landlords may be drawn to slightly smaller or older housing stock, potentially shaping the extent and type of reinvestment they undertake. In Charlotte, owner-occupied parcels are smaller in lot size than those owned by small landlords (16,376 vs. 22,276 square feet), but they maintain a larger heated area (2,357 vs. 2,164 square feet). Rentals owned by small landlords also have more bedrooms and bathrooms on average than those owned by large landlords, suggesting that large landlords in Charlotte, as observed for Minneapolis, gravitate towards smaller houses with smaller lot sizes, fewer bedrooms and bathrooms, although the age of SFHs owned by large landlords is younger than those owned by owners or small landlords. Overall, these descriptive patterns align for Minneapolis and Charlotte, indicating

Table 10. Mean Property Characteristics by Owner Type

Panel A: Minneapolis

Ownership Type	**Owner Occupied**	**Rental, $<$ 20 Properties**	**Rental, $\geq$ 20 Properties**
Lot Size (sqft)	6,030	5,891	5,505
	(1,985)	(2,229)	(1,210)
Building Size (sqft)	1,404	1,382	1,249
	(572)	(726)	(326)
Bedrooms	3.06	3.13	3.26
	(0.92)	(1.12)	(0.99)
Bathrooms	1.87	1.84	1.45
	(0.91)	(1.00)	(0.63)
Building Age	92.29	94.13	91.71
	(24.14)	(27.03)	(28.46)
Number of Properties	62,463	11,159	1,229

Panel B: Charlotte

Ownership Type	**Owner Occupied**	**Rental, $<$ 20**	**Rental, $\geq$ 20**
Lot Size (sqft)	16,376	22,276	11,698
	(20,766)	(68,435)	(73,507)
Building Size (sqft)	2,357	2,164	1,799
	(1,082)	(1,153)	(587)
Bedrooms	3.50	3.31	3.20
	(1.14)	(0.89)	(0.68)
Bathrooms	2.23	2.10	1.94
	(0.79)	(0.86)	(0.47)
Building Age	33.7	38.1	28.6
	(20.7)	(26.8)	(19.6)
Number of Properties	158,802	89,171	20,584

Notes: Means are computed for property characteristics of single-family homes (SFHs). Standard deviations appear in parentheses. For Minneapolis, Owner-occupied parcels are identified via homestead exemption. For Charlotte, Owner-occupied parcels are identified where the tax-bill mailing address matches the property address. "Rental, $<$ 20" indicates small landlords with fewer than 20 single-family homes (SFHs); "Rental, $\geq$ 20" indicates large landlords with 20+ SFHs. Lot size is parcel area (sqft); Building size is heated living area (sqft). Minneapolis statistics refer to 2024. Charlotte statistics refer to 2023.

that properties owned by large landlords often differ from both owner-occupied homes and from small-scale rental holdings. In our analysis, we control for these differences in properties to isolate the difference in home reinvestment coming from ownership type, regardless of the variation in other property characteristics.

2.5.2 Spatial Distribution of Single-Family Homes

Panel (A) of Figure 10 shows the spatial distribution of single-family homes (SFHs) in Minneapolis by ownership type. Figure 10a illustrates the broad coverage of owner-occupied versus rental properties, revealing that owner-occupied units dominate many southern neighborhoods, while rentals are loosely interspersed throughout the city, with greater proximity to central Minneapolis or downtown compared to owner-occupied properties. Figure 10b highlights SFHs owned by large landlords, defined under various thresholds in Section 2.3. These SFHs appear less numerous but exhibit pockets of concentration in select areas, such as the north-western part of the city. Together, these maps underscore that while the rental share is sizable, large-landlord ownership remains a smaller but distinct segment of the Minneapolis SFH market.

Panel (B) of Figure 10 shows the spatial distribution of single-family homes (SFHs) in Charlotte by ownership type. Figure 10c shows the distribution of SFHs by tenure status, revealing a broad mix of owner-occupants and rentals across the city. Figure 10d focuses on SFHs owned by large landlords, illustrating that such holdings are much more dispersed in Mecklenburg County.

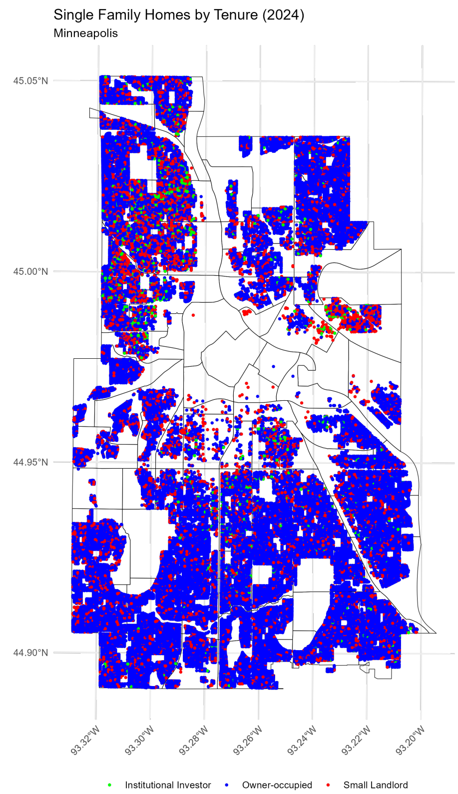

(a) SFHs by Tenure status

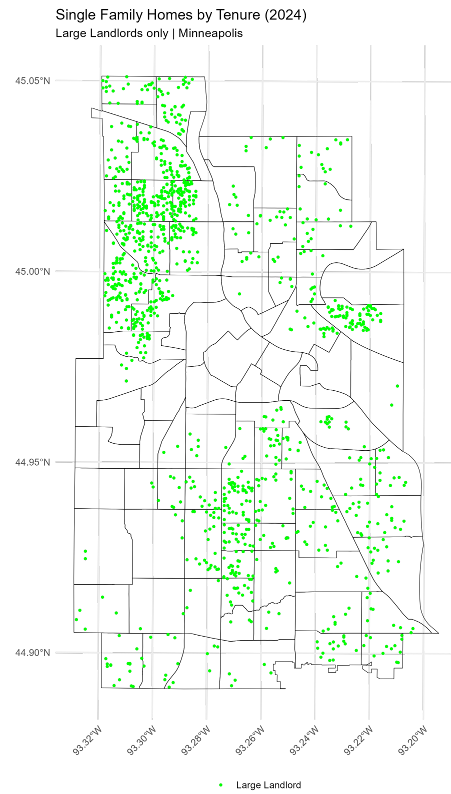

(b) SFHs owned by large landlords

Panel A: Minneapolis

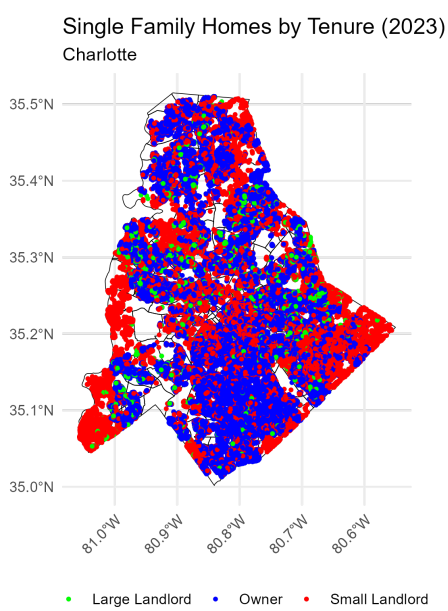

(c) SFHs by Tenure status

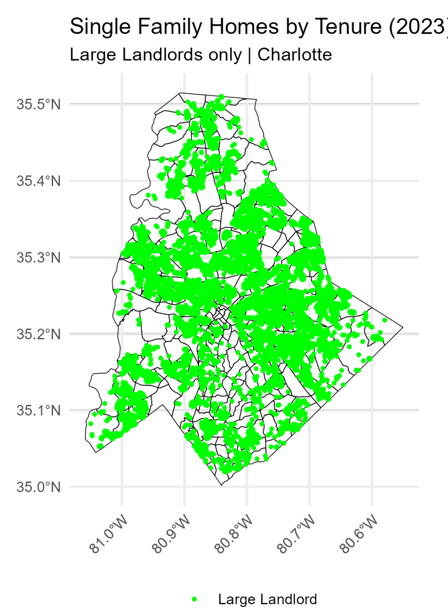

(d) SFHs owned by large landlords

Panel B: Charlotte

Figure 10. Single-Family Homes (SFH) in Minneapolis and Charlotte

2.5.3 Large Landlord Ownership over Time

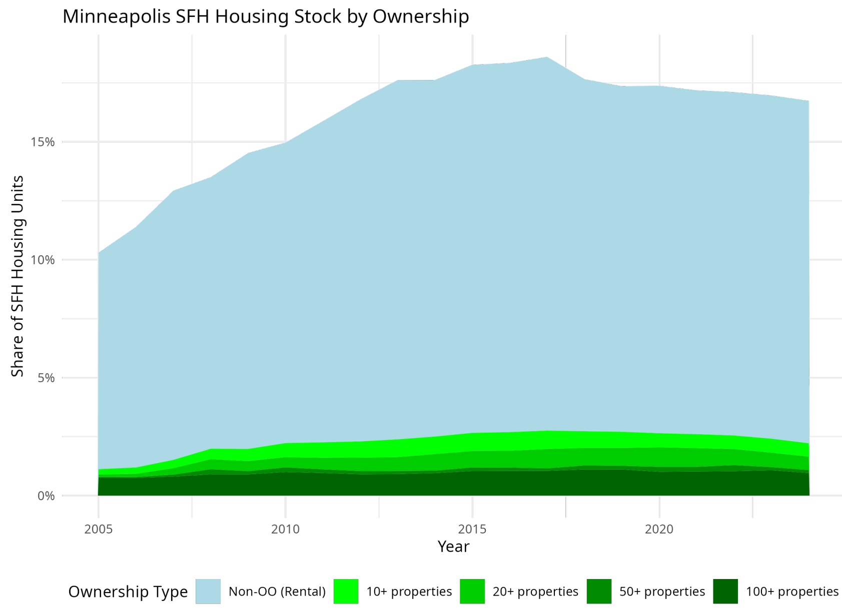

(a) Minneapolis: 2005–2025

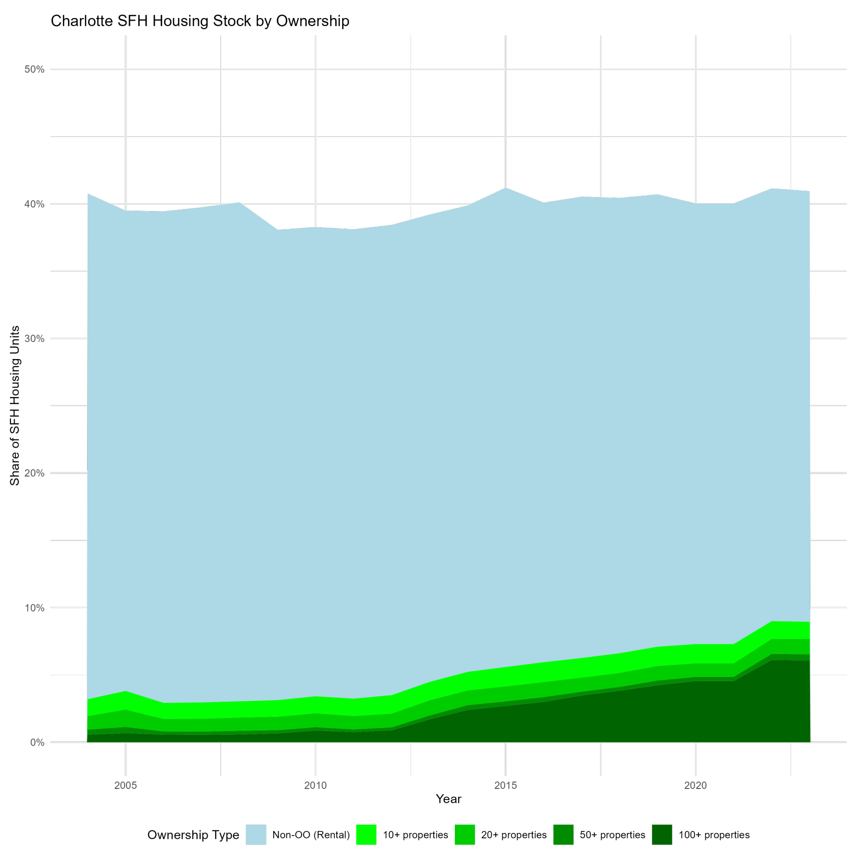

(b) Charlotte: 2004–2023

Figure 11. Large-Landlord Ownership over time in Minneapolis and Charlotte

Figure 11 traces the share of single-family homes (SFHs) held by large rental portfolios in Minneapolis (2005–2024) and Charlotte (2004–2023). Across both markets the trajectory is unmistakably upward, yet the pace and scale differ in ways that foreshadow our later regression results.

In Minneapolis, the overall share of non-owner occupied single-family properties has grown over time. Moreover, beginning from a small base in 2005, the proportion of SFHs owned by landlords with greater than 10 properties more than doubles by 2024. Growth is most pronounced between 2020 and 2022, coinciding with the pandemic-era surge in single-family rental demand and cheap mortgage credit. Even at the end of the sample, however, very large owners (with more than 50 or 100 properties) still account for only a small sliver—roughly 1–2 percent—of the city's detached-housing stock. The slow, steady climb suggests a market where large investors are active but have yet to displace the traditional mix of owner-occupants and small landlords.

Charlotte's time series reveals a steeper expansion. Although the share of non-owner occupied single family properties remains relatively flat, large-landlord penetration accelerates in the wake of the Great Financial Crisis, pauses mid-decade, and then surges again after 2018. By 2023, owners with portfolios of more than 20 properties control close to nine percent of all SFHs, and even the *more than 100-property* group grows substantially over this time period. The pattern aligns with Charlotte's faster population growth and the city's attractiveness to institutional single-family REITs, as noted in a recent GAO report (Naamane, 2024).

Overall, Charlotte's descriptive patterns are consistent with the study's broader framing: the prevalence of single-family rental housing has grown, led in part by large landlords whose portfolios now comprise a noteworthy share of the market. Coupled with the observed variation in building age and lot size, these data provide essential context for our analysis of how tenure status and landlord size relate to the quantity and dollar value of reinvestment in SFHs.

2.5.4 Descriptive Evidence

Table 11. Differences in Mean Permit Activity for Single-Family Homes

	Mean		Difference	
	(1)	(2)	(1)–(2)	t-stat
Panel A: Minneapolis				
Owner-Occupied vs. Rental Properties				
	Owner	**Rental**		
Any permit count	0.343	0.270	0.073***	30.94
Mechanical permit count	0.072	0.058	0.014***	15.30
Plumbing permit count	0.136	0.117	0.019***	14.23
Building permit count	0.135	0.093	0.042***	35.44
Building permit value (USD)	2,397	1,780	617***	11.08
Rental Properties: Small vs. Large Landlords				
	Small	**Large**		
Any permit count	0.283	0.163	0.120***	23.56
Mechanical permit count	0.062	0.031	0.031***	16.35
Plumbing permit count	0.122	0.084	0.038***	12.34
Building permit count	0.098	0.049	0.049***	19.47
Building permit value (USD)	1,919	674	1,245***	13.19
Panel B: Charlotte				
Owner-Occupied vs. Rental Properties				
	Owner	**Rental**		
Building permit count	0.0121	0.0115	0.0006***	5.51
Building permit value (USD)	418	516	–98***	–8.13
Rental Properties: Small vs. Large Landlords				
	Small	**Large**		
Building permit count	0.0121	0.0061	0.0060***	26.83
Building permit value (USD)	557	144	412***	24.74

Notes: Means are annualised at the parcel–year level. Panel A uses Minneapolis permitting data, which include mechanical, plumbing, and building permits. Panel B uses Charlotte data, where only building-permit information is available. In each city, Panel A compares owner-occupied with rental properties, while Panel B compares small landlords (fewer than 20 parcels) with large landlords (20+ parcels). Differences are computed as column (1) minus column (2). Welch two-sample t-statistics are reported in the final column. *** $p < 0.01$.

Table 11 reports differences in mean annual permit activity for single-family homes in Minneapolis (Panel A) and Charlotte (Panel B). In Minneapolis, where detailed permitting records are available, owner-occupied properties are consistently more likely to show activity across all permit categories—mechanical, plumbing, and building—than rentals. For example, owners average 0.343 permits per parcel–year compared with 0.270 among rentals, with significant differences across categories. Within the rental sector, small landlords (fewer than 20 properties) exhibit significantly higher permitting activity than large landlords (20+ properties), both in terms of permit counts and permit values. This pattern suggests that large landlords reinvest less frequently and at lower dollar amounts, even after normalizing by property–year. As noted previously, permit data has a much smaller coverage, yet we observe a similar gradient in permit activity by the size of landlord holdings. Owner-occupied homes pull slightly more building-permit, though their permit values are lower on average than those of rental properties. At the same time, small landlords invest more heavily than larger landlords. Small-landlords apply for nearly twice as many (0.0121 vs 0.0061) and spend $557 per year compared to large-landlord parcels, who only spend $144 per year. Overall, Charlotte confirms that ownership structure matters for housing reinvestment, but the city's permit regime leads to far lower levels of observable activity.

2.5.5 Regression Results

Table 12. Home Reinvestment by Tenure Status and Landlord Scale: Minneapolis and Charlotte

	Dependent variable:			
	Mechanical Permits (Count)	Plumbing Permits (Count)	Building Permits (Count)	Building Permits (Dollars)
	negative binomial	*negative binomial*	*negative binomial*	*OLS*
	(1)	(2)	(3)	(4)
Panel A: Minneapolis				
Owner-Occupied	0.069***	−0.002	0.227***	0.437***
	(0.015)	(0.011)	(0.012)	(0.024)
Large Landlord (20+)	−0.471***	−0.172***	−0.503***	−0.595***
	(0.054)	(0.034)	(0.043)	(0.064)
Housing Characteristics	✓	✓	✓	✓
Neighborhood F.E.	✓	✓	✓	✓
Year F.E.	✓	✓	✓	✓
Observations	597,248	597,248	597,248	597,248
Panel B: Charlotte				
Owner-Occupied			−0.045***	−0.005**
			(0.011)	(0.002)
Large Landlord (20+)			−0.429***	−0.075***
			(0.033)	(0.006)
Housing Characteristics			✓	✓
Tract Group 1 F.E.			✓	✓
Year F.E.			✓	✓
Observations			4,123,141	4,123,141

Notes: Each column reports regressions of permitting outcomes on tenure status and landlord scale, using small landlord category as the reference group and controlling for housing characteristics where indicated. Housing characteristics include bedrooms, full bathrooms, heated area, lot size, age, and age squared. Minneapolis (Panel A) includes mechanical, plumbing, and building permits (counts via negative binomial; value via OLS). Charlotte (Panel B) has only building permits. Minneapolis models include neighborhood and year fixed effects; Charlotte models include Tract Group 1 and year fixed effects. Census Tract Group 1 fixed effects are defined by grouping tracts by the leading five digits of the six-digit census tract identifiers. Standard errors are in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12 presents results from our regression models. In Panel A, we show the differences in reinvestment by ownership type for Minneapolis. Comparing owner-occupied SFHs with similar single-family rentals suggests that owner-occupied units receive more reinvestment, though this is not true for all types of reinvestment. After controlling for housing characteristics, year and neighborhood of a SFH, owner-occupied single family homes have a 7.1%[31] higher expected count rate for mechanical permits than otherwise similar rental single-family homes, and a 25.5%[32, 33] higher rate of building-permit filings. Furthermore, building permits for owner-occupied units represent approximately 54.8%[34] greater annual dollar values of work performed than for rental units, all else held equal[35]. Comparing large-landlord owned single-family homes to otherwise similar units not owned by large landlords, homes with large-landlord ownership have a 39.5%[36] lower expected count of building permits with 44.8%[37] lower annual dollar values for building permits, all else held equal. Comparing owner-occupied to rental SFHs, there is no statistical difference in filings for plumbing permits after controlling for property and neighborhood characteristics. Although we do find some differences in plumbing permit filings between small and large landlords, these differences are much lower in magnitude compared to other permit types.

In Panel B, we show results for Charlotte using two specifications that are directly comparable to the Minneapolis building-permit outcomes: Column (3) reports a negative binomial model of building-permit counts with Census Tract Group 1 fixed effects (FE), and Column (4) reports an OLS model of natural log of building-permit values. As in Minneapolis, we are interested in whether owner-occupied single-family homes exhibit systematically different reinvestment behavior than similar rental properties, and whether large landlords

[31] $\exp(0.069) - 1 = 0.071$, where 0.069 comes from Table 12, Panel A, Col. (1).

[32] $\exp(0.227) - 1 = 0.255$, where 0.227 comes from Table 12, Panel A, Col. (3).

[33] This translates to about 20.3% fewer building-permit counts for rental SFHs relative to owner-occupied homes, since $\frac{25.5\%}{1+25.5\%} \approx 20.3\%$.

[34] $\exp(0.437) - 1 = 0.548$, where 0.437 comes from Table 12, Panel A, Col. (4).

[35] This corresponds to about 35.4% lower building-permit values for rental SFHs relative to owner-occupied homes, computed as $\dfrac{54.8\%}{1 + 54.8\%} \approx 35.4\%$.

[36] $\exp(-0.503) - 1 = -0.395$, where -0.503 comes from Table 12, Panel A, Col. (3).

[37] $\exp(-0.595) - 1 = -0.448$, where -0.595 comes from Table 12, Panel A, Col. (4).

differ from smaller-scale rental owners. We first focus on differences in reinvestment between rental and owner-occupied single-family homes. In Column (3), the negative-binomial coefficient on *Owner-Occupied* is -0.045 and statistically significant. Column (4) reports natural log of building permit values and finds that owner-occupied properties have a 0.5%[38] lower total dollar value of permitted work than rental properties. Next, Table 12 highlights reinvestment differences for large landlords in Charlotte. The coefficients are negative and statistically significant: *Large Landlord (20+)* shows a coefficient of -0.429 in Column (3), implying that large landlords have about a 34.9%[39] lower expected count of building permits than smaller landlords, and a coefficient of -0.075 in Column (4), implying roughly a 7.2%[40] lower total dollar value of permitted work[41].

Additional Tests One salient feature of the process of re-investment is the fact that the largest single year of permit activity is the year of sale (Figure 26). In the first column of Appendix Table 32, we include a Sale Year indicator which confirms the elevated rates of permit activity in Minneapolis across all permit types. The inclusion of this variable also does increase the coefficient for owner occupied housing to indicate that owner occupiers pull more permits for all reinvestment activity. As before, the largest differences in pulling permits are for what we have termed proactive reinvestment activity.

In column 2 of Table 32, we interact Sale Year with our categorical variable for owner-occupied and large landlord owners of single-family homes. These results indicate that even though large landlords pull fewer permits overall, they are more likely to pull permits in the year of sale. This is likely due to the fact any reinvestment activity is easier when the house is empty. For owner occupiers, the results indicate that they pull fewer permits in the sale

[38] $e^{-0.005} - 1 \approx -0.5\%$, where -0.005 comes from Table 12, Panel B, Col. (4).

[39] $e^{-0.429} - 1 \approx -34.9\%$, Table 12, Panel B, Col. (3).

[40] $e^{-0.075} - 1 \approx -7.2\%$, Table 12, Panel B, Col. (4).

[41] Table 34 in the Appendix presents our results using alternative thresholds to define large landlords—10+, 50+, and 100+ properties—the coefficients remain negative and statistically significant, indicating fewer permits and lower dollar values for these owners.

year than small investors; presumably, this is because there is a smaller need for reinvestment in the sale year if they are reinvesting more in their homes throughout the years of ownership.

In the main results, we used a threshold of holding 20+ properties to differentiate between large and small landlords. However, much of the recent policy discussion concerned with impacts on the single family rental market focuses on landlords that own 100 or more properties. Because Minneapolis has a much smaller sample of landlords holding more than 50 or 100 properties, we focus this analysis on Mecklenburg County. As is evident in Tables 33 and 34 in the Appendix, there is a monotonic relationship between the number of holdings and the amount of permit activities. Consistent with the main results of the paper, the largest landlords pull the fewest permits and re-invest the smallest amounts in each single-family property.

While we control for the age of the property, if some landlords focus on newer properties in particular neighborhoods, then there could be a relationship between age and the neighborhood level fixed effects that is non-linear. As shown in Figure 27 in the Appendix, we first confirm that homes with ages below the median require much few permits than those that are in the oldest parts of the distribution. We find that no matter the part of the age distribution, owner occupiers consistently reinvest more in their homes. Further, the main results are robust to different levels of geographic specificity in the neighborhood fixed effects[42].

2.6 Discussion and Conclusion

This paper investigates whether there are systematic differences in how the housing stock is maintained across different types of ownership and housing tenure, using novel datasets that combine parcel-level ownership records with reinvestment records for the same buildings. We compiled two such datasets: one for all single-family homes in the city of Minneapolis,

[42]For details, see Table 35 in the Appendix.

Minnesota from 2017 to 2024, and another for Charlotte, North Carolina from 2004 to 2023. In both cities, we find that owner-occupiers perform significantly more work on their homes than landlords. In Minneapolis, owner-occupiers perform significantly more permitted work than otherwise comparable rental units. In Charlotte, the owner–renter gap is smaller and not statistically distinguishable from zero. Additionally, in both cities, we find that differences exist between small and larger landlords, where larger landlords perform less work requiring permits. Importantly, differences in reinvestment behavior cannot be explained solely by those properties' characteristics or the neighborhoods where single-family rentals are concentrated.

As we observe in Minneapolis - where data on multiple permit types are available - the greatest differences across owner and tenure types are in building permits. Importantly, we find little difference in the numbers of plumbing permits that are issued across single family homes owned by individual or large landlords. Such permits likely cover work that is arguably not optional for the livability of a unit; whereas investments that rejuvenate aging properties are more discretionary.[43] In Charlotte - where large owners are more prevalent - we find additional evidence that even among large owners, reinvestment into properties decreases with the number of properties owned.

The data employed in this study only cover a small fraction of all homes in the United States and only include work large enough to require permits under the respective rules of Minneapolis and Charlotte. However, our findings are consistent between the two study areas and are broadly consistent with the results of studies that use a different approach compared to ours - i.e., those Galster (1983) found via a survey or Shilling, Sirmans and Dombrow (1991); Gatzlaff, Green and Ling (1998) by systematically documenting differential rates of housing between owner-occupied homes and rental housing units. Further, unlike prior literature, our empirical design allows us to control for property and neighborhood characteristics, permitting us to rule out that differences in observed reinvestment are due to

[43]In Minneapolis, as is the case in most other cities, city codes spell out requirements for deeming a property inhabitable. Such requirements often specify minimum standards for plumbing, heating, and ventilation. Landlords who fail to comply with such requirements are required to provide renters with relocation assistance. See, City of Minneapolis (2026)

other observable differences between the types of units used as rentals or occupied by owners. Until the American Housing Survey (AHS) expands its set of questions about reinvestment to also cover rental units, this approach may be the most viable way to systematically answer our research question.

Coupled with the observation of an increase in prevalence of permanent single-family rentals and the increase in large-landlord ownership of SFHs – as documented by Naamane (2024); Ky, Starling and Tran (2021); Ky and Starling (2023) and observed within our own sample – our finding that rentals receive less reinvestment than otherwise similar owner-occupied SFRs implies that the aggregate amount of reinvestment performed across the housing stock may decline if the shift towards permanent rental units continues. Any such decline would bear consequences for the physical condition of the housing stock, as well as for employment among contractors, handymen, and other professionals involved in housing reinvestment work.

Policymakers may be able to counteract such trends via several policy levers that encourage reinvestment. First, with proactive code enforcement, cities may be able to identify under-maintained properties at an earlier stage of dilapidation. Second, cities and counties may be able to use fiscal policy measures such as tax incentives to encourage building reinvestment. At the same time, research needs to better understand the business models of various investors in single family homes. Presumably, all are interested in maximizing profit. While reinvesting less in a property is one channel to improve profits, resale of the portfolio of single-family rental properties is also an important channel. Accounting for these various mechanisms of generating return will provide insights into how policy makers might best intervene to encourage reinvestment among all investors.

References

An, Brian Y., Andrew Jakabovics, Anthony W. Orlando, Seva Rodnyansky,

and Eunjee Son. 2024. "Who Owns America? A Methodology for Identifying Landlords' Ownership Scale and the Implications for Targeted Code Enforcement." *Journal of the American Planning Association*, 90(4): 627–641. Publisher: Routledge _eprint: https://doi.org/10.1080/01944363.2023.2292674.

Austin, Neroli. 2022. "Keeping Up with the Blackstones: Institutional Investors and Gentrification."

Billings, Stephen B., and Adam Soliman. 2023. "The Social Spillovers of Homeownership: Evidence from Institutional Investors."

Chinloy, Peter. 1980. "The effect of maintenance expenditures on the measurement of depreciation in housing." *Journal of Urban Economics*, 8(1): 86–107.

City of Minneapolis. 2024. "Construction work exempt from permit."

City of Minneapolis. 2026. "Title 12 - HOUSING | Code of Ordinances | Minneapolis, MN | Municode Library."

Coven, Joshua. 2023. "The Impact of Institutional Investors on Homeownership and Neighborhood Access."

Davidoff, Thomas. 2004. "Maintenance and the Home Equity of the Elderly."

Decker, Nathaniel. 2023. "The Prevalence, Profitability, and Risks of Milking Among Low-End Small Rental Properties." *Housing Policy Debate*, 33(6): 1536–1553. Publisher: Routledge _eprint: https://doi.org/10.1080/10511482.2023.2210560.

Desmond, M. 2016. *Evicted: Poverty and Profit in the American City.* Crown Publishing Group.

Galster, George C. 1983. "Empirical Evidence on Cross-Tenure Differences in Home Maintenance and Conditions." *Land Economics*, 59(1): 107–113. Publisher: [Board of Regents of the University of Wisconsin System, University of Wisconsin Press].

Galster, George C. 1987. *Homeowners and Neighborhood Reinvestment.* Duke University Press. Google-Books-ID: jnyCNXIGYXMC.

Gatzlaff, Dean H., Richard K. Green, and David C. Ling. 1998. "Cross-Tenure Differences in Home Maintenance and Appreciation." *Land Economics*, 74(3): 328–342. Publisher: [Board of Regents of the University of Wisconsin System, University of Wisconsin Press].

Giacoletti, Marco, Rawley Heimer, Wenli Li, and Edison G. Yu. 2025. "Single-Family REITs and Local Housing Markets."

Gorback, Caitlin, Franklin Qian, and Zipei Zhu. 2025. "Impact of Institutional Owners on Housing Markets."

Harding, John P., Stuart S. Rosenthal, and C. F. Sirmans. 2007. "Depreciation of housing capital, maintenance, and house price inflation: Estimates from a repeat sales model." *Journal of Urban Economics*, 61(2): 193–217.

Helms, Andrew C. 2003. "Understanding gentrification: an empirical analysis of the determinants of urban housing renovation." *Journal of Urban Economics*, 54(3): 474–498.

Ihlanfeldt, Keith. 2021. "Property Tax Homestead Exemptions: An Analysis of the Variance in Take-Up Rates Across Neighborhoods." *National Tax Journal*, 74(2): 405–430. Publisher: The University of Chicago Press.

Ky, Kim-Eng, and Libby Starling. 2023. "Very large investors increase their share of the Twin Cities rental-home market." *Federal Reserve Bank of Minneapolis.*

Ky, Kim-Eng, Libby Starling, and Tu-Uyen Tran. 2021. "New property-data tool reveals patterns of investor ownership in the Twin Cities area." *Federal Reserve Bank of Minneapolis.*

Mecklenburg County GIS. 2025*a*. "Mecklenburg County Building Permits."

Mecklenburg County GIS. 2025*b*. "Mecklenburg County Tax Parcels."

Minnesota Geospatial Commons. 2017. "MetroGIS Regional Parcel Dataset - (Year End 2017) - Minnesota Geospatial Commons."

Minnesota Geospatial Commons. 2018. "Metro Regional Parcel Dataset - (Year End 2018) - Minnesota Geospatial Commons."

Minnesota Geospatial Commons. 2019. "MetroGIS Regional Parcel Dataset - (Year End 2018) - Minnesota Geospatial Commons."

Minnesota Geospatial Commons. 2020. "MetroGIS Regional Parcel Dataset - (Year End 2019) - Minnesota Geospatial Commons."

Minnesota Geospatial Commons. 2021. "MetroGIS Regional Parcel Dataset - (Year End 2020) - Minnesota Geospatial Commons."

Minnesota Geospatial Commons. 2022. "MetroGIS Regional Parcel Dataset - (Year End 2021) - Minnesota Geospatial Commons."

Minnesota Geospatial Commons. 2023. "Metro Regional Parcel Dataset - Year End 2023 - Minnesota Geospatial Commons."

Minnesota Geospatial Commons. 2024. "Metro Regional Parcel Dataset - Year End 2024 - Minnesota Geospatial Commons."

Naamane, Jill. 2024. "Rental Housing: Information on Institutional Investment in Single-Family Homes." Government Accountability Office GAO-24-106643.

OpenData Minneapolis. 2024*a*. "Assessing Department Parcel Data 2024."

OpenData Minneapolis. 2024*b*. "CCS Permits."

Polimeni, Nicholas, and Brian An. 2024. "Uncovering Neighborhood-level Portfolios of Corporate Single-Family Rental Holdings and Equity Loss."

Rose, Geoff, and Richard Harris. 2022. "The three tenures: A case of property maintenance." *Urban Studies*, 59(9): 1926–1943. Publisher: SAGE Publications Ltd.

Shannon, Jerry, Kim Skobba, and Jermaine Durham. 2023. "Landlords and Housing Quality in Rural Georgia: Assessing the Relationship." *Housing Policy Debate*, 0(0): 1–22. Publisher: Routledge _eprint: https://doi.org/10.1080/10511482.2023.2273345.

Shilling, James D., C. F. Sirmans, and Jonathan F. Dombrow. 1991. "Measuring depreciation in single-family rental and owner-occupied housing." *Journal of Housing Economics*, 1(4): 368–383.

Wilhelmsson, Mats. 2008. "House price depreciation rates and level of maintenance." *Journal of Housing Economics*, 17(1): 88–101.

Chapter 3

Does Local Urban Governance Status Matter? Evidence from India

3.1 Introduction

Urban local bodies (ULBs) are responsible for a wide range of public goods in developing countries such as schools, primary health facilities, sanitation, streets, lighting, and local regulation. Accordingly, whether and how local urban governance status matters for local economic development is a long-standing question, with implications for many rapidly urbanizing economies (Brollo et al., 2013; Faguet, 2012; Patrick and Mothorpe, 2017). However, a key empirical challenge is that governance structures are endogenous: places that are bigger and more developed are more likely to be recognized as urban towns and receive municipal status, while smaller or less developed settlements may not receive such recognition and therefore stagnate even further by being outside the urban system (Gadenne, 2017). As a result, simple cross-sectional comparisons in such contexts will conflate pre-existing differences with the effects of policy. Put differently, governance status is chosen rather than randomly assigned, making causal inference difficult. To address this challenge, researchers have often relied on varying identification strategies. For instance, Lee (2021) finds a positive effect of urban governance on access to water facilities in India by developing a satellite-based measure of confounders such as urbanization and also controlling for population density and other demographic variables in an OLS framework with fixed effects. By contrast, Mukhopadhyay (2017) finds no significant effect of governance status on public goods, such as water and sanitation, in India by comparing the provision of public goods across different types of urban settlements as categorized by their governance status. Meanwhile, although most papers study the impact on the extensive margin, more recently, Narasimhan and Weaver (2024) focus on the intensive margin of public goods provision, examining how governance status may have differential effects on public goods provision depending on the size of the settlement. More broadly, other studies have also explored this question in different contexts and using various methods (Mukhopadhyay, 2015; Mitra and Nagar, 2018), yet the empirical evidence remains mixed.

This paper leverages a sharp institutional feature in India to address this identification problem: the Census Town (CT) classification thresholds that determine eligibility for an area to be considered urban. The Indian Census defines Census Towns as settlements meeting all three criteria: (i) population of at least 5,000 (ii) population density of at least 400 per square kilometer and (iii) at least 75% of the male main workforce in non-agricultural activities. Crossing these multi-dimensional cutoffs sharply increases the probability that a settlement will be formally recognized by state governments as a Statutory Town (ST), i.e. granted a municipality or other ULB status, but notably, it *does not* guarantee it. We exploit this quasi-random variation in recognition likelihood around the thresholds to construct a local fuzzy regression discontinuity (RD) design. In effect, meeting the Census Town criteria in 2001 serves as an instrument for obtaining statutory recognition in 2011, allowing us to isolate the effect of local urban governance status on public goods offered by the settlements.

The core idea is a fuzzy multi-threshold RD with multiple running variables. Using rich settlement-level microdata from the Socioeconomic High-resolution Rural-Urban Geographic (SHRUG) platform (Asher, Novosad et al., 2019), we create a running variable called frontier distance using the three CT cutoffs and restrict attention to settlements near these thresholds. We first confirm a strong first-stage relationship: becoming marginally eligible as a Census Town in 2001 leads to a jump in the probability of statutory recognition in 2011. We then estimate the effect of statutory recognition, i.e., ULB status, on a broad set of development indicators measured by the 2011 Indian Census. Our outcome data are assembled by harmonizing the 2011 Village Directory (VD) and Town Directory (TD), which enumerate local infrastructure and amenities. We link these to demographic data from the 2011 Primary Census Abstract (PCA). This provides a comprehensive settlement-level panel of outcomes including educational facilities, health infrastructure, financial institutions, and other community amenities.

Preview of results. We find that crossing the urban eligibility thresholds leads to a pronounced increase in the likelihood of statutory recognition. Meeting all three Census Town criteria in 2001 raises the probability of obtaining statutory recognition by 7.1 percentage points in our preferred local specification, with a first-stage F-statistic of 18.05. Using CT eligibility as an instrument for statutory town status, we find large local average treatment effects on public goods provision: ULB status increases the number of government primary schools by 13.86 per settlement, middle schools by 7.72, and secondary schools by 4.89. We also find positive effects on health and financial infrastructure: statutory recognition increases hospitals by 2.53, family welfare centers by 3.00, cooperative banks by 4.09, and agricultural credit societies by 2.84. Community amenities also improve, with positive estimates for public libraries and reading rooms, while sports infrastructure declines by 5.71 facilities, consistent with land reallocation in formalizing settlements[44]. Results are robust to bandwidth choices, and robustness checks support the validity of the design. We find no evidence of manipulation around the thresholds and detect balance in baseline characteristics[45].

Our work builds on a growing literature that examines how local institutional capacity and governance shape development outcomes in low- and middle-income countries. Besley and Persson (2010) emphasize that investments in fiscal and legal capacity are key determinants of long-run development, and that variations in local administrative effectiveness can shape the quality of service delivery. Relatedly, Gadenne (2017) shows that greater fiscal autonomy can improve accountability and service delivery when citizens can observe how funds are used, while Burgess et al. (2015) highlights how political incentives affect the allocation of public goods across regions. We contribute to this literature by providing causal evidence on a different but complementary margin: the formal transition of villages to local urban governance status. Rather than focusing on local fiscal resources, we study how the administrative act of granting municipal status itself changes public goods provision. By leveraging the multi-threshold eligibility criteria for Census Town classification as a quasi-experiment, we isolate exogenous variation in the likelihood of statutory recognition and estimate its local

[44] Due to the small magnitude of the first-stage coefficient, the IV estimates scale reduced-form effects by a relatively large factor. The corresponding reduced-form estimates are directionally consistent and reported in the appendix.

[45] Our final regressions control for any remaining differences in these variables.

average treatment effect on infrastructure outcomes. This approach provides new evidence on how the extension of local urban governance status affects provision of local public goods in urbanizing economies.

This paper contributes to the literature via two main fronts. First, it adds to the literature by quantifying the effect of local urban governance status on local public goods provision in large developing economies. Although Denis and Marius-Gnanou (2010), Pradhan (2013) and Tewari (2020) have noted that many Census Towns remain under rural governance despite meeting Census's urban criteria, our study is the first to estimate the loss in public goods provision for these areas that are functionally urban but rurally governed. Second, by leveraging the multi-threshold criteria for Census Town classification, our paper uses a novel instrument and highlights the potential for using multi-dimensional cutoffs in a fuzzy RD design setting. This approach can be applied to other contexts where policy eligibility is determined by multiple threshold-based criteria, allowing researchers to exploit quasi-random variation in policy exposure.

3.2 Institutional Background

India's urban classification rests on two pillars. The first consists of settlements officially designated as urban by state governments via notification under municipal law. These include municipal corporations, municipalities, and nagar panchayats[46], and are each known as *Statutory Towns* (STs). The second consists of *Census Towns* (CTs), which are settlements that satisfy Census thresholds: total population at least 5,000; population density of at least 400 persons per km^2; and at least 75 percent of the male main workforce in non-agricultural activities. CTs are urban for statistical purposes but may continue under rural governance unless notified as STs by the state governments. Crossing the CT thresholds increases the

[46] A type of town council for smaller urban areas.

chance of statutory recognition but *does not guarantee it* as states differ in municipalization policy and timing. Some states have proactive policies to incorporate new urbanizing areas, while others delay or avoid creating new ULBs due to budgetary or political considerations.

3.3 Local Governance in India

The Nagarpalika Act, which was the 74th Amendment to the Indian Constitution enacted in 1992, provides the legal framework for urban local governance in India (Ministry of Home Affairs, Government of India, 2024). This constitutional amendment is supported by state-level municipal acts, such as Government of Goa, 1968; Government of Gujarat, 1963; Government of Karnataka, 1964; Government of Maharashtra, 1965; Government of Telangana, 2019 etc., that outline the specific procedures for forming Urban Local Bodies (ULBs). The entire process, from initial identification to operational ULB, can take several years and is subject to state-specific variations in policy, political will, and administrative capacity. Some states proactively municipalize urbanizing settlements, while others delay recognition due to fiscal or political constraints, leading to the phenomenon of Census Towns—settlements that are functionally urban but remain under rural governance.

3.3.1 What comes with ULB recognition?

When a settlement receives statutory recognition as an Urban Local Body (ULB), it undergoes a fundamental transformation in governance structure, fiscal arrangements, and administrative responsibilities. The 74th Constitutional Amendment Act of 1992 mandates that ULBs assume control over 18 functions listed in the Twelfth Schedule, including urban planning, regulation of land use, roads and bridges, water supply, public health and sanitation, fire services, urban poverty alleviation, and provision of urban amenities (Ministry of Home Affairs, Government of India, 2024). Crucially, ULB recognition brings changes in three key dimensions:

Governance and Administrative Structure: Statutory towns transition from rural governance under gram panchayats to elected municipal councils or corporations. This shift creates a dedicated urban administrative apparatus with specialized departments for engineering, health, education, and revenue collection. ULBs are headed by elected mayors or municipal chairpersons, with ward-level representation ensuring political accountability for urban service delivery.

Fiscal Autonomy and Revenue Sources: ULBs gain the authority to levy and collect municipal taxes, most notably property taxes, as well as user charges for water, sanitation, and other services. They also become eligible for devolved funds from state finance commissions and centrally-sponsored urban schemes such as the Jawaharlal Nehru National Urban Renewal Mission (JNNURM) and later the Smart Cities Mission and AMRUT (Atal Mission for Rejuvenation and Urban Transformation). While many ULBs remain fiscally dependent on state transfers, statutory recognition opens access to dedicated urban funding streams unavailable to rural settlements.

Service Delivery Mandates: Upon municipalization, ULBs inherit responsibility for providing and maintaining urban infrastructure and services. This includes establishing and operating primary schools, health dispensaries, water supply systems, sewerage and drainage networks, street lighting, solid waste management, and fire protection services.

However, the practical impact of ULB recognition is heterogeneous. Well-functioning municipalities in states with strong urban governance traditions may deliver substantially better services, while newly formed ULBs in resource-constrained states may struggle with capacity limitations, inadequate transfers, and weak revenue collection.

3.3.2 How does Census Town classification affect statutory recognition?

Although Census Town (CT) classification and statutory recognition are distinct processes that are built on different criteria and governed by different authorities, there is a strong relationship between the two and one often inspires the other (Tewari, 2020; Roy and Pradhan, 2018). Meeting the CT criteria signals that a settlement has urban characteristics and may warrant municipal governance. When a settlement crosses the CT thresholds, it draws attention from state urban development authorities and policymakers, who may then consider it for notification as a statutory town. The CT designation provides an objective benchmark indicating that the settlement has reached a level of population size, density, and economic activity consistent with urban areas. This can prompt state governments to evaluate whether the settlement is ready for municipal governance and whether it can sustain the administrative and fiscal responsibilities of a ULB. In fact, the central government also monitors CTs and encourages states to municipalize them, as exhorted by Ministry of Housing and Urban Affairs, Government of India (2016) where the Ministry of Housing and Urban Affairs (MoHUA) formally urged states to consider CTs for statutory recognition.

3.4 Data Sources and Construction

Using SHRUG data platform[47], we collect data at the village and town level[48] for our study (Asher, Novosad et al., 2019). We also collect state-level town directories from the Government of India's Open Government Data (OGD) platform and list of Statutory Towns (STs) from National Housing Bank's (NHB) website. We highlight the main datasets used in our analysis:

[47] version 2.1 was available at the time of our study.

[48] This is the most granular level at which information is available and is identified using SHRUG IDs.

- **Population Census Data (2001 & 2011):** We draw on comprehensive census data for all Indian settlements across two decennial rounds. The Primary Census Abstract (PCA) provides basic demographic characteristics including total population, gender and age composition, Scheduled Caste and Scheduled Tribe populations, literacy rates, and detailed workforce composition. These data allow us to construct the three Census Town eligibility criteria: population size, population density, and share of male main workers in non-agricultural activities.

- **Village and Town Directories:** We obtain infrastructure and service provision data from SHRUG's harmonized Census Village and Town Directories for 2001 and 2011. These include measures of education facilities such as primary through senior secondary schools, health infrastructure such as hospitals, dispensaries, primary health centers, family welfare centers, financial access variables such as cooperative banks, agricultural credit societies, and various community amenities. When definitions differ between the Village and Town Directories, we create harmonized outcome variables to ensure comparability across rural and urban settlements.[49]

- **Statutory Town Lists:** We obtain comprehensive lists of all statutorily recognized urban local bodies from the National Housing Bank and state-level census town directories from the Government of India's Open Government Data (OGD) platform.

Our final sample consists of over 500,000 settlements in India, which includes most inhabited villages and towns included in the 2001 and 2011 Census. Of these, over 3,700 have statutory urban status i.e. municipalities of some form as of 2011, while the remaining are governed as rural villages. However, many of those rural-governed settlements are urban in character. In total, about 9,000 settlements meet the Census definition of urban. This implies that roughly 2% of all Indian settlements were considered urban in 2011 by this criterion, and the rest were rural villages. These figures highlight that India's urban population is

[49] See the chapter appendix for the full variable crosswalk.

concentrated in a relatively small number of settlements – which tend to be much larger on average than the multitude of villages.

3.5 Running Variables and Treatment Assignment

The key running variables for our analysis are the three Census Town eligibility criteria: population size, population density, and share of male main workers in non-agricultural activities. We construct these variables using the 2001 Census data for each settlement. Specifically, we consider the following definitions for the running variables[50]:

- **Population Size (P):** Total population of the settlement as recorded in the 2001 Census.
- **Population Density (D):** Total population in 2001 divided by the area of the settlement in square kilometers, computed from the area measures provided in the Village and Town Directories.
- **Non-Agricultural Share (N):** The proportion of male main workers engaged in non-agricultural activities in 2001, calculated as the number of male main workers in non-agricultural sectors divided by the total number of male main workers in the settlement.

We then define the treatment variable for statutory recognition based on whether a settlement was officially designated as a statutory town (ST) in 2011. This is a binary indicator that takes the value of 1 if the settlement was recognized as an ST in 2011, and 0 otherwise. The treatment assignment is not deterministic based on the running variables alone, as some settlements that meet the Census Town criteria may not receive statutory

[50] Density plots of these variables are presented in the Appendix.

recognition due to state-level policy decisions, while others that do not meet the criteria may be recognized due to historical or political reasons. Therefore, we use a fuzzy regression discontinuity design that exploits the discontinuities in the probability of receiving statutory recognition at the thresholds defined by the running variables to identify the causal effect of local urban governance status on local public goods provision.

3.6 Summary Statistics

3.6.1 Urban/Rural Classification vs Statutory Recognition

We define statutory recognition using the state-level files available on Indian government's Open Government Data (OGD) platform[51]. These files report the list of statutory towns (municipalities, corporations, nagar panchayats, etc.) for each state as of 2011. We merge these files with the PCA data using state and settlement names to create an indicator for statutory recognition.

Table 13 provides a descriptive sense of this process by cross-tabulating threshold eligibility in 2001 and statutory status by 2011. Several patterns stand out. First, 17,509 settlements had populations above 5,000 yet were not statutory towns by 2011, compared to 3,570 settlements that were both above 5,000 and statutorily recognized (Panel A of Table 13). A similar discrepancy is seen for the density criterion (Panel B), where 194,146 settlements had density above 400 per km^2 in 2001 but were not statutory towns by 2011, while only 3,560 settlements met the density criterion and had statutory town status. For the non-agricultural workforce criterion (Panel C), 35,855 settlements had at least 75% non-agricultural male workers but were not statutory towns, compared to 2,673 that met this criterion in 2001 and were statutory towns by 2011. Second, looking at the combined criteria (Panel D), a total of

[51]State-level files were collected via https://data.gov.in/.

Table 13. Statutory Recognition and Urban Thresholds

Panel A: Population $\geq$ 5,000		
	Pop $<$ 5,000	**Pop $\geq$ 5,000**
Not Statutory	485,580	17,509
Statutory	159	3,570

Panel B: Density $\geq$ 400 per km^2		
	Density $<$ 400	**Density $\geq$ 400**
Not Statutory	308,943	194,146
Statutory	169	3,560

Panel C: Non-Agricultural Male Workers $\geq$ 75%		
	Non-Ag $<$ 75%	**Non-Ag $\geq$ 75%**
Not Statutory	467,234	35,855
Statutory	1,056	2,673

Panel D: All Three Thresholds Combined		
	Did not Meet Thresholds	**Met Thresholds**
Not Statutory	499,708	3,381
Statutory	1,194	2,535

5,916 settlements met all three urban criteria in 2001 (3,381 not statutory + 2,535 statutory). However, only 2,535 of these were actually recognized as statutory towns by 2011, while 3,381 settlements met the criteria but remained governed as rural. Conversely, there were 1,194 statutory towns that did not meet all the Census criteria. This highlights that satisfying the Census definition is not sufficient for urban governance as political implementation by states is crucial. It also underscores the quasi-random aspect of our design: among settlements around the thresholds, some received statutory recognition by 2011 and others did not, not purely based on merit but partly due to state-specific policies and timing. Nevertheless, as we show below, crossing the thresholds greatly increases the likelihood of statutory recognition, providing the basis for our fuzzy regression discontinuity approach.

3.6.2 Treatment Assignment and Treatment Status

Table 14 presents the distribution of settlements by Census Town (CT) classification in 2001 and statutory town (ST) status by 2011. Panel A shows the full sample of all settlements, while Panel B restricts to settlements close to the CT thresholds.

Table 14. Treatment Assignment and Status

Panel A: Full Sample			
CT in 2001	**Statutory by 2011**	**N**	**Share (%)**
No	No	495,215	98.61
No	Yes	1,199	0.24
Yes	No	3,223	0.64
Yes	Yes	2,542	0.51
Total		**502,179**	**100.00**

Panel B: Close to Threshold Sample			
CT in 2001	**Statutory by 2011**	**N**	**Share (%)**
No	No	36,726	98.86
No	Yes	100	0.27
Yes	No	297	0.80
Yes	Yes	28	0.08
Total		**37,151**	**100.00**

Notes: "CT in 2001" indicates whether a settlement met all three Census Town criteria (population $\geq$ 5,000, density $\geq$ 400 per km^2, non-agricultural workers $\geq$ 75%) in 2001. "Statutory by 2011" indicates whether the settlement had statutory town status in 2011. Panel B restricts to settlements within the bandwidth around the CT thresholds.

Several patterns emerge from Table 14. In the full sample (Panel A), the vast majority of settlements (98.6%) neither met the CT criteria in 2001 nor received statutory recognition by 2011. Among the 5,765 settlements that met the CT thresholds in 2001, only 2,542 (44%) gained statutory status by 2011, while 3,223 (56%) remained under rural governance despite being functionally urban. This substantial gap between eligibility and recognition illustrates the discretionary nature of state municipalization policy. Conversely, 1,199 settlements that did not meet the CT criteria nevertheless received statutory recognition, typically reflecting historical municipal status or special circumstances.

The close-to-threshold sample (Panel B) reveals similar patterns on a smaller scale, with 297 settlements meeting CT criteria but not receiving statutory status, and 100 settlements gaining statutory status despite not meeting the criteria. This variation around the thresholds provides the identifying variation for our fuzzy regression discontinuity design, where crossing the CT thresholds substantially increases but does not guarantee statutory recognition.

In this section, we lay out the identification strategy. The goal is to estimate the causal effect of statutory recognition and ULB status on settlement outcomes.

3.7 Quasi-random variation from Census Towns

In this section, we show that Census Town (CT) designation leads to a discontinuous increase in the probability of statutory recognition at the thresholds. Furthermore, becoming a CT does not directly change governance status as the decision to grant a ULB is made by the state government. So theoretically, we should not expect a direct effect of CT designation on local outcomes through the local urban governance status channel unless the settlement is also statutorily recognized. One could argue that CT may be capturing the differences in agglomeration economies or other unobserved characteristics that may be correlated with local outcomes. However, close to the thresholds, these characteristics should be similar on either side of the cutoff. For a country of 1.5 billion people, a settlement with a population of 4,900 and 5,100, or density of 390 and 410, is not likely to be systematically different in terms of agglomeration economies. Furthermore, even though we cannot rule out differences in unobserved characteristics, we can test for balance in observed characteristics (e.g., literacy rate, caste composition) and run placebo tests at pseudo-thresholds to check if there are any discontinuities in outcomes at these points. First, we show that there is a strong first-stage relationship between meeting the CT thresholds and statutory recognition. We then use this quasi-random variation to instrument for statutory recognition and identify the effect of statutory recognition on local outcomes.

Below we describe how we implement this empirical strategy in practice. Let P_i, D_i, and N_i denote the population, density (per km^2), and non-agricultural employment share of settlement i. Define three binary indicators for meeting the CT thresholds:

$$
\begin{aligned}
Z_{P,i} &= \mathbb{1}\left\{\frac{P_i}{5000} - 1 \geq 0\right\} \\
Z_{D,i} &= \mathbb{1}\left\{\frac{D_i}{400} - 1 \geq 0\right\} \\
Z_{N,i} &= \mathbb{1}\left\{\frac{N_i}{0.75} - 1 \geq 0\right\}
\end{aligned}
$$

We define CT eligibility as $Z_i = Z_{P,i} \times Z_{D,i} \times Z_{N,i}$, which equals 1 if settlement i meets all three thresholds and 0 otherwise, which we use as an instrument for statutory recognition. The idea is that settlements that meet the CT thresholds are more likely to receive statutory recognition, but the decision to grant statutory recognition is not deterministic based on the CT thresholds alone. Therefore, we can use the variation in CT eligibility to identify the causal effect of statutory recognition on local outcomes.

Next, let ST_i be an indicator for statutory recognition by 2011 i.e. whether settlement i gets an Urban Local Body (ULB) or not, by 2011. Then, our object of interest is the probability of statutory recognition $P(ST_i = 1|Z_i, X_i)$, where X_i is a vector of controls. Our data enables us to estimate the following first stage regression:

$$ST_i = \pi_0 + \pi_1 Z_i + X_i'\beta + \delta_{s(i)} + \varepsilon_i \tag{6}$$

where Z_i is CT eligibility in 2001 (instrument), X_i is a vector of smooth functions of the running variables and other controls (e.g. $Z_{P,i}, Z_{D,i}, Z_{N,i}$, literacy, caste shares), and $\delta_{s(i)}$ are district-level fixed effects.

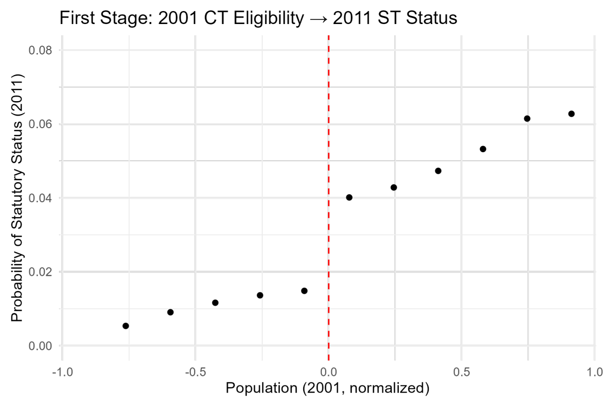

(A) Population threshold

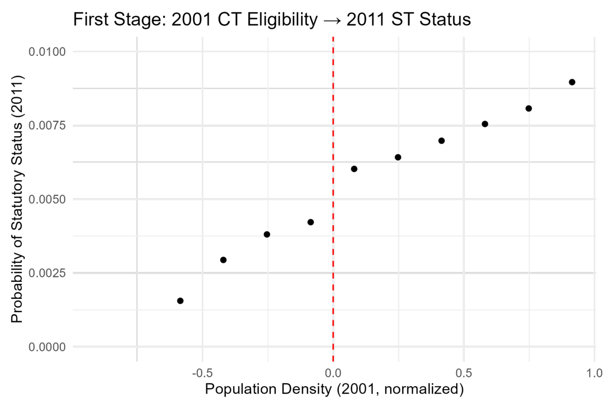

(B) Density threshold

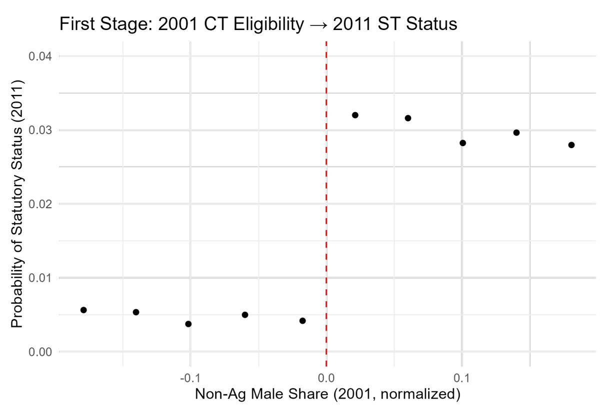

(C) Non-agricultural share threshold

Figure 12. First-stage Discontinuities and Statutory Recognition

Notes: Each panel shows a binned scatter plot of the predicted probability of statutory recognition in 2011 against one of the three normalized running variables from 2001. Predicted probabilities are fitted values from the global first-stage OLS regression of statutory status on Census Town eligibility, controls, and district fixed effects (Equation 6). The dashed red line marks the normalized threshold.

Figure 12 presents graphically the first stage of our fuzzy RD design explained in equation 6 close to the thresholds[52]. The figure shows binned scatter plots of the probability of statutory status in 2011 against each running variable in 2001. Sub-figure (A) of Figure 12 shows the relationship between population and statutory recognition. Probability of statutory recognition is positively related to population, which is expected since an increase in population signifies a larger settlement that is more likely to be recognized as a town. Despite this positive relationship, there is a clear discontinuity in the probability of statutory recognition at the population threshold of 0 after normalization. This discontinuity indicates that settlements that just meet the population threshold are significantly more likely to receive statutory recognition than those that just miss it. Sub-figure (B) of Figure 12 shows the relationship between density and statutory recognition. Similar to population, there is a positive relationship between density and statutory recognition, but the discontinuity at the density threshold, while present, is less pronounced than for population. Sub-figure (C) of Figure 12 shows the relationship between non-agricultural share and statutory recognition. The relationship is flat, but there is a clear discontinuity at the non-agricultural share threshold. Settlements that just meet the non-agricultural share threshold are more likely to receive statutory recognition than those that just miss it. Overall, these figures provide evidence of a strong first-stage relationship between meeting the CT thresholds and statutory recognition, which supports our identification strategy for estimating the causal effect of statutory recognition on local outcomes. In Section 3.9, we present a table that reports the first-stage coefficient and F-statistics for global and local estimates.

[52] We present results up to 100% away from the thresholds.

3.8 Fuzzy RD with multiple running variables

India's CT thresholds create jointly-binding rules on three dimensions: population $P \geq 5{,}000$, density $D \geq 400$ per km², and male main non-agricultural share $N \geq 0.75$. Define running variables as normalized distances to the cutoffs: $r_P = P/5000 - 1$, $r_D = D/400 - 1$, and $r_N = N/0.75 - 1$. Following the logic of multi-cutoff RD, we focus on observations near the cutoffs and develop a frontier using a soft-min functional specification.

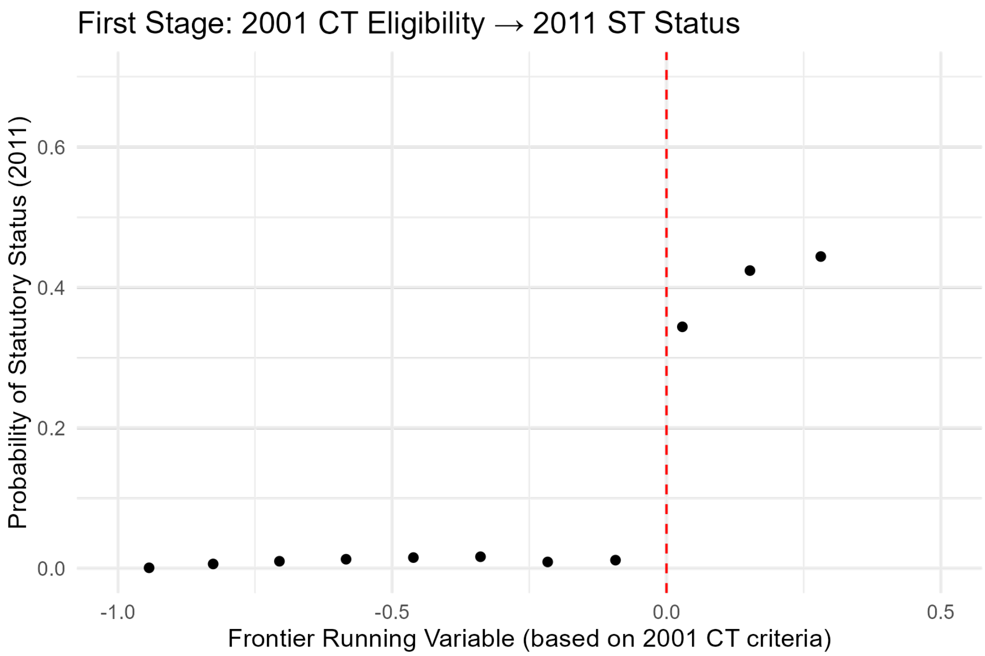

Figure 13. Probability of Statutory Recognition by Frontier Distance

Notes: Binned scatter plot of the predicted probability of statutory recognition in 2011 against the soft-min frontier running variable constructed from the three 2001 Census Town criteria. The frontier variable is defined as $r_i^* = -\frac{1}{\alpha}\ln\big(\exp(-\alpha\, r_P) + \exp(-\alpha\, r_D) + \exp(-\alpha\, r_N)\big)$ with $\alpha = 10$, which smoothly approximates the minimum of the three normalized distances. Negative values indicate that a settlement falls short of at least one threshold; positive values indicate that all three thresholds are met. Each dot represents the mean predicted probability within an equal-width bin. Predicted probabilities are fitted values from the global first-stage regression[53] (Equation 6).

Figure 13 plots the predicted probability of statutory recognition against the frontier running variable, which collapses the three Census Town dimensions into a single scalar. The figure reveals a clear discontinuity at the eligibility frontier, near the zero on the horizontal axis. To the left of the cutoff, where settlements fall short of at least one threshold, the predicted probability of receiving a municipal body is close to zero and hovers near 1%. Immediately to the right, where settlements satisfy all three criteria, the probability jumps sharply to approximately 35%, rising further to above 40% for settlements that comfortably exceed the thresholds. This pattern is precisely what the fuzzy RD design requires: a large, discrete change in the probability of treatment at the eligibility boundary, embedded within an otherwise smooth relationship. The magnitude of the jump is substantial and visually unambiguous, providing strong graphical evidence that Census Town eligibility is a predictor of statutory recognition. This reinforces the relevance condition of our instrument and is consistent with the formal first-stage F-statistics reported in Table 15.

The contrast between Figure 13 and the individual-threshold plots in Figure 12 is also informative. When projected onto any single running variable, the discontinuity appears modest because most of the variation in eligibility is driven by the joint interaction of all three criteria. The frontier variable captures this joint variation and motivates our use of a multi-dimensional instrument defined by the intersection of all three thresholds, rather than relying on any single cutoff.

3.9 First stage: Effects on Statutory Recognition

The first-stage results in Table 15 reveal important patterns regarding the strength and validity of our instrumental variable strategy. In the global sample, the instrument exhibits substantial explanatory power with a coefficient of 0.431 and a first-stage F-statistic of 749.43, well exceeding conventional thresholds for weak instrument concerns. However, this

Table 15. First-stage: Effect of Census-Town Eligibility on Statutory Status

	Global	**Local**
Dependent variable: Statutory town = 1		
CT eligibility in 2001	0.431***	0.071***
	(0.0157)	(0.0167)
Controls	Yes	Yes
District fixed effects	Yes	Yes
SEs clustered at district level	Yes	Yes
First-stage F-statistic	749.43	18.05
Partial R^2 of instrument	0.2522	0.0127
Adj. R^2	0.3493	0.1007
Observations	502,179	37,151
Districts (FE groups)	619	558

Notes: First-stage estimates of statutory recognition on the census-threshold instrument with controls and district fixed effects. Controls include population, density, non-agricultural male workforce share, literacy rate, and caste composition. Standard errors are clustered by district (in parentheses). "Local" restricts to settlements near the 2001 census thresholds (population ≈ 5,000, density ≈ 400, non-agricultural male workforce share ≈ 0.75). The single-IV first-stage F is t^2 from the clustered specification.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

strong statistical relationship may be misleading for causal identification purposes. The global specification, while statistically powerful, is potentially vulnerable to endogeneity. Settlements across the entire population and density spectrum may differ systematically in ways that correlate with both census threshold compliance and subsequent outcomes. For instance, larger settlements may have greater political influence or administrative capacity that simultaneously makes them more likely to meet census criteria and more likely to attract infrastructure investment, regardless of statutory status.

In contrast, the local specification of restricting analysis to settlements near the census cutoffs[54] provides more credible identification despite the smaller coefficient magnitude (0.071) and relatively lower F-statistic (18.05). This specification focuses on settlements that are observationally similar except for small differences in census characteristics that push them just above or below the thresholds. The lower adjusted R^2 (0.1007) and partial

[54]We use the following restrictions around the thresholds: (i) population ±5000 (ii) density ±400 (iii) non-ag male main work share ±0.2. Table 38 in the appendix presents bandwidth-sensitivity results.

R^2 of instrument (0.0127) indicate that we are capturing variation that is more plausibly exogenous, as settlements just above and below the cutoffs should be similar in unobservable characteristics in the absence of treatment, as long as no manipulation occurs.

The F-statistic of 18.05 for our local specification is above the conventional threshold of 16.38 for strong instruments in the just-identified case (Staiger and Stock, 1997) and well above the critical value of 8.96 for 15% maximal IV bias (Stock and Yogo, 2002), suggesting the instrument maintains reasonable strength for our identification strategy. The local estimates thus provide our preferred specification for downstream analysis, offering a credible foundation for interpreting the causal effects of statutory recognition on settlement outcomes.

3.10 Main effects on outcomes

We next present 2SLS estimates of the effect of statutory recognition on settlement outcomes. We group outcomes into categories most relevant to local public goods and services: education, health, financial access, and community infrastructure. The reduced-form estimates provide the numerator of the Wald ratio underlying our IV specification and show that crossing the CT eligibility frontier improves public goods provision across the same outcome categories[55].

Education. Table 16 presents the estimated effects of statutory recognition on educational infrastructure provision. The results demonstrate substantial positive impacts across all levels of schooling. Statutory recognition leads to an increase of 13.86 additional primary schools, 7.72 additional middle schools, and 4.89 additional secondary schools. All estimates are statistically significant at the 1% level, with robust standard errors clustered at the district level. The largest absolute effect occurs at the primary level, consistent with primary education being a foundational service that local governments can expand most readily. Primary schools

[55]The full reduced-form table is reported in Appendix Table 36.

Table 16. Effect of Local Urban Governance Status on School Provision

	Dependent variable: Number of government schools		
	Primary	Middle	Secondary
Effect of ULB status	13.86***	7.72***	4.89***
	(4.00)	(2.25)	(1.30)
Controls	Yes	Yes	Yes
District FE	Yes	Yes	Yes
Observations	37,143	37,099	37,148

Notes: Table reports two-stage least squares estimates where the endogenous regressor is statutory town status (ULB = 1) and the excluded instrument is Census Town eligibility in 2001. The local sample restricts to settlements near the 2001 census thresholds (population ±5,000; density ±400; non-agricultural male share ±0.20). All specifications include controls for log population, log density, non-agricultural male workforce share, literacy rate, main worker share, and caste composition (SC and ST shares), as well as district fixed effects. Robust standard errors clustered at the district level are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

are also typically smaller in size and easier to open than middle or secondary schools, facilitating rapid increases in provision following statutory recognition. As we move up the educational ladder, the estimated effects become smaller, which is plausible given the higher capital requirements, larger catchment areas, and greater regulatory demands associated with middle and secondary schools. Even so, the fact that all three levels show significant positive effects suggests that ULB status facilitates broad improvements in educational infrastructure rather than a narrow shift at only one school level.

Health. Table 17 presents the estimated effects of local urban governance status on health infrastructure provision. The results demonstrate substantial positive impacts across both health facilities. Local urban governance status leads to an increase of 2.53 additional hospitals and 3.00 additional family welfare centers. Both estimates are statistically significant at the 1% level, with robust standard errors clustered at the district level. The positive effects on both general hospitals and family welfare centers suggest that local urban governance status facilitates improvements in core health infrastructure. The somewhat larger effect on family welfare centers is noteworthy because these facilities play a central role in maternal and child health, family planning, and preventive care within India's public health system (Dhingra

Table 17. Effect of Local Urban Governance Status on Hospital Provision

	Dependent variable	
	All hospitals	Family welfare centers
Effect of ULB status	2.53^{***}	3.00^{***}
	(0.69)	(0.88)
Controls	Yes	Yes
District FE	Yes	Yes
Observations	37,091	37,077

Notes: Two-stage least squares estimates where the endogenous regressor is statutory town status (ULB = 1) and the excluded instrument is Census Town eligibility in 2001. The local sample restricts to settlements near the 2001 census thresholds (population ±5,000; density ±400; non-agricultural male share ±0.20). All specifications include controls for log population, log density, non-agricultural male workforce share, literacy rate, main worker share, and caste composition (SC and ST shares), as well as district fixed effects. Robust standard errors clustered at the district level are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

and Dutta, 2011). This pattern is also plausible given that family welfare centers typically require less capital and can be expanded more quickly than full-service hospitals.

Table 18. Effect of Local Urban Governance Status on Financial Access

	Dependent variable	
	Cooperative banks	Agricultural credit societies
Effect of ULB status	4.09^{***}	2.84^{***}
	(0.97)	(1.01)
Controls	Yes	Yes
District FE	Yes	Yes
Observations	37,089	37,093

Notes: Two-stage least squares estimates where the endogenous regressor is statutory town status (ULB = 1) and the excluded instrument is Census Town eligibility in 2001. The local sample restricts to settlements near the 2001 census thresholds (population ±5,000; density ±400; non-agricultural male share ±0.20). All specifications include controls for log population, log density, non-agricultural male workforce share, literacy rate, main worker share, and caste composition (SC and ST shares), as well as district fixed effects. Robust standard errors clustered at the district level are reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Financial Access. Table 18 presents the estimated effects of local urban governance status on financial infrastructure provision. The results demonstrate substantial positive impacts across both types of financial institutions. Local urban governance status leads to an increase of 4.09 additional cooperative banks and 2.84 additional agricultural credit societies, both statistically significant at the 1% level with robust standard errors clustered at the district level. The positive effects across both institution types indicate that municipalization enhances multiple forms of financial intermediation serving different community needs. Cooperative banks can support broader commercial and household financial activity, while agricultural credit societies remain relevant in peri-urban settlements where agricultural livelihoods persist alongside urbanization. The expansion of financial infrastructure may therefore reflect both greater demand in newly recognized towns and the possibility that local urban governance status signals administrative capacity and development potential to formal financial institutions (Burgess and Pande, 2005).

Table 19. Effect of Local Urban Governance Status on Community Infrastructure

	Dependent variable			
	Public libraries	Public reading rooms	Cinema halls	Sports infrastructure
Effect of ULB status	1.05**	1.44**	0.79	−5.71***
	(0.50)	(0.63)	(0.49)	(1.48)
Controls	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Observations	37,147	37,147	37,146	37,146

Notes: Two-stage least squares estimates where the endogenous regressor is statutory town status (ULB = 1) and the excluded instrument is Census Town eligibility in 2001. The local sample restricts to settlements near the 2001 census thresholds (population ±5,000; density ±400; non-agricultural male share ±0.20). All specifications include controls for log population, log density, non-agricultural male workforce share, literacy rate, main worker share, and caste composition (SC and ST shares), as well as district fixed effects. Robust standard errors clustered at the district level are reported in parentheses.
Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Community Infrastructure. Table 19 presents the estimated effects of local urban governance status on community infrastructure and access. The results show positive effects on several cultural and educational amenities. Local urban governance status leads to a statistically significant increase of 1.05 public libraries and 1.44 public reading rooms. The estimate for cinema halls is positive at 0.79 but imprecisely estimated. By contrast, sports infrastructure declines by 5.71 facilities. This negative effect is consistent with spatial constraints and land-use tradeoffs in urban development: as settlements formalize into statutory towns, open spaces previously used for sports facilities may be repurposed for schools, health facilities, banking infrastructure, roads, or residential development. Overall, these findings suggest that formal statutory recognition may enhance access to community facilities, though the transition may involve restructuring of certain types of infrastructure, particularly those requiring substantial land area.

3.11 Robustness Checks

3.11.1 Density Plots and McCrary Test for Manipulation

We conduct checks around the statutory recognition thresholds to assess the validity of our RD design. Specifically, we plot the density of the running variables and implement the McCrary (2008) test for manipulation of the running variable, which examines the continuity of the density of the running variable at the cutoff.

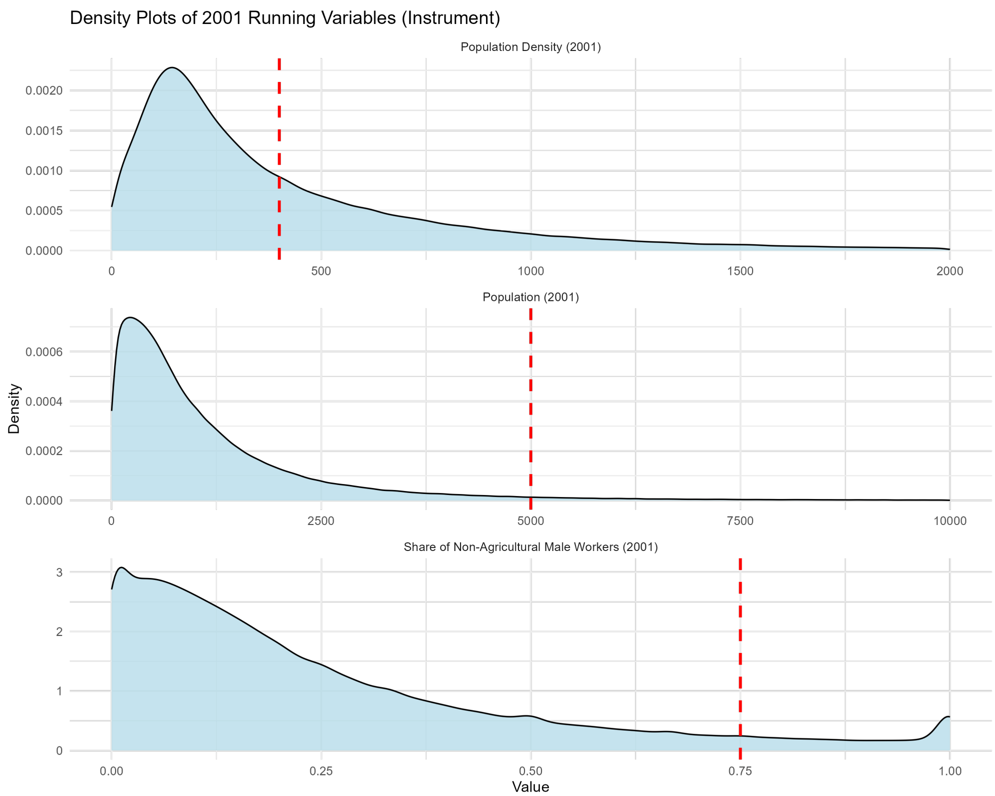

Figure 14. Density plots of CT-definition variables and thresholds

Figure 14 shows the density plots of the three running variables (population density, population, and share of non-agricultural male workers) around their respective cutoffs of 400 persons per km^2, 5,000 persons, and 0.75. The plots indicate that the density of the running variables is continuous and smooth around the cutoffs, with no clear signs of manipulation or strategic sorting.

Table 20. McCrary Density Tests at CT Thresholds

Running Variable	Cutoff	T-statistic	P-value
Population (2001)	5,000	0.754	0.451
Population Density (2001)	400	0.245	0.807
% Non-Agricultural Male Workers (2001)	0.75	−1.101	0.271

Notes: This table reports results from the McCrary (2008) manipulation test which examines whether there is a discontinuity in the density of the running variable at the statutory threshold. The T-statistic tests the null hypothesis of no discontinuity. All tests use a triangular kernel with bandwidth selection and jackknife variance estimation. P-values greater than 0.10 indicate no evidence of manipulation at conventional significance levels.

The McCrary (2008) density test formally assesses whether there is evidence of manipulation of the running variable around the cutoff. The test estimates the density of the running variable separately on each side of the threshold and tests whether there is a discontinuity in the density at the cutoff. A significant discontinuity would suggest strategic sorting around the threshold, which would violate the identifying assumptions of the RD design.

Table 20 presents the results of the McCrary (2008) test for each of our three running variables used to define census town status. For population in 2001 (cutoff at 5,000), the test statistic is 0.754 with a p-value of 0.451. For population density in 2001 (cutoff at 400 persons per km^2), the test statistic is 0.245 with a p-value of 0.807. For the share of non-agricultural male workers (cutoff at 0.75), the test statistic is -1.101 with a p-value of 0.271. In all three cases, we fail to reject the null hypothesis of no discontinuity in the density at conventional significance levels. This provides reassuring evidence that there is no systematic manipulation of these running variables around the statutory thresholds. Remaining robustness tests can be found in the Appendix.

3.12 Conclusion

This paper provides evidence that local urban governance status leads to meaningful improvements in local public goods at India's urbanizing fringe. Using a novel multi-threshold fuzzy RD design that leverages Census Town eligibility criteria, we isolate the effect of local urban governance status on public goods provision from underlying confounding factors.

Our local IV estimates show that statutory recognition increases the number of government schools, hospitals, family welfare centers, cooperative banks, and agricultural credit societies, while decreasing the number of sports facilities, consistent with the idea of reallocation of space and land constraints in formalizing settlements. These local average treatment effects indicate that ULB status materially shifts the stock of public infrastructure in settlements near the Census Town thresholds. The corresponding reduced-form estimates are directionally consistent and confirm that crossing the eligibility frontier improves public goods provision even before scaling by the first stage.

The policy implications are twofold. First, there are clear benefits to timely municipalization of emerging urban areas. Our results show that delaying the recognition of urban settlements may imply missed opportunities for public infrastructure. Proactively converting eligible large villages into statutory towns and providing them with fiscal support and a governance framework could help improve their overall infrastructure. Second, our results highlight the importance of state capacity and support for new ULBs. We emphasize that simply declaring a settlement as a municipality is not a panacea. The improvements we document likely come from a combination of federal, state, and local initiatives, the ability to levy local taxes, greater funding availability, and improved local governance. State governments could enhance the benefits of municipalization by ensuring that newly formed ULBs receive adequate resources, technical assistance, and training to fulfill their functions.

Our study opens avenues for future research. One extension would be to examine longer-term outcomes beyond 2011, using household survey data such as subsequent rounds of NFHS to assess impacts on education, health, and income. A clear need remains to understand *why* and *how* local urban governance status leads to improved public goods provision. Is it primarily through increased fiscal resources, improved administrative capacity, greater political representation, or a combination of these factors? We encourage future work to unpack the mechanisms through which municipalization translates into better infrastructure and services. Another avenue is to explore the fiscal channel: how do municipal finances change at the threshold, and to what extent do new ULBs rely on own-source revenue versus transfers? Understanding heterogeneity in statutory recognition across states, such as why some are reluctant to create new towns, possibly to avoid sharing revenue or due to political patronage networks, could inform policymakers seeking to overcome barriers to timely municipalization. Our findings underscore that formally bringing settlements into the urban administrative fold has measurable benefits, and that ensuring institutional change keeps pace with demographic change will be crucial for sustainable development in rapidly urbanizing economies.

References

Asher, Sam, Paul Novosad, et al. 2019. "The Socioeconomic High-resolution Rural-Urban Geographic Dataset on India (SHRUG)." *Development Data Lab (data platform documentation)*, Version details as used in the paper.

Besley, Timothy, and Torsten Persson. 2010. "State capacity, conflict, and development." *Econometrica*, 78(1): 1–34.

Brollo, Fernanda, Tommaso Nannicini, Roberto Perotti, and Guido Tabellini. 2013. "The Political Resource Curse." *American Economic Review*, 103(5): 1759–1796.

Burgess, Robin, and Rohini Pande. 2005. "Do rural banks matter? Evidence from the Indian social banking experiment." *American economic review*, 95(3): 780–795.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel. 2015. "The value of democracy: evidence from road building in Kenya." *American Economic Review*, 105(6): 1817–1851.

Denis, Eric, and Kamala Marius-Gnanou. 2010. "Toward a better appraisal of urbanization in India." *Cybergeo: European Journal of Geography.* Publisher: CNRS-UMR Géographie-cités 8504.

Dhingra, Bhavna, and Ashok Kumar Dutta. 2011. "National rural health mission." *Indian Journal of Pediatrics*, 78(12): 1520–1526.

Faguet, Jean-Paul. 2012. *Decentralization and Popular Democracy: Governance from Below in Bolivia.* Ann Arbor:University of Michigan Press.

Gadenne, Lucie. 2017. "Tax Me, but Spend Wisely? Sources of Public Finance and Government Accountability." *American Economic Journal: Applied Economics*, 9(1): 274–314.

Government of Goa. 1968. "Goa Municipalities Act, 1968." *https://www.indiacode.nic.in/bitstream/123456789/10078/1/goa_municipalities_act%2C_1968.pdf*, Accessed: 2025-10-30.

Government of Gujarat. 1963. "The Gujarat Municipalities Act, 1963." *https://udd.gujarat.gov.in/pdf/L/Act_and_rules/TheGujaratMunicipalitiesAct-1963.pdf*, Accessed: 2025-10-30.

Government of Karnataka. 1964. "Karnataka Municipalities Act, 1964." *https://www.indiacode.nic.in/bitstream/123456789/8130/1/22_of_1964_%28e%29.pdf*, Accessed: 2025-10-30.

Government of Maharashtra. 1965. "Maharashtra Municipalities Act, 1965." *https://prsindia.org/files/bills_acts/acts_states/maharashtra/1965/1965MH40.pdf*, Accessed: 2025-10-30.

Government of Telangana. 2019. "Telangana Municipalities Act, 2019." *https://www.cgg.gov.in/wp-content/uploads/2022/05/CGG-publication-on-TM-ACT.pdf*, Accessed: 2025-10-30.

Lee, Hae Nim. 2021. "Essays on Local Governance in India - ProQuest."

McCrary, Justin. 2008. "Manipulation of the running variable in the regression discontinuity design: A density test." *Journal of econometrics*, 142(2): 698–714.

Ministry of Home Affairs, Government of India. 2024. "74th Amendment and Municipalities in India." *https://secforuts.mha.gov.in/74th-amendment-and-municipalities-in-india/*, Accessed: 2025-10-30.

Ministry of Housing and Urban Affairs, Government of India. 2016. "States asked to convert 3,784 urban areas into statutory Urban Local Bodies." *Press Information Bureau Press Release*, Accessed: 2025-10-30.

Mitra, Arup, and Jay Prakash Nagar. 2018. "City size, deprivation and other indicators of development: Evidence from India." *World Development*, 106: 273–283.

Mukhopadhyay, Partha. 2015. "Republic of India Understanding India's Urban Frontier."

Mukhopadhyay, Partha. 2017. "Does Administrative Status Matter for Small Towns in India?" In *Subaltern Urbanisation in India: An Introduction to the Dynamics of Ordinary Towns.* , ed. Eric Denis and Marie-Hélène Zérah, 443–469. New Delhi:Springer India.

Narasimhan, Veda, and Jeffrey Weaver. 2024. "Polity Size and Local Government Performance: Evidence from India." *American Economic Review*, 114(11): 3385–3426.

Patrick, Carlianne, and Christopher Mothorpe. 2017. "Demand for New Cities: Property Value Capitalization of Municipal Incorporation." *Regional Science and Urban Economics*, 67: 78–89.

Pradhan, Kanhu. 2013. "Unacknowledged Urbanisation: New Census Towns of India."

Roy, Shamindra Nath, and Kanhu Charan Pradhan. 2018. "Predicting the future of census towns." *Economic and Political Weekly*, 53(49): 70–79.

Staiger, Douglas, and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica*, 65(3): 557–586.

Stock, James H., and Motohiro Yogo. 2002. "Testing for Weak Instruments in Linear IV Regression." National Bureau of Economic Research Technical Working Paper 0284, Cambridge, MA.

Tewari, Surya. 2020. "Inter-State Variation in Statutory and Census Towns." National Institute of Urban Affairs, New Delhi. Accessed: 2025-10-30.

Appendix

Appendix for Chapter 1

A1.1 Road Tax Renewal Elections in Ohio

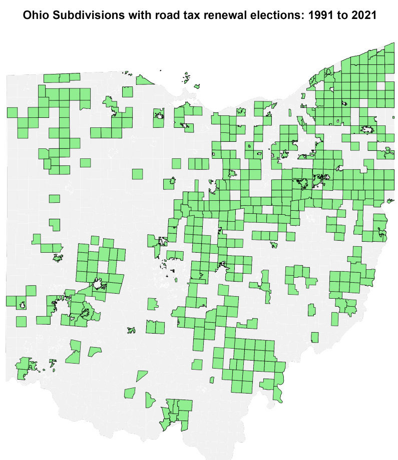

Figure 15. All Road Tax Renewal Elections in Ohio (1991-2021)

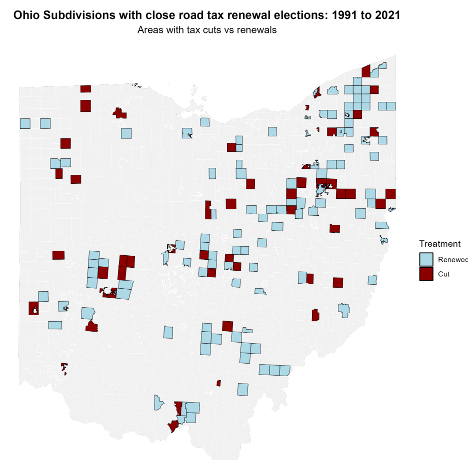

Figure 16. Close Road Tax Renewal Elections in Ohio (1991-2021)

Figure 15 displays the distribution of county subdivisions in Ohio that held at least one road tax renewal election between 1991 and 2021. The map reveals that although elections are relatively dispersed, they are more common in metropolitan areas such as Cleveland, Columbus, Cincinnati, Toledo etc. For jurisdictions that held multiple elections during this period, Figure 16 focuses on the subset of elections that were decided by narrow margins, specifically those closest to the 50% approval threshold that forms the basis of our regression discontinuity analysis. These close elections provide the quasi-experimental variation we exploit to identify the causal effects of road tax levy decisions on housing prices, and we do not observe any spatial clustering of the appearance of these close elections on the ballot, as well as the results of these close elections. The geographic distribution of close elections largely mirrors that of all elections, suggesting that narrow electoral outcomes are not systematically concentrated in particular types of communities, which supports the validity of our identification strategy.

A1.2 Additional Tables

A1.2.1 Full set of Treatment Effects

Tables below present the full set of treatment effects for the results presented in our heterogeneity analysis and robustness checks. Table 21 presents the treatment effects after applying 1% Winsorization to the data, ensuring robustness against outliers. Table 22 breaks down the treatment effects by urban and rural categories, highlighting the differential impacts on housing prices based on the urbanization level.

Table 21. Full set of estimates – Median Housing Price (after 1% Winsorization)

Year relative to vote	**Estimate**	**Std. error**	***p*-value**	**Confidence interval**
$t-3$	258	6,217	0.967	[−11,927, 12,443]
$t-2$	1,870	6,257	0.765	[−10,393, 14,134]
$t-1$	3,489	7,202	0.628	[−10,627, 17,605]
t	−1,626	6,730	0.809	[−14,818, 11,566]
$t+1$	−4,312	8,906	0.628	[−21,767, 13,143]
$t+2$	−8,190	8,099	0.312	[−24,063, 7,683]
$t+3$	−13,275	7,588	0.080	[−28,146, 1,597]
$t+4$	−22,582	7,904	0.004	[−38,074, −7,090]
$t+5$	−20,952	9,554	0.028	[−39,678, −2,225]
$t+6$	−18,793	8,789	0.033	[−36,020, −1,566]
$t+7$	−17,918	7,180	0.013	[−31,990, −3,845]
$t+8$	−21,146	8,921	0.018	[−38,631, −3,660]
$t+9$	−17,789	7,770	0.022	[−33,018, −2,561]
$t+10$	−13,147	7,526	0.081	[−27,898, 1,604]

Notes: Supplements Figure 9 in main text. This table reports the treatment-effect estimates of cutting road tax levies (relative to renewing them) from three years before to ten years after the vote. Data are winsorized at the 1% level. All regressions include the covariates listed in Table 3. The outcome is the median house price in constant 2010 U.S. dollars (city-year observations). A coefficient of −22,582 at $t+4$ means that, four years after the vote, cities that cut their levies had median house prices $22,582 lower than those that renewed. Standard errors are in parentheses.

Table 22. Treatment Effects on Housing Prices by Urban vs. Rural Categories

Panel A: Urban

Year	Estimate	Std. Error	p-value	Conf. Interval
$t-3$	-2,636	8,066	0.744	[-18,446, 13,173]
$t-2$	-9,607	7,310	0.189	[-23,935, 4,722]
$t-1$	1,045	6,496	0.872	[1,045, 1,045]
t	458	7,873	0.954	[-14,973, 15,889]
$t+1$	-5,087	7,617	0.504	[-20,016, 9,843]
$t+2$	-3,675	9,077	0.686	[-21,465, 14,115]
$t+3$	-11,657	7,667	0.128	[-26,684, 3,370]
$t+4$	-8,846	9,162	0.334	[-26,804, 9,112]
$t+5$	-8,967	9,311	0.336	[-27,217, 9,284]
$t+6$	-24,476	7,127	0.001	[-38,446, -10,507]
$t+7$	-14,457	7,869	0.066	[-29,880, 966]
$t+8$	-26,174	8,921	0.003	[-43,659, -8,688]
$t+9$	-19,469	8,221	0.018	[-35,582, -3,357]
$t+10$	-23,969	10,364	0.021	[-44,284, -3,655]

Panel B: Rural

Year	Estimate	Std. Error	p-value	Conf. Interval
$t-3$	6,384	9,962	0.522	[-13,142, 25,910]
$t-2$	1,609	7,758	0.836	[-13,597, 16,814]
$t-1$	-835	8,354	0.920	[-835, -835]
t	-8,727	9,351	0.351	[-27,056, 9,601]
$t+1$	5,731	7,458	0.442	[-8,886, 20,349]
$t+2$	-2,970	7,107	0.676	[-16,899, 10,960]
$t+3$	-1,866	6,538	0.775	[-14,681, 10,949]
$t+4$	-10,505	8,641	0.224	[-27,441, 6,431]
$t+5$	-16,897	7,564	0.026	[-31,722, -2,072]
$t+6$	-4,561	6,551	0.486	[-17,402, 8,280]
$t+7$	-7,632	8,936	0.393	[-25,146, 9,882]
$t+8$	-469	8,063	0.954	[-16,273, 15,335]
$t+9$	-11,492	8,383	0.170	[-27,923, 4,940]
$t+10$	1,851	7,790	0.812	[-13,417, 17,119]

Notes: Supplements Figure 6 in text. Each panel reports separate regressions of median house price on a referendum "cut vs. maintain" indicator, broken down by Urban and Rural classifications. Columns show the year relative to the referendum vote, estimated treatment effect, standard error, p-value, and 95% confidence interval. Standard errors are robust. A negative estimate indicates lower house prices in areas that cut their road taxes relative to areas that maintain them.

Table 23. Treatment Effect on Housing Prices for Top-Quartile Tax Cuts

Year relative to vote	Estimate	Std. error	p-value	Confidence interval
$t-3$	-2,850	13,435	0.832	[−29,183, 23,482]
$t-2$	-13,984	13,290	0.293	[−40,032, 12,064]
$t-1$	447	12,148	0.971	[−23,363, 24,257]
t	-10,456	15,381	0.497	[−40,603, 19,692]
$t+1$	-6,199	16,738	0.711	[−39,006, 26,608]
$t+2$	-18,657	17,844	0.296	[−53,631, 16,317]
$t+3$	-30,307	21,968	0.168	[−73,365, 12,751]
$t+4$	-31,679	13,072	0.015	[−57,300, −6,058]
$t+5$	-37,063	26,790	0.167	[−89,571, 15,446]
$t+6$	-29,830	15,478	0.054	[−60,167, 508]
$t+7$	-47,933	13,968	0.001	[−75,310, −20,557]
$t+8$	-38,800	16,894	0.022	[−71,913, −5,687]
$t+9$	-44,279	13,410	0.001	[−70,563, −17,995]
$t+10$	-38,659	17,087	0.024	[−72,150, −5,168]

Notes: Supplements Figure 7 in the text. Estimates compare cities that failed to renew a large (top-quartile) road maintenance tax levy with similar cities that renewed. The outcome is median house price in constant 2010 U.S. dollars; the unit of observation is the city–year. Covariates from Table 3 are included in every regression. A treatment effect of −\$31,679 in year $t+4$ means that four years after the vote, treated cities' median sale prices are \$31,679 lower than those of renewing cities.

A1.3 Additional Robustness Tests

A key RD assumption is that observations just above and below the threshold are comparable in all aspects except for treatment status. Table 24 presents balance tests across community characteristics at the voting threshold. As shown, demographic and socioeconomic variables—including population size, poverty rates, educational attainment, employment status, and racial composition—exhibit no statistically significant discontinuities at the cutoff. The absence of any discontinuity suggests that our estimated effects on housing prices represent the impact of failing to renew road tax levies rather than pre-existing community differences. This balance verification via a formal RD test, using covariates as the outcome variable, strengthens our conclusion that the observed housing price effects stem from decisions regarding road tax levies, not from underlying differences in community characteristics.

A1.3.1 Covariate Discontinuity Table

A1.3.2 Covariate Discontinuity Plots

Table 24. Covariate Discontinuity Test Results

Variable	Estimate	Std. error	p-value	Confidence interval
Population	−388	1,094	0.722	[−2,532, 1,755]
Poverty Rate	0.017	0.014	0.234	[−0.011, 0.045]
% with Kids	−0.007	0.012	0.539	[−0.030, 0.015]
% Households with Children < 18	0.0001	0.007	0.981	[−0.014, 0.014]
% Less than High School Education	−0.004	0.020	0.834	[−0.043, 0.035]
% Some College Education	−0.012	0.011	0.274	[−0.034, 0.009]
% Renters	−0.005	0.015	0.754	[−0.035, 0.025]
Unemployment Rate	−0.002	0.006	0.733	[−0.013, 0.009]
% White	−0.007	0.011	0.499	[−0.028, 0.014]
% Black	−0.004	0.009	0.685	[−0.021, 0.014]
% Married	−0.013	0.015	0.374	[−0.042, 0.016]
% Separated	0.001	0.002	0.485	[−0.002, 0.004]
Income Heterogeneity Index	0.007	0.013	0.566	[−0.018, 0.033]
Median Family Income ($)	−3,147	2,963	0.288	[−8,953, 2,660]
% Under 5 Years Old	−0.006	0.004	0.126	[−0.013, 0.002]
% Aged 5 to 17	−0.007	0.007	0.284	[−0.021, 0.006]
% Aged 18 to 64	0.004	0.008	0.611	[−0.012, 0.021]
% Racial Minority	0.007	0.011	0.499	[−0.014, 0.028]

Notes: Each row reports the estimated discontinuity at the levy-approval cutoff for the specified covariate, using robust local-linear RD with the mean effective bandwidth. Confidence intervals are 95%.

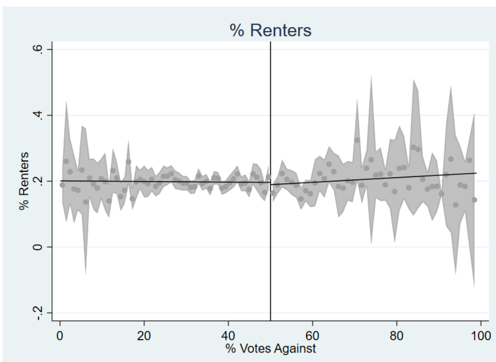

% Renters

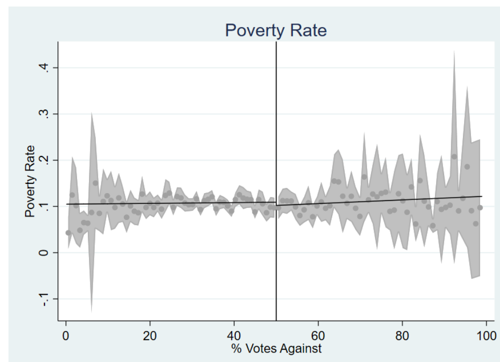

Poverty Rate

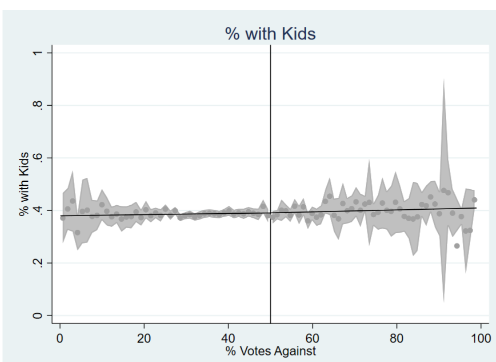

% With Kids

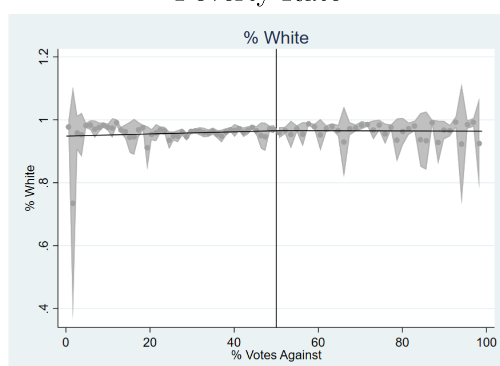

% White

Figure 17. Covariate Discontinuity Plots - Part 1A

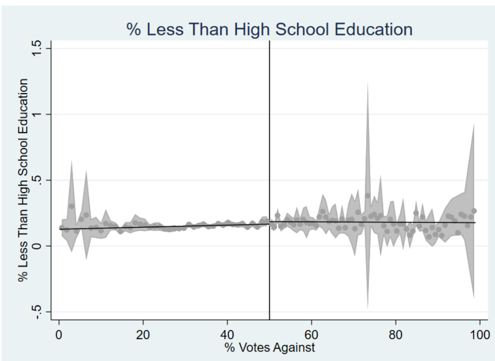

% with Less than High School education

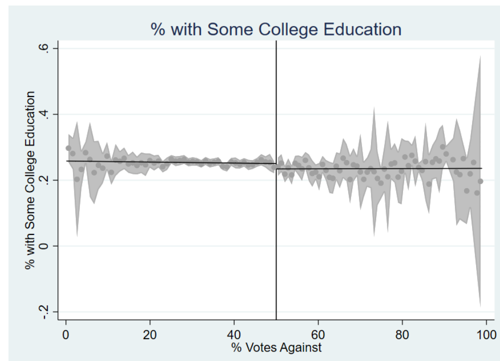

% Attended Some College

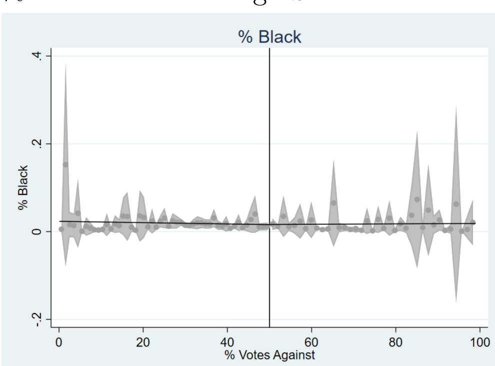

% Black

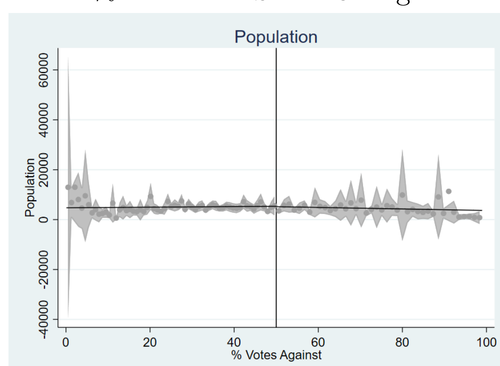

Population

Figure 18. Covariate Discontinuity Plots - Part 1B

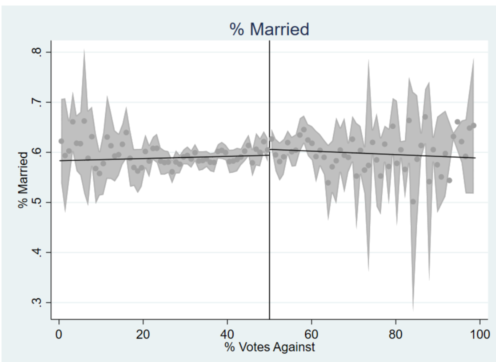

% Married

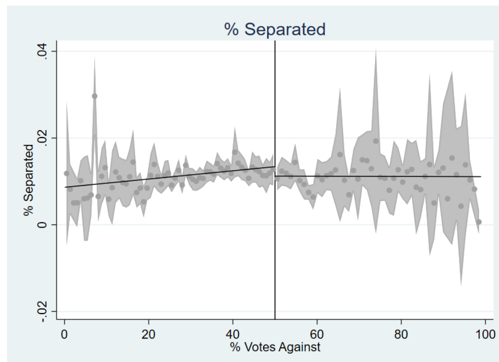

% Separated

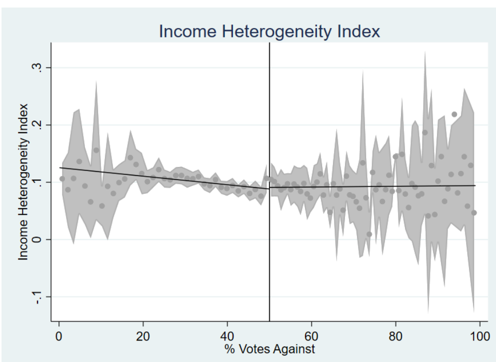

Income Heterogeneity Index

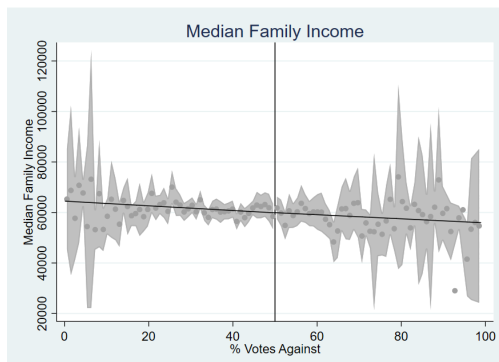

Median Family Income

Figure 19. Covariate Discontinuity Plots - Part 2A

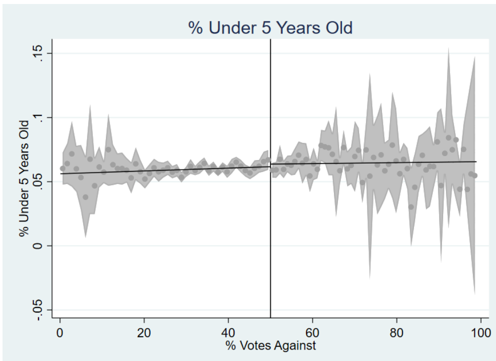

% Less than 5 years old

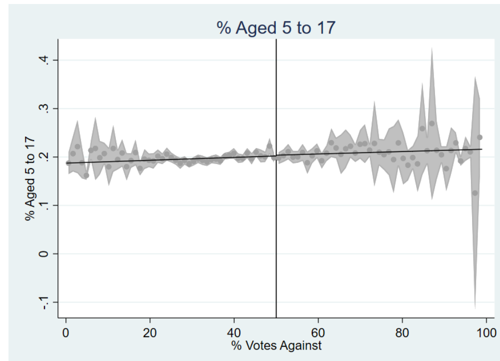

% between 5 to 17 years

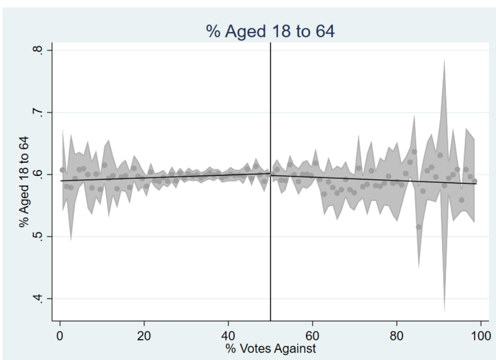

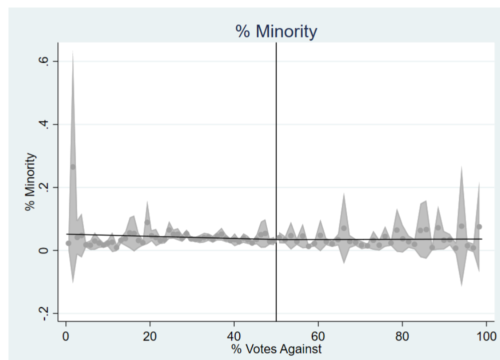

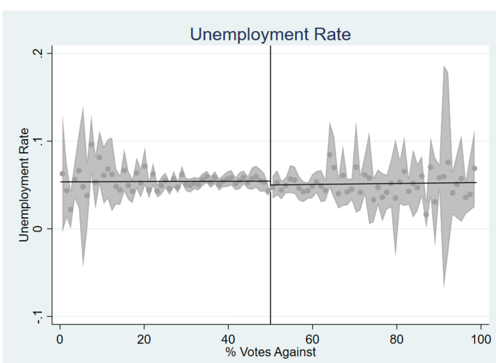

ariate Discontinuity

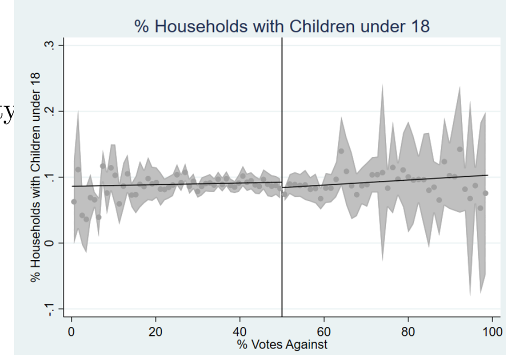

Unemployment Rate

% Households with Children 18 years old or less

Figure 21. Covariate Discontinuity Plots - Part 3

A1.3.3 Hedonic Balance Check

Table 25. Hedonic Balance: Full Sample vs. Close to Cutoff

	Full Sample		Close to Cutoff	
	Treated	Control	Treated	Control
Universal building square feet	1800	1940	1762	1907
	(794)	(874)	(740)	(878)
Acres	1.5	1.6	1.6	1.6
	(5.5)	(5.8)	(6.2)	(5.0)
Total rooms	6.4	6.6	6.4	6.5
	(2.5)	(1.7)	(3.3)	(1.7)
Total baths	2.1	2.3	2.1	2.2
	(1.0)	(1.0)	(1.1)	(1.0)
Age of house	39	38	40	37
	(33)	(32)	(33)	(32)
Air Conditioning	0.91	0.92	0.90	0.91
	(0.28)	(0.27)	(0.30)	(0.28)
Condition: Excellent	0.02	0.02	0.03	0.02
	(0.15)	(0.13)	(0.16)	(0.13)
Condition: Good	0.37	0.34	0.37	0.31
	(0.48)	(0.47)	(0.48)	(0.46)
Condition: Fair	0.06	0.06	0.05	0.06
	(0.23)	(0.24)	(0.22)	(0.24)
Condition: Poor	0.01	0.01	0.01	0.01
	(0.11)	(0.10)	(0.10)	(0.11)
One-story	0.56	0.52	0.52	0.53
	(0.50)	(0.50)	(0.50)	(0.50)

Notes: This table compares hedonic characteristics of houses in treated and control areas for the full sample and for the subsample close to the cutoff (within 40-60% votes against). Standard deviations are in parentheses. Hedonic characteristics include house size (in square feet), lot size (in acres), number of rooms, We examine whether housing characteristics are balanced across the treatment threshold. Table 25 compares hedonic features—building square footage, lot size, rooms, bathrooms, age, air conditioning, condition ratings, and architectural style—between treated (failed renewals) and control (successful renewals) jurisdictions. The results demonstrate balance both in the full sample and the subsample of close elections. The similarity in characteristics near the threshold supports the continuity assumption underlying our identification and validates using narrow election outcomes as quasi-random treatment assignment.

A1.3.4 Bandwidth Sensitivity Analysis

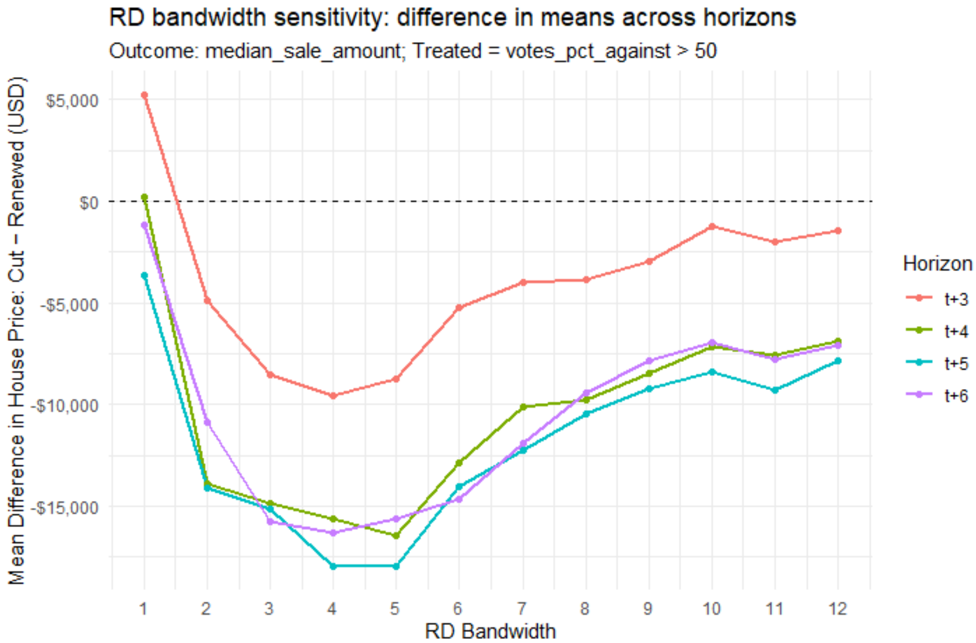

Figure 22. Bandwidth Sensitivity Analysis

Figure 22 plots the differences in mean estimate of the outcome variable, median house prices, between treated and control groups across a grid of bandwidth choices to observe how sensitive the estimates are to the choice of bandwidth. Estimates are consistently negative starting year t+4 and of comparable magnitude over a broad interval around the MSERD optimal bandwidth used in the main analysis. The relative flatness of the curve near the optimal bandwidth indicates that our main result is not driven by a finely tuned bandwidth choice.

A1.4 A Dynamic General Equilibrium Model of Roads

A1.4.1 Representative Household

The economy is populated by a representative, infinitely–lived household that maximizes utility over consumption[56], housing services, and local amenities while disliking labor (including commuting time). In its most general form, the household's problem is to choose sequences of consumption $\{c_t\}_{t=0}^{\infty}$, housing stock $\{h_{t+1}\}_{t=0}^{\infty}$, private capital $\{K_{t+1}\}_{t=0}^{\infty}$, and labor supply $\{n_t\}_{t=0}^{\infty}$ to maximize lifetime utility:

$$\max_{\{c_t,h_t,n_t,K_{t+1},h_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \,\tilde{u}\Big(c_t, h_t, n_t, G_t\Big), \qquad 0 < \beta < 1, \tag{7}$$

We consider the following utility function,

$$\tilde{u}\Big(c_t, h_t, n_t, G_t\Big) \;=\; c_t\; h_t^{\alpha_h A(G_t)}\;\exp\Big[-\frac{\kappa}{1+\theta}\Big(n_t\,[1+\phi(G_t)]\Big)^{1+\theta}\Big] \tag{8}$$

where c_t denotes non–housing consumption, h_t the flow of housing services, and $A(G_t)$ the amenity derived from the public capital stock G_t, which represents the quality of local public goods.

After taking logs, we can write the utility function as a convenient separable specification

$$u\Big(c_t, h_t, n_t, G_t\Big) = \ln c_t + \alpha_h A(G_t)\ln h_t - \frac{\kappa}{1+\theta}\Big[n_t\Big(1+\phi(G_t)\Big)\Big]^{1+\theta}, \tag{9}$$

with n_t hours of work[57], $\phi(G_t) \geq 0$ a commuting factor that rises as road quality falls, $\kappa > 0$, and $\theta > 0$ is the inverse Frisch elasticity of labor supply i.e. $\frac{1}{\theta}$ is the Frisch elasticity of

[56]Since we don't have any random variables or distributions, we are not dealing with any expectations or expected utility.

[57]We model labor disutility using standard isoelastic form, as used in King, Plosser and Rebelo (1988) and Chetty et al. (2011).

labor supply with respect to real wages. The parameter α_h captures the household's preference for housing services, and $A(G_t)$ is a function that captures how public infrastructure quality (e.g., roads) affects the utility derived from housing. The term $\phi(G_t)$ captures the disutility of commuting, which increases as road quality deteriorates.

Budget Constraint and Capital/Housing Accumulation:

Let w_t be the real wage, r_t the rental rate on private capital K_t, $p_{h,t}$ the *pre–tax* purchase price of one unit of housing stock, τ_L the labor–income tax rate, and τ_h the property-tax rate levied on the *value of the housing stock*. The household's per-period budget can be written in the following way:

$$c_t \;+\; i_t \;+\; p_{h,t}\, x_t \;=\; (1-\tau_L) w_t n_t \;+\; r_t K_t \;-\; \underbrace{\tau_h\, p_{h,t}\, h_t}_{\text{property tax on holdings}} \tag{10}$$

where x_t is net housing investment and i_t is private–capital investment. The capital and housing asset laws of motion are

$$K_{t+1} = (1-\delta_K) K_t + i_t, \qquad \delta_K \in (0,1), \tag{11}$$

$$h_{t+1} = (1-\delta_h) h_t + x_t, \qquad \delta_h \in (0,1). \tag{12}$$

Equivalently, we can substitute out i_t and x_t to rewrite the budget constraint as,

$$c_t + K_{t+1} - (1-\delta_K)K_t + p_{h,t}[h_{t+1} - (1-\delta_h)h_t] = (1-\tau_L)w_t n_t + r_t K_t - \tau_h p_{h,t} h_t \tag{13}$$

Optimality Conditions:

Maximization of (9) subject to (10)–(13) yields:

1. **Consumption–Saving Euler Equation:** The household equalizes the marginal utility of consuming an additional unit today to the discounted utility of saving it until tomorrow. This yields:

$$U_c(c_t, h_t, A(G_t)) = \beta U_c(c_{t+1}, h_{t+1}, A(G_{t+1}))\Big[(1-\tau_L)w_{t+1}\frac{\partial n_{t+1}}{\partial K_{t+1}} + r_{t+1} + 1 - \delta_K\Big]. \quad (14)$$

which simplifies (under certainty and taking labor choices as separate) to

$$U_c(c_t, \cdot) = \beta U_c(c_{t+1}, \cdot)\Big(1 + r_{t+1} - \delta_K\Big) \quad (15)$$

In other words, the marginal rate of intertemporal substitution equals the $(1+r_{t+1}-\delta_K)$ factor. This is the standard Euler equation ensuring optimal capital accumulation.

Moreover, given that $U_c = \frac{1}{c_t}$ from our utility specification, we have

$$\frac{1}{c_t} = \beta\,\frac{1}{c_{t+1}}\Big(1 + r_{t+1} - \delta_K\Big); \quad (16)$$

2. **Labor–leisure trade–off:** The household chooses labor n_t such that the marginal disutility of working (including commuting) equals the after-tax marginal benefit of earning wages. Formally,

$$-U_n(c_t, h_t, A(G_t)) = (1-\tau_L)w_t U_c(c_t, h_t, A(G_t)) \quad (17)$$

where U_n is the partial derivative of utility with respect to labor (negative due to disutility) and U_c is the partial derivative with respect to consumption. Using our utility form, we get

$$U_n = -\kappa\Big[n_t(1+\phi(G_t))\Big]^{\theta}(1+\phi(G_t)), \quad U_c = \frac{1}{c_t}. \tag{18}$$

Thus, the labor optimality condition 17 becomes:

$$\kappa\Big[n_t(1+\phi(G_t))\Big]^{\theta}(1+\phi(G_t)) = \frac{(1-\tau_L)w_t}{c_t}. \tag{19}$$

$$\kappa\Big(1+\phi(G_t)\Big)^{1+\theta} n_t^{\theta} = \frac{(1-\tau_L)w_t}{c_t}; \tag{20}$$

$$n_t \;=\; \left[\frac{(1-\tau_L)\,w_t}{\kappa\,c_t\,\Big(1+\phi(G_t)\Big)^{1+\theta}}\right]^{1/\theta} \tag{21}$$

Intuitively, the marginal rate of substitution between leisure (or time) and consumption equals the real after-tax wage. An increase in commuting factor $\phi(G_t)$ (worse roads) raises the left-hand side (effective disutility of labor), leading the household to supply less labor for a given wage. This aligns with the idea that poor road quality reduces effective labor supply as households require higher compensation to work the same hours because commuting erodes their usable time.

3. **Euler equation for Housing:** From our utility function, we have $u_{c,t} = 1/c_t$ and $u_{h,t} = \alpha_h A(G_t)/H_t$. The FOC with respect to H_{t+1} gives

$$p_{h,t}\,u_{c,t} \;=\; \beta\Big[u_{h,t+1} \;+\; u_{c,t+1}\,(1-\delta_h-\tau_h)\,p_{h,t+1}\Big], \tag{22}$$

or, in consumption units,

$$p_{h,t} \;=\; \beta\,\frac{c_t}{c_{t+1}}\Big[(1-\delta_h-\tau_h)\,p_{h,t+1}+\alpha_h\,A(G_{t+1})\,\frac{c_{t+1}}{H_{t+1}}\Big] \tag{23}$$

$$p_{h,t} \;=\; \beta\Big[(1-\delta_h-\tau_h)\,\frac{c_t}{c_{t+1}}\,p_{h,t+1}+\alpha_h\,A(G_{t+1}(\tau_h))\,\frac{c_t}{H_{t+1}}\Big]. \tag{24}$$

The equation 24 encapsulates how the current price of housing $p_{h,t}$ reflects the discounted value of future benefits from owning housing and the benefits associated with holding it. The bracketed terms are, respectively, the *after–tax, after–depreciation resale value* next period and the *service dividend* (marginal value of one more unit of housing services) next period. The first term inside the brackets represents the discounted future resale value of the housing asset after accounting for depreciation and property taxes. The second term captures the flow of housing services (amenities) derived from owning the house, adjusted for the quality of public infrastructure $A(G_{t+1}(\tau_h))$ and scaled by consumption per unit of housing stock. Notice that the amenity value $A(G_{t+1}(\tau_h))$ depends on the future road quality, which in turn is influenced by the current property tax τ_h through its effect on maintenance spending, as explained in subsequent subsections. A cut in τ_h today leads to lower maintenance M_t, causing G_{t+1} to decline over time, which reduces $A(G_{t+1}(\tau_h))$ and thus lowers the service dividend from housing. This mechanism captures how changes in public infrastructure quality, driven by tax policy, affect housing prices through their impact on the flow of housing services.

A1.4.2 Representative Firm

A competitive firm produces output Y_t with Cobb–Douglas technology

$$Y_t = K_t^{\alpha_k} (\mathcal{A}_t L_t)^{1-\alpha_k}, \qquad 0 < \alpha_k < 1, \tag{25}$$

enforcing factor–price equalization

$$w_t = (1-\alpha_k)\frac{Y_t}{L_t}, \qquad r_t = \alpha_k \frac{Y_t}{K_t}. \tag{26}$$

Profit Maximization. The firm operates in a competitive market for inputs and output, so it hires labor and rents capital until factor prices equal their marginal products. In each period, the firm solves

$$\max_{K_t, L_t} \Big\{ F(K_t, L_t) - w_t L_t - r_t K_t \Big\}, \tag{27}$$

taking wage w_t and capital rental r_t as given. The first-order conditions are:

1. **Labor demand:**

$$w_t = (1-\alpha_k) K_t^{\alpha_k} (\mathcal{A}_t L_t)^{-\alpha_k} = (1-\alpha_k)\frac{Y_t}{L_t}. \tag{28}$$

   This is the usual result that the real wage equals the marginal product of labor (MP_L).

2. **Capital demand:**

$$r_t = \alpha_k K_t^{\alpha_k - 1} (\mathcal{A}_t L_t)^{1-\alpha_k} = \alpha_k \frac{Y_t}{K_t}. \tag{29}$$

   Thus, the rental rate on capital equals the marginal product of capital (MP_K).

Under constant returns and competitive markets, the firm earns zero economic profit in equilibrium (all output is paid out to factors). This means

$$Y_t = w_t L_t + r_t K_t \tag{30}$$

in equilibrium. Because poor local roads can indirectly reduce effective labor supplied (households may choose smaller L_t) or even the productivity of labor (if $\mathcal{A}_t$ were to depend on G_t), the equilibrium wage and output will adjust accordingly. In our model, we focus on the household-side mechanism; we do not explicitly insert G_t into $F(\cdot)$ so $\mathcal{A}_t$ is treated as exogenous, but one could imagine an extension where $\mathcal{A}_t = \mathcal{A}(G_t)$ similar to household amenities function, making public infrastructure a productive externality that boosts TFP. That would reinforce the mechanism: a decline in G_t would lower productivity and thus lower wages from the firm side as well. For now, wages change endogenously mainly due to labor supply shifts in the simple model.

A1.4.3 Government and Public Infrastructure

Taxation and Spending. The government collects taxes and uses the revenue to finance road maintenance. Taxes have two components: an exogenous property tax τ_h (subject to voter approval) and an endogenous tax (such as τ_L on labor income, or potentially a lump-sum tax) that provides additional revenue. Denote by T_t^{exo} the exogenous tax revenue from the voted levy, and T_t^{endo} the endogenous tax revenue from other sources. For example, if τ_L is a labor tax, then $T_t^{\text{endo}} = \tau_L w_t n_t$, which varies with the economy and is therefore "endogenous". The property tax revenue is $T_t^{\text{exo}} = \tau_h p_{h,t} h_t$ each period, directly tied to the housing market outcome.

The government budget constraint is:

$$T_t^{\text{endo}} + T_t^{\text{exo}} = M_t, \tag{31}$$

assuming all tax revenue is spent on road maintenance M_t. We abstract from any other government consumption or transfers.

In a scenario where the road tax is cut (e.g., τ_h drops to zero after a failed levy vote), T_t^{exo} falls exogenously. For a balanced budget, maintenance spending M_t must be reduced by an equal amount. This captures the immediate fiscal impact of the tax cut, such as an observed 11% loss in road maintenance funding empirically.

Public Capital (Roads) Accumulation. The stock of public road infrastructure G_t evolves over time depending on maintenance. We adopt a law of motion in line with Rioja (2003) to formalize how insufficient maintenance leads to infrastructure depreciation, whereas sufficient maintenance can maintain road quality[58]. Let δ_G be the "normal" depreciation rate of roads (wear-and-tear if adequately maintained). Actual effective depreciation $\delta_{G,t}$ can exceed δ_G if maintenance is below the requirement. A convenient formulation is:

$$G_{t+1} = G_t(1 - \delta_{G,t}), \tag{32}$$

where

$$\delta_{G,t} = \varphi \cdot \max\left\{0, 1 - \frac{M_t}{\delta_G G_t^{\psi}}\right\}. \tag{33}$$

In words, if maintenance spending M_t falls short of the amount needed to cover normal depreciation ($\delta_G G_t$), the depreciation rate increases proportionally to the shortfall. The parameter $\varphi \geq 0$ governs how sensitive the infrastructure is to under-maintenance. For example, if M_t is only half of $\delta_G G_t$, then $1 - \frac{M_t}{\delta_G G_t} = 0.5$, and $\delta_{G,t} = 0.5\varphi$; this means roads wear out. If M_t is zero (no upkeep), depreciation could jump substantially (if φ is large, G_t may deteriorate very quickly). On the other hand, if M_t is at least $\delta_G G_t$ (maintenance fully covers yearly wear), then $\delta_{G,t} = 0$ (no deterioration). We also assume $\delta_{G,t}$ cannot go below 0 even if M_t exceeds $\delta_G G_t$ (any extra maintenance beyond replacing depreciation might marginally improve roads but with diminishing returns). This Riojas-style law of motion captures the intuition that lack of maintenance shortens the life of existing public capital. Maintenance spending effectively slows down depreciation, preserving the infrastructure stock; cutting maintenance causes G_t to decline faster over time. Importantly, a reduction in the road tax τ_h translates to lower M_t and thus a gradual decline in G_t over several periods.

[58]In this setup, roads can only be maintained, not improved.

The deterioration in road quality will typically become noticeable after a few periods of under-maintenance, consistent with evidence that the negative effects on road conditions (and subsequently house prices) emerge with a lag. In our model, that lag is endogenously determined by the above law of motion – if M_t drops at $t = 0$, G_t will depreciate a bit faster each period, compounding into a significant drop in infrastructure quality after some years.

A1.4.4 Housing Market

Housing is a tradable asset in *fixed* aggregate supply $\bar{h}$[59]. With a representative household, market clearing is

$$h_t \;=\; \bar{H} \quad \text{for all } t. \tag{34}$$

Given (34), the price $p_{h,t}$ is determined by the housing Euler (23) and the paths of $\{c_t, G_t\}$. In steady state, substituting $c_{t+1}=c_t=c$, $h_{t+1}=h_t=h=\bar{H}$, $G_{t+1}=G_t=G$, $p_{h,t+1}=p_{h,t}=p_h$ into (23) yields the user–cost formula:

$$p_h \;=\; \frac{\beta}{1-\beta(1-\delta_h-\tau_h)}\alpha_h \, A(G)\, \frac{c}{h}. \tag{35}$$

Better roads i.e. higher $A(G)$ higher raise p_h. Larger h reduces marginal service value (log utility), lowering p_h ceteris paribus.

The steady-state housing price formula in equation (35) reveals how the discount factor β interacts with depreciation δ_h and property taxes τ_h to determine asset valuation. The term $\frac{\beta}{1-\beta(1-\delta_h-\tau_h)}$ represents the present value multiplier that converts the flow of housing services into an asset price.

To understand this multiplier, consider that $1 - \delta_h - \tau_h$ represents the net retention rate of housing value after accounting for physical depreciation and tax obligations. The household

[59]This means that the total quantity of housing is constant over time, and any changes in demand will only affect prices. We can also allow extensions to the model that incorporate changes in housing supply.

effectively loses fraction $\delta_h + \tau_h$ of the housing asset's value each period through wear and tax payments. The remaining fraction $1 - \delta_h - \tau_h$ carries forward to the next period, where it generates both service flows and continuation value.

The economic intuition is clear: higher property taxes τ_h reduce the net retention rate, making housing a less attractive store of value and lowering equilibrium prices. A patient household (high β) values future service flows more highly, increasing the multiplier and thus housing prices. Conversely, rapid depreciation δ_h erodes the asset's continuation value, reducing prices through the same channel.

This formulation also shows why property tax cuts can have ambiguous effects on housing prices in our model. While lower τ_h directly increases the multiplier and boosts prices, it also reduces maintenance funding, leading to infrastructure decay that lowers the service flow $\alpha_h A(G)$. The net effect depends on the relative magnitudes of these opposing forces and their timing.

A1.4.5 Competitive Equilibrium

We now define a recursive competitive equilibrium for this economy. Given an initial private capital stock K_0 and initial public capital (road quality) G_0, and a sequence of exogenous tax policy $\{\tau_{h,t}\}$ (with a drop in τ_h at the time of the tax cut shock), a competitive equilibrium is a sequence of allocations $\{c_t, h_{t+1}, n_t, K_{t+1}, M_t, G_t\}_{t\geq 0}$ and prices $\{w_t, r_t, p_{h,t}\}_{t\geq 0}$ such that:

1. **Household Optimization:** Given prices and taxes, the representative household chooses $\{c_t, h_{t+1}, n_t, K_{t+1}\}$ to maximize its utility subject to its budget and accumulation constraints. The household's choice satisfies the first-order optimality conditions described above (Consumption and housing Euler equations, labor supply condition, and housing demand condition, etc.), and the transversality condition on capital holding.

2. **Firm Optimization:** The representative firm chooses inputs $\{K_t, L_t\}$ each period to maximize profit. In equilibrium, $L_t = n_t$ (labor market clears) and $K_t = K_t$ held by the household (capital market clears), and the wage and rental rates satisfy $w_t = F_L(K_t, L_t)$ and $r_t = F_K(K_t, L_t)$ as given by the Cobb-Douglas marginal products. This ensures the firm's first-order conditions are met and profit is zero.

3. **Government Budget and Public Capital:** The government budget constraint holds each period: $T_t^{\text{endo}} + T_t^{\text{exo}} = M_t$, with $T_t^{\text{exo}} = \tau_{h,t} p_{h,t} h_t$ and $T_t^{\text{endo}} = \tau_L w_t n_t$. The road maintenance M_t feeds into the law of motion for G_t: given G_t at the start of period t, the next period's public capital G_{t+1} is determined by $G_{t+1} = G_t(1 - \delta_{G,t})$ with $\delta_{G,t}$ increasing if M_t is insufficient (as specified above). The government sets M_t according to the available tax revenue each period.

4. **Market Clearing:** All markets clear in equilibrium:

   - **Goods market:** Output is used for private consumption, private investment, and public maintenance. The resource constraint each period is:

     $$Y_t = c_t + i_t + M_t. \tag{36}$$

     Total output Y_t produced by the firm equals the sum of household consumption c_t, private investment i_t, and government demand for road maintenance M_t.

   - **Labor market:** $L_t = n_t$. The labor supplied by the household equals labor demanded by the firm. The real wage adjusts such that this holds.

   - **Capital market:** The firm's capital input equals the household's capital stock: K_t (used in production) $= K_t$ (owned by household). The rental rate r_t adjusts to clear this market.

- **Housing market:** $h_t = \bar{H}$. The total housing demanded by the household equals the fixed housing stock. The house price $p_{h,t}$ adjusts each period to clear the housing market.

5. **Consistency:** The expectations of the household are consistent with realized outcomes. In a deterministic steady-state scenario, this means the household correctly anticipates the path of prices and G_t. In a perfect-foresight or rational expectations equilibrium, the household foresees that a cut in τ_h implies lower future M_t and thus a gradually declining G_t, factoring this into its decisions. The sequence $\{p_{h,t}\}$ reflects the present value of housing services given the expected decline in amenities.

A1.4.6 Solving the model

The steady-state equilibrium can be characterized as a system of 9 equations in 9 unknowns. We solve for the following endogenous variables: $(c, n, K, Y, w, r, q, M, G)$, where $q \equiv (1+\tau_h)p_h$ is the effective housing price and housing quantity is fixed at $h = \bar{H}$.

A1.4.7 System of Equilibrium Equations

The steady-state equilibrium is determined by the following 9 independent equations:

1. **Euler equation (steady-state):** From the intertemporal optimality condition:

$$1 = \beta(1 + r - \delta_K) \quad \Rightarrow \quad r = \delta_K + \beta^{-1} - 1 \tag{37}$$

This pins down the rental rate r directly from parameters.

2. **Firm capital pricing:** From profit maximization:

$$r = \alpha_k \frac{Y}{K} \tag{38}$$

This relates the output-capital ratio to the rental rate.

3. **Firm wage pricing:** With labor market clearing $L = n$:

$$w = (1 - \alpha_k) \frac{Y}{n} \tag{39}$$

This relates the wage to output and labor supply.

4. **Production technology:** The Cobb-Douglas production function:

$$Y = K^{\alpha_k} (A \cdot n)^{1-\alpha_k} \tag{40}$$

This ties together output, capital, and labor.

5. **Labor-leisure trade-off:** From household optimization:

$$\kappa [1 + \phi(G)]^{1+\theta} n^{\theta} = \frac{(1 - \tau_L) w}{c} \tag{41}$$

This links labor supply to wages, consumption, and road quality through commuting costs.

6. **Housing price determination:** From the housing Euler equation in steady state:

$$q = \frac{\beta}{1 - \beta(1 - \delta_h - \tau_h)} \cdot \alpha_h A(G) \frac{c}{\bar{H}} \tag{42}$$

This determines the housing price from the present value of housing services, incorporating the discount factor and net retention rate.

7. **Government budget constraint:** Balanced budget condition:

$$M = \tau_L w n + \tau_h p_h \bar{H} \tag{43}$$

This defines maintenance spending from tax revenues.

8. **Road quality steady-state:** From the law of motion $G_{t+1} = G_t(1 - \delta_{G,t})$:

$$G = G(1 - \delta_G) \quad \text{where} \quad \delta_G = \varphi \cdot \max\left\{0, 1 - \frac{M}{\delta_G^{\text{norm}} G^{\psi}}\right\} \tag{44}$$

In the minimal-upkeep steady state, this simplifies to $M = \delta_G^{\text{norm}} G^{\psi}$.

9. **Resource constraint:** Goods market clearing with capital accumulation:

$$c = Y - \delta_K K - M \tag{45}$$

Note that $i = K - (1 - \delta_K)K = \delta_K K$ in steady state. This accounts for the allocation of output between consumption, investment, and maintenance.

A1.4.8 Transitional Dynamics and Welfare Analysis

To analyze the transitional dynamics following a tax cut, we solve the model numerically two different ways: under perfect foresight[60] and partial adjustment[61]. The key steps are:

1. **Calibration:** Assign parameter values based on empirical estimates or literature benchmarks. For example, β, α_h, η, κ, θ, δ_K, δ_G, and φ are calibrated to match observed household behavior, infrastructure depreciation rates, and labor supply elasticities.

[60] Households are forward-looking and adjust their behavior immediately to changes in policy and economic conditions.

[61] Households adjust their behavior gradually in response to changes in policy and economic conditions.

2. **Initial Steady State:** Compute the pre-tax-cut steady state by solving the equilibrium conditions under the initial τ_h. This provides baseline values for $\{M_t, G_t, p_{h,t}, c_t, n_t, Y_t\}$.

3. **Shock Implementation:** Introduce the tax cut by reducing τ_h at $t = 0$. Solve the model forward to trace the path of endogenous variables $\{M_t, G_t, p_{h,t}, c_t, n_t, Y_t\}$ as the economy transitions to the new steady state.

4. **Dynamic Path:** Use one of the methods stated above to compute the full transition path of $\{M_t, G_t, p_{h,t}, c_t, n_t, Y_t\}$. Households anticipate the future decline in G_t due to lower maintenance and adjust their consumption, labor supply, and housing demand accordingly. The path of G_t is updated each period using the law of motion in equation 32.

5. **Welfare Analysis:** Compute the household's lifetime utility before and after the tax cut. Decompose the welfare change into short-run gains (higher disposable income) and long-run losses (lower amenities and higher commuting costs).

6. **Sensitivity Analysis:** Vary key parameters (e.g., φ, α_h, θ) to test the robustness of the results. Examine how the magnitude of the tax cut or alternative government policies affect outcomes.

The numerical solution involves iterating on the household's Euler equations, firm conditions, and government budget constraint while updating G_t using its law of motion. This allows us to simulate the gradual decline in road quality and its impact on housing prices, labor supply, and overall welfare.

A1.4.9 Discussion of Policy Shock

A permanent cut in the voted levy lowers τ_h exogenously. In the short run, disposable income rises as tax burden decreases, but G_t is initially unchanged. This leads to a sharp increase in consumption but also a sharp decrease in maintenance budget. Over time, reduced maintenance spending accelerates depreciation in equation 32, lowering $A(G_t)$ and raising $\phi(G_t)$; this depresses labor supply, housing demand, and therefore equilibrium house prices, as demonstrated in Section A1.4.11.

A1.4.10 Model Calibration

The model parameters are calibrated based on empirical evidence and standard values from the literature. Table 26 summarizes the calibrated values used in the analysis. These values are chosen to reflect realistic economic conditions. For example, $\beta = 0.96$ corresponds to a 4% annual discount rate, while $\delta_K = 0.05$ reflects typical depreciation rates for private capital. Other parameters such as labor disutility parameters κ and θ are calibrated to match observed labor supply elasticities, and φ captures the empirical sensitivity of road quality to maintenance shortfalls.

A1.4.11 Model Results

The model is solved numerically using both a perfect foresight and a partial adjustment algorithm to simulate the economy's response to a permanent cut in the road maintenance tax τ_h. Below, we present the key results from partial adjustment simulations[62].

- A reduction in τ_h from 1% to 0% leads to an decrease in maintenance budget M_t and road quality G_t, which declines gradually over time due to the law of motion for public capital.

[62] We focus on partial adjustment results here as they are better equipped for illiquid housing markets and gradual behavioral responses. Perfect foresight results are qualitatively similar but exhibit sharper short-run dynamics.

Table 26. Baseline calibration and sensitivity ranges

Parameter	Model block	Baseline	Sensitivity range	Key empirical source
β	Household intertemporal utility	0.96	0.94–0.99	Cooley (1995)
α_h	utility weight on housing services	0.35	0.25–0.40	(U.S. Bureau of Labor Statistics, 2024)
η	Amenity elasticity	0.88		†
θ	Inverse Frisch elasticity labor supply	2	1.5–4.0	Chetty et al. (2011)
κ	Labor-disutility scale	2.9	2.9	Chatterjee, Gibson and Rioja (2017)
α_k	Private-capital share	0.30	0.25–0.40	Chatterjee, Gibson and Rioja (2017)
δ_K	Private Capital depreciation (annual)	0.06	0.04–0.08	Cooley (1995)
δ_h	Housing depreciation (annual)	0.05	0.04–0.08	
δ_G^{norm}	Road stock depreciation	0.25		†
φ	Under-maintenance sensitivity	0.64	0.25–0.75	Rioja (2003)
γ	commuting disutility sensitivity	0.47		†
ψ	maintenance requirement elasticity	0.47		†
τ_L	Labor income tax share for roads	0.35		†

Note: † Calibrated to match empirical moments

- The house price $p_{h,t}$ first increases and then falls significantly as the amenity value $A(G_t)$ declines, reflecting the reduced desirability of housing in areas with deteriorating roads

The figures 23 illustrates the dynamics of key variables over time following the tax cut. The model's transitional dynamics following a permanent reduction in the property tax from 1 mills to 0 mills.

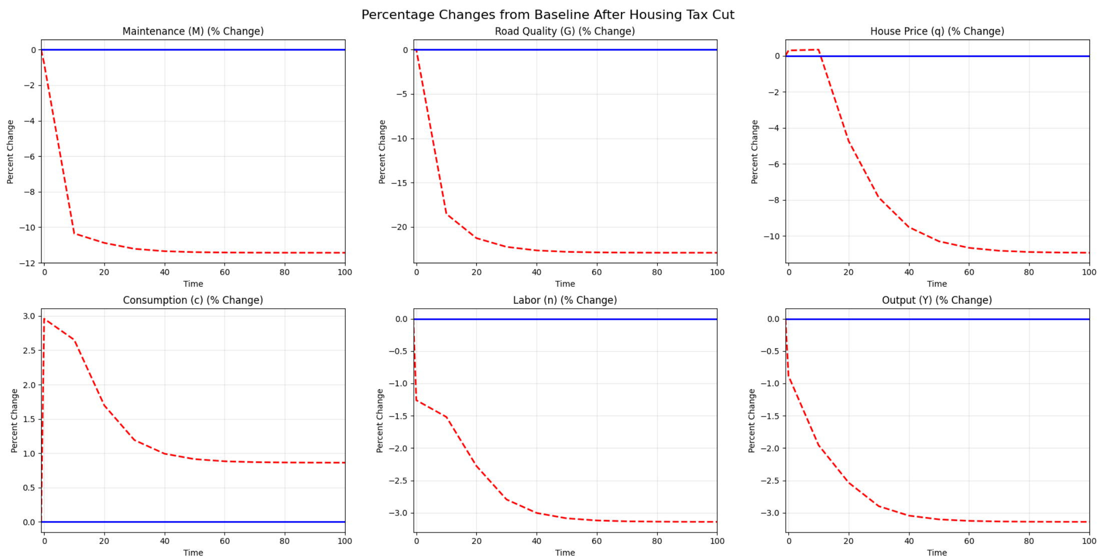

Figure 23. Dynamics of variables following the tax cut.

Figure 23 plots the full transition paths (percentage deviations from baseline). Most of the adjustment in M_t and q_t occurs early, while G_t converges more gradually as under-maintenance compounds over time.Findings from Figure 23 are summarized below:

- **Roads and maintenance.** Following the road tax cut, the road maintenance budget drops by about 10% and then drifts to roughly -11% (top-left panel of Figure 23). Given the law of motion for roads, G_t declines gradually and monotonically, approaching a long-run decline to 21% by the end of the simulated horizon.

- **House prices.** The housing price q_t initially shows a small blip (about +0.3%) and then declines smoothly toward a long-run drop of an average of 9% over the simulated horizon as the amenity value $A(G_t)$ deteriorates.

- **Aggregates.** The macro aggregates move modestly: labor n_t falls by about -3.0%, output also Y_t by about -3%, and consumption c_t increases by about 3% in the long run. Consumption exhibits a slight short-run overshoot (of +3%) as households initially respond to higher disposable income from the tax cut, before settling slightly above baseline (bottom row). This reflects the household's initial response to higher disposable income from the tax cut, followed by adjustments and substitution effect as road quality declines.

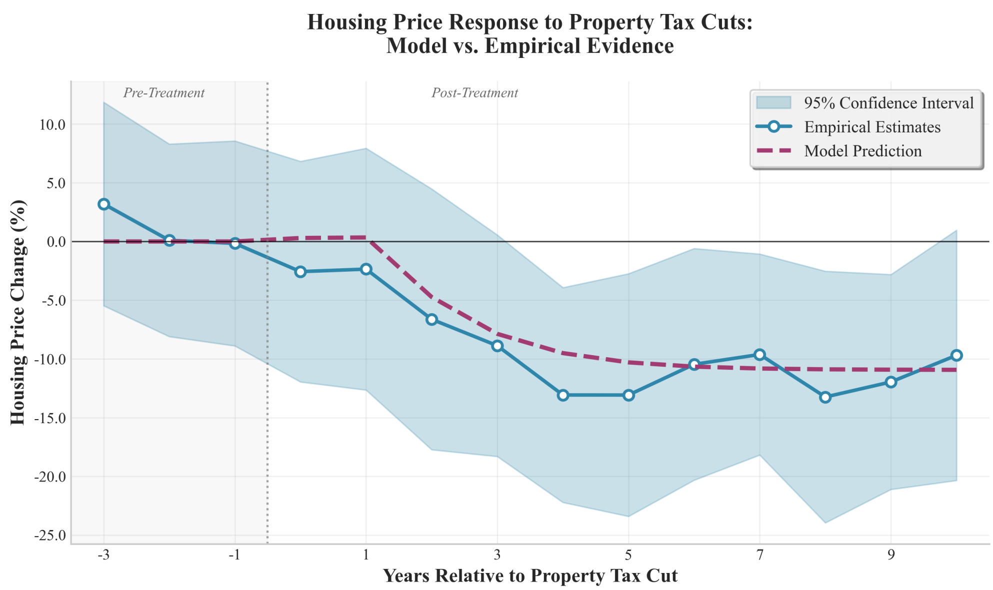

Figure 24. Model vs Empirical Estimates of House Prices Following the Tax Cut.

Model vs. empirical house-price path. Figure 24 compares the model-implied house-price response to the event-study estimates in Figure 5. After aligning the time units, the model tracks the medium- and long-run decline well, squarely within the empirical confidence band. The model is slightly conservative initially as the drop in housing prices occurs more gradually than the empirical estimates. Nevertheless, the overall magnitude and timing of the decline remain consistent with the data as the model starts to deviate from the baseline around year 2 and approaches a long-run average drop of about 9% by year 10. This pattern is consistent with the mechanism presented in the paper: eliminating the tax reduces maintenance revenues, which gradually lowers road quality, depresses amenity $A(G_t)$, and in turn pulls down housing valuations.

Overall, the calibrated model reproduces the targeted steady-state effects (about -21% in roads, -9% in house prices, and -11% in maintenance) and generates transitional dynamics for house prices that closely follow the empirical path while remaining slightly less sharp at the short-run trough.

A1.4.12 Model Limitations

1. **Representative agent assumption.** Heterogeneity in income or commuting preferences could attenuate price effects. Sensitivity analyses with quasi-linear preferences or idiosyncratic $\phi(G_t)$ show qualitative robustness.

2. **No spatial heterogeneity.** We currently ignore spatial heterogeneity. The model intentionally abstracts from space for tractability. Future work can embed G_t in a Rosen–Roback framework to capture cross-location sorting.

3. **Fixed housing stock.** Allowing new construction would introduce a supply elasticity that dampens price responses. Empirical evidence suggests housing supply in many U.S. jurisdictions is quite inelastic.

4. **No direct productivity role for roads.** Incorporating G_t into firm TFP $\mathbb{A}_t$ would further enrich the model. However, this would require additional assumptions about how road quality affects firm location and productivity, which is beyond our current scope.

The DGE framework formalizes the intuition that road–maintenance tax cuts provide immediate relief but, by undermining infrastructure quality, erode local amenities and raise commuting costs, culminating in lower house values and welfare.

A1.4.13 Derivations of Key Equations

A1.4.14 Derivation of the Euler Equation for consumption 16

Below $U_c \equiv \partial u/\partial c$ and λ_t is the Lagrange multiplier on the period-t budget constraint.

Set up the Lagrangian:

$$L = \sum_{t=0}^{\infty} \beta^t \Big[u(c_t, h_t, A(G_t)) + \lambda_t \Big((1-\tau_L) w_t n_t + r_t K_t - \tau_h p_{h,t} h_t - c_t - p_{h,t} h_t - K_{t+1} + (1-\delta_K) K_t \Big) \Big].$$

FOC w.r.t. consumption c_t:

$$\frac{\partial L}{\partial c_t} = 0 \quad \Longrightarrow \quad \beta^t U_c(c_t, h_t, A(G_t)) - \lambda_t = 0 \quad \Longrightarrow \quad \lambda_t = \beta^t U_c(c_t, h_t, A(G_t)). \tag{A1}$$

FOC w.r.t. next-period capital K_{t+1}:

$$\frac{\partial L}{\partial K_{t+1}} = 0 \quad \Longrightarrow \quad -\lambda_t + \lambda_{t+1} \Big[(1-\tau_L) w_{t+1} \frac{\partial n_{t+1}}{\partial K_{t+1}} + r_{t+1} + 1 - \delta_K \Big] = 0. \tag{A2}$$

Combine (A1) and (A2):

Insert λ_t and λ_{t+1} from (A1) into (A2):

$$\beta^t U_c(c_t, h_t, A(G_t)) = \beta^{t+1} U_c(c_{t+1}, h_{t+1}, A(G_{t+1})) \Big[(1-\tau_L) w_{t+1} \frac{\partial n_{t+1}}{\partial K_{t+1}} + r_{t+1} + 1 - \delta_K \Big].$$

Divide both sides by β^t to obtain:

$$U_c(c_t, h_t, A(G_t)) = \beta U_c(c_{t+1}, h_{t+1}, A(G_{t+1})) \Big[(1-\tau_L) w_{t+1} \frac{\partial n_{t+1}}{\partial K_{t+1}} + r_{t+1} + 1 - \delta_K \Big].$$

Why the derivative $\frac{\partial n_{t+1}}{\partial K_{t+1}}$? Capital chosen today affects tomorrow's labor choice only through the household's optimal plan. If, as we assume, labor in each period is chosen independently of the inherited capital stock, this derivative is zero and the term drops out, giving the familiar

$$U_c(c_t, \cdot) = \beta U_c(c_{t+1}, \cdot)\Big(1 + r_{t+1} - \delta_K\Big),$$

which is the standard Euler equation for intertemporal consumption smoothing.

Equation 10 still contains the marginal utility of consumption $U_c(\cdot)$. Because the one-period utility function chosen is:

$$u(c_t, h_t, A(G_t)) = \ln c_t + \alpha_h \ln h_t + \eta \ln A(G_t) - \frac{\kappa}{1+\theta}\Big[n_t(1 + \phi(G_t))\Big]^{1+\theta}, \tag{6}$$

its partial derivative with respect to consumption is:

$$U_c(c_t, h_t, A(G_t)) = \frac{\partial u}{\partial c_t} = \frac{1}{c_t}. \quad \text{(log utility)}$$

Substituting U_c into the Euler Equation:

Start from Eq. 10:

$$U_c(c_t, \cdot) = \beta\, U_c(c_{t+1}, \cdot)\Big(1 + r_{t+1} - \delta_K\Big). \tag{10}$$

Replace U_c with $\frac{1}{c}$ (log utility):

$$\frac{1}{c_t} = \beta\, \frac{1}{c_{t+1}}\Big(1 + r_{t+1} - \delta_K\Big).$$

This is exactly Eq. 11:

$$\boxed{\frac{1}{c_t} = \beta\, \frac{1}{c_{t+1}}\Big(1 + r_{t+1} - \delta_K\Big)}. \tag{11}$$

A1.4.15 Labor–leisure trade–off

To derive the labor-leisure trade-off condition, we start by computing the partial derivatives of the utility function with respect to consumption and labor:

For consumption:

$$u_c(c_t, h_t, n_t, G_t) = \frac{1}{c_t} \tag{B1}$$

For labor, we have:

$$u_n(c_t, h_t, n_t, G_t) = -\,\kappa\Big[n_t\Big(1+\phi(G_t)\Big)\Big]^{\theta}\Big(1+\phi(G_t)\Big). \tag{B2}$$

Optimal labour supply equates the marginal disutility of work to the after-tax marginal benefit of wages

$$-\,u_n(c_t, h_t, n_t, G_t) \;=\; (1-\tau_L)\,w_t\,u_c(c_t, h_t, n_t, G_t)$$

Using B1 and B2, we can rewrite this as:

$$\kappa\Big[n_t\Big(1+\phi(G_t)\Big)\Big]^{\theta}\Big(1+\phi(G_t)\Big) \;=\; \frac{(1-\tau_L)\,w_t}{c_t}$$

This is exactly the condition we wanted to derive:

$$\kappa\Big(1+\phi(G_t)\Big)^{1+\theta}\,n_t^{\theta} \;=\; \frac{(1-\tau_L)\,w_t}{c_t}$$

$$n_t \;=\; \left[\frac{(1-\tau_L)\,w_t}{\kappa\,c_t\,\Big(1+\phi(G_t)\Big)^{1+\theta}}\right]^{1/\theta}$$

A1.4.16 Euler equation for Housing

We derive the housing Euler starting from the household's problem.

1. Primitives. Per–period utility:

$$u(c_t, h_t, n_t, G_t) = \ln c_t + \alpha_h A(G_t) \ln h_t - \frac{\kappa}{1+\theta}\Big[n_t(1+\phi(G_t))\Big]^{1+\theta}.$$

Marginal utilities (holding G_t fixed):

$$u_{c,t} = \frac{1}{c_t}, \qquad u_{h,t} = \frac{\partial u}{\partial h_t} = \alpha_h \frac{A(G_t)}{h_t}.$$

2. Budget constraint in stock form. Using $H_{t+1} = (1-\delta_h)H_t + x_t$ and substituting $x_t = H_{t+1} - (1-\delta_h)H_t$, the period-$t$ budget is

$$c_t + i_t + p_{h,t}\Big(H_{t+1} - (1-\delta_h)H_t\Big) = (1-\tau_L)w_t n_t + r_t K_t - \tau_h p_{h,t} H_t.$$

Rearrange the housing terms:

$$-p_{h,t}H_{t+1} + p_{h,t}(1-\delta_h)H_t - \tau_h p_{h,t} H_t = -p_{h,t}H_{t+1} + p_{h,t}(1-\delta_h-\tau_h)H_t.$$

3. Lagrangian. Let λ_t be the multiplier on the budget constraint:

$$\begin{aligned}\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \Big\{ & u(c_t, h_t, n_t, G_t) \\ & + \lambda_t\Big[(1-\tau_L)w_t n_t + r_t K_t - c_t - i_t - p_{h,t}(H_{t+1} - (1-\delta_h)H_t) \\ & - \tau_h p_{h,t} H_t - (K_{t+1} - (1-\delta_K)K_t)\Big]\Big\}\end{aligned} \tag{46}$$

4. FOCs for c_t and H_{t+1}. Consumption:

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 \;\Rightarrow\; \lambda_t = u_{c,t}.$$

For H_{t+1} (differentiating with respect to the housing stock chosen for next period):

$$\frac{\partial \mathcal{L}}{\partial H_{t+1}} = 0 : \quad \beta^{t+1} u_{h,t+1} + \beta^{t+1} \lambda_{t+1} p_{h,t+1}(1-\delta_h-\tau_h) - \beta^t \lambda_t p_{h,t} = 0.$$

Divide by β^t and rearrange:

$$p_{h,t}\lambda_t = \beta\Big[u_{h,t+1} + \lambda_{t+1} p_{h,t+1}(1-\delta_h-\tau_h)\Big].$$

5. Replace multipliers with marginal utility of consumption. Since $\lambda_t = u_{c,t}$:

$$p_{h,t} u_{c,t} = \beta\Big[u_{h,t+1} + u_{c,t+1}(1-\delta_h-\tau_h) p_{h,t+1}\Big],$$

which is equation (22).

6. Convert to consumption units. Divide both sides by $u_{c,t+1}$:

$$p_{h,t}\frac{u_{c,t}}{u_{c,t+1}} = \beta\Big[(1-\delta_h-\tau_h)p_{h,t+1} + \frac{u_{h,t+1}}{u_{c,t+1}}\Big]$$

With $u_{c,t} = 1/c_t$ we have $\frac{u_{c,t}}{u_{c,t+1}} = \frac{c_{t+1}}{c_t}$, so

$$p_{h,t} = \beta\,\frac{c_t}{c_{t+1}}\Big[(1-\delta_h-\tau_h)p_{h,t+1} + \frac{u_{h,t+1}}{u_{c,t+1}}\Big].$$

Substitute $u_{h,t+1}/u_{c,t+1} = \Big(\alpha_h A(G_{t+1})/H_{t+1}\Big)\cdot c_{t+1}$:

$$p_{h,t} = \beta\,\frac{c_t}{c_{t+1}}\Big[(1-\delta_h-\tau_h)p_{h,t+1} + \alpha_h A(G_{t+1})\frac{c_{t+1}}{H_{t+1}}\Big]$$

which is equation (23).

A1.5 Road-Quality Vision Model: Fine-Tuning and Evaluation

In this section, we document our pipeline to classify road-surface quality from satellite imagery into three classes: poor (0), medium (1), and high (2). We begin from an ImageNet-pretrained vision model called *ConvNeXt V2*, available on Hugging Face and supported by Meta AI (Woo et al., 2023). We use the ConvNeXt V2 backbone and fine-tune the base version of the model on labeled road imagery from Brewer et al. (2021). The resulting model produces categorical ratings which can be converted into a continuous Road Quality Score (RQS), used in Section 1.5.1.

Fine-tuning Workflow. Figure 25 summarizes the end-to-end workflow. We begin with a corpus of 53,677 satellite tiles (224x224 pixels) labeled for road quality from Brewer et al. (2021). We feed these images into a ConvNeXt V2 model, replacing the original classification head with a 3-way softmax layer. We then fine-tune the entire model end-to-end using cross-entropy loss and the AdamW optimizer. The training process is monitored via a validation set, and we select the best-performing checkpoint based on minimum validation loss. Finally, we evaluate the fine-tuned model on a held-out test set to report accuracy, F1, and class-wise precision and recall.

Data and Training. We fine-tune ConvNeXt V2 trained on ImageNet-1k dataset. The classification head is replaced by a 3-way softmax. We fine-tune the model for 3 epochs and use cross-entropy loss, AdamW optimizer with learning rate of 5×10^{-5} and weight decay of 0.01, and a batch size of 16 for training and 32 for evaluation. Inputs use the model's

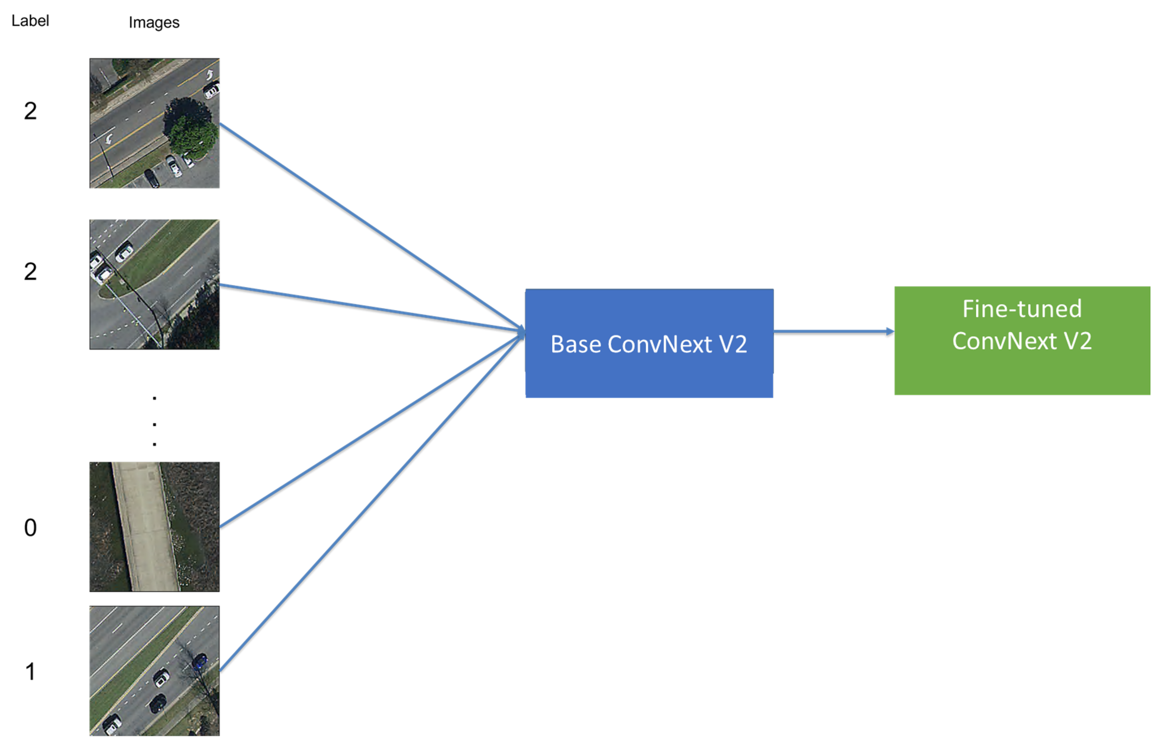

Figure 25. Fine-tuning workflow for road-quality classification

ImageNet normalization; no explicit augmentation. We use a weighted, stratified 70/15/15 split and select the checkpoint with the lowest validation loss.

Split	**Share**	**Stratified**	**Weights**	**Rationale**
Train	70%	By label	3/2/1	Mitigates class imbalance by allocating more minority-class samples to training, improving gradient signal and stability.
Validation	15%	By label	3/2/1	Class-aware monitoring consistent with train; used for model selection (min. val. loss).
Test	15%	By label	3/2/1	Hold-out evaluation with comparable class mix for fair assessment across splits.

Table 27. Split design and class-aware allocation. Weights are used only to allocate per-class counts into splits, *not* as loss weights.

Model Evaluation Metrics. We report accuracy, macro F1, and class-wise precision/recall to account for any residual imbalance. We also provide a confusion matrix to diagnose systematic confusions between adjacent quality classes.

Table 28. Overall accuracy by split (fine-tuned ConvNeXt V2).

Split	Overall Accuracy
Train	0.9115
Validation	0.8525
Test	0.8337

Table 29. Confusion matrix on the held-out Test set (N=1,539). Low/Medium/High correspond to classes 0/1/2.

		Predicted		
		Low	Medium	High
	Low	33	11	9
Actual	Medium	12	78	181
	High	2	41	1172

Table 30. Classification metrics on the Test set derived from Table 29.

Class	Precision	Recall	F1	Support
0	0.702	0.623	0.660	53
1	0.600	0.288	0.389	271
2	0.861	0.965	0.910	1,215
Macro avg	0.721	0.625	0.653	1,539
Weighted avg	0.810	0.834	0.810	1,539
Accuracy		0.8337		1,539

Table 28 presents model accuracy results on training, validation and test datasets. Accuracy is 0.9115 on train, 0.8525 on validation, and 0.8337 on test, indicating a good fit without overfitting on training data. As expected, accuracy declines slightly from train to validation to test. However, the gap between validation and test is less than 10 percentage points, suggesting good generalization across unseen data.

Table 29 shows the confusion matrix on the held-out test set and is useful to understand the types of errors the model makes when classifying road quality. As presented in the table, most errors are between adjacent classes: medium (1) is often predicted as high (2), and poor (0) occasionally as medium (1). Confusions between poor (0) and high (2) are rare, while high (2) is predominantly classified correctly.

Table 30 presents per-class metrics. Class 2 attains high precision and recall, driving weighted averages near overall accuracy, which suggests the model is very reliable at identifying high-quality roads. Class 1 shows low recall, lowering macro F1 and reflecting difficulty on borderline surfaces. This means the model often misses medium-quality roads, misclassifying them as high quality. Class 0 performs moderately, with balanced precision/recall relative to its smaller support, indicating reasonable reliability on poor-quality roads despite fewer examples. Overall, the model excels at identifying high-quality roads but nevertheless struggles more with medium and poor classes, especially in recall for medium quality.

Base vs. Fine-Tuned ConvNeXt V2. Table 31 summarizes the main gains from end-to-end adaptation. We compare our fine-tuned ConvNeXt V2 against a baseline where the pretrained backbone is frozen and only the classification head is trained. This linear-probe baseline on frozen ImageNet features approach is common in transfer learning.

Table 31. Fine-tuned vs. baseline ConvNeXt V2.

Model	Overall Accuracy
Fine-tuned ConvNeXt V2	0.8337
Baseline ConvNeXt V2	0.2586

Fine-tuning improves test accuracy by 57.51 (from 0.2586 to 0.8337) percentage points, a more than 220% relative gain. This indicates that adapting the pretrained backbone to road imagery is crucial for learning task-specific features and achieving reliable generalization, justifying end-to-end fine-tuning over a frozen baseline.

A1.6 MVPF Calculation

As stated in Section 1.7, we use the following formula to calculate the Marginal Value of Public Funds (MVPF):

$$\text{MVPF} = \frac{\text{WTP}}{\text{Net Cost to Government}} = \frac{\text{Tax Savings - House Value Decline}}{\text{Direct Long Run Cost + Fiscal Externality}} \quad (47)$$

We combine our estimated treatment effects on house prices with expected loss in tax revenue to calculate the MVPF of road maintenance tax levies. The expected loss in tax revenue, as calculated in Table 1, is $76 per household per year. Using a capitalization rate of 5% (Hilber, 2011; Cellini, Ferreira and Rothstein, 2010), the lifetime present value of this annual tax saving is $1,520 per household, which can also be interpreted as the tax savings from cutting the renewal road tax levy. Using the standard property tax rate of 1.5% in Ohio, and the estimated decline in house prices from the tax cut in Table 5, we compute the 10-year tax base decline of $1,911, which represents the fiscal externality from cutting the road tax levy. The net cost to the government is then calculated as the direct long-run cost of $1,520 plus the fiscal externality of $1,911, totaling $3,431. The willingness to pay (WTP) is calculated as the tax savings of $1,520 minus the yearly capitalization of house price estimates over 10 years, which amounts to $6,370 per household[63], giving a WTP of $-\$4,850$. Finally, we compute the MVPF as $-\$4,850/\$3,431 = -\$1.41$.

References

Brewer, Ethan, Jason Lin, Peter Kemper, John Hennin, and Dan Runfola. 2021. "Predicting road quality using high resolution satellite imagery: A transfer learning

[63] This calculation is based on the idea that the house price estimate in year t reflects future disutility from decline in public amenity.

approach." *Plos one*, 16(7): e0253370.

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

Chatterjee, Santanu, John Gibson, and Felix Rioja. 2017. "Optimal Public Debt Redux." *Journal of Economic Dynamics and Control*, 83: 162–174. Originally circulated as ICePP Working Paper No. 16-13 at Georgia State University.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review*, 101(3): 471–475.

Cooley, Thomas F. 1995. *Frontiers of business cycle research.* Princeton University Press.

Hilber, Christian AL. 2011. "The economic implications of house price capitalization: A survey of an emerging literature." *Spatial Economics Research Centre Discussion Paper*. Discussion Paper 91.

King, Robert G, Charles I Plosser, and Sergio T Rebelo. 1988. "Production, growth and business cycles: I. The basic neoclassical model." *Journal of monetary Economics*, 21(2-3): 195–232.

Rioja, Felix K. 2003. "Filling potholes: Macroeconomic effects of Maintenance versus New Investments in public infrastructure." *Journal of Public Economics*, 87(9-10): 2281–2304.

U.S. Bureau of Labor Statistics. 2024. "Consumer Expenditures in 2023." U.S. Bureau of Labor Statistics BLS Report. Published on BLS website.

Woo, Sanghyun, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. 2023. "Convnext v2: Co-designing and scaling convnets with masked autoencoders." 16133–16142.

Appendix

Appendix for Chapter 2

A2.1 Permit Timing Analysis

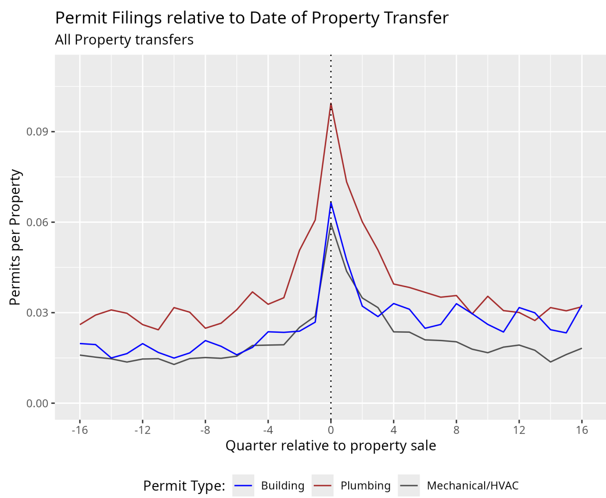

Panel A: Minneapolis

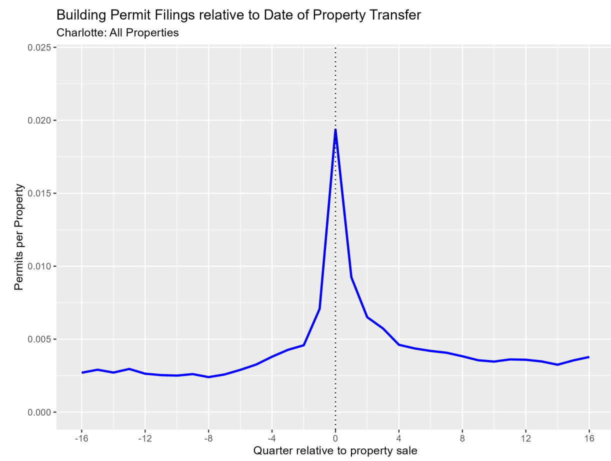

Panel B: Charlotte

Figure 26. Permit Issuance Rates Surrounding Property Transfers

We study the timing of the permit around the sale date to understand the reinvestment behavior of property buyers and sellers as they are about to either sell or purchase a property. Panel A of Figure 26 presents results for Minneapolis for this permit timing analysis. Permits appear to largely reflect reinvestment performed by incumbent property owners: More than 85% of all permits occur more than a year away from any property transfer. Nonetheless, as shown in Figure 26, permit issuance rates increase surrounding property transfers, particularly in the quarters immediately following a property changing hands. Panel B of Figure 26 displays permit issuance behavior surrounding property transfers in Charlotte, North Carolina. Despite Charlotte's less stringent permit requirements and data availability limitations, the observed pattern is qualitatively similar to that of Minneapolis.

A2.2 Parcel Age and Permits

Figure 27 illustrates how the average number of permits filed by a single-family home (SFH) in Charlotte and Minneapolis vary with property age. We group building ages into deciles from 1 to 120 years. As properties mature, they typically experience more wear and tear, which corresponds to an increase in reinvestment activity. Indeed, the figure shows a rising trend in permits for older deciles, followed by a slight dip at the later deciles (around 100 years). Although the median SFH in Minneapolis is much older than median SFH in Charlotte, we see similar positive relationship between permit count and age in Minneapolis. While this relationship is not strictly monotonic, it highlights the expectation that older homes often require more upkeep, consistent with findings in the broader housing literature (Davidoff, 2004).

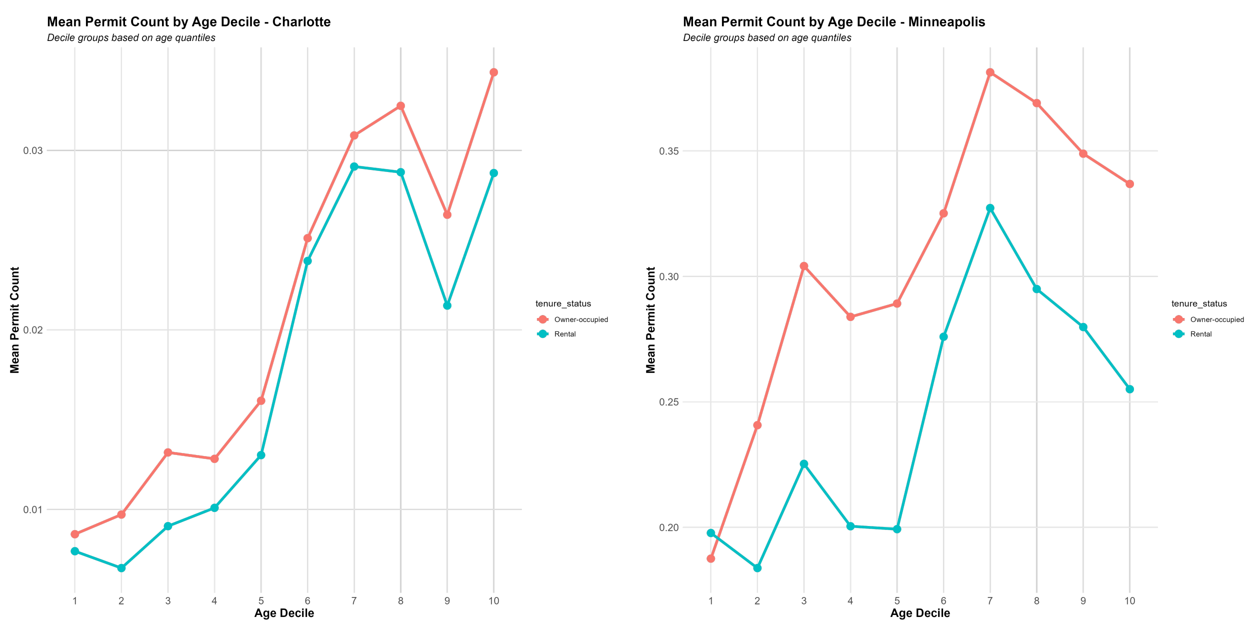

Figure 27. Mean Permit Count by Age and Tenure

A2.3 Supplementary Tables

Table 32. Home Reinvestment by Tenure Status and Landlord Scale: Minneapolis (Sale-Year Effects)

	(1)	(2)
Panel A: Mechanical Permits		
Owner-Occupied	0.127***	0.177***
	(0.015)	(0.017)
Large Landlord (20+)	−0.382***	−0.388***
	(0.054)	(0.057)
Sale-Year indicator	1.034***	1.270***
	(0.017)	(0.034)
Sale Year × Owner-Occupied		−0.310***
		(0.039)
Sale Year × Large Landlord (20+)		0.589***
		(0.182)
Panel B: Plumbing Permits		
Owner-Occupied	0.052***	0.076***
	(0.011)	(0.012)
Large Landlord (20+)	−0.089***	−0.102***
	(0.034)	(0.035)
Sale-Year indicator	0.983***	1.119***
	(0.013)	(0.026)
Sale Year × Owner-Occupied		−0.199***
		(0.030)
Sale Year × Large Landlord (20+)		0.436***
		(0.132)
Panel C: Building Permits		
Owner-Occupied	0.263***	0.314***
	(0.012)	(0.013)
Large Landlord (20+)	−0.452***	−0.463***
	(0.043)	(0.045)
Sale-Year indicator	0.754***	1.070***
	(0.015)	(0.031)
Sale Year × Owner-Occupied		−0.444***
		(0.035)
Sale Year × Large Landlord (20+)		0.713***
		(0.160)
Panel D: Building Permit Value		
Owner-Occupied	0.507***	0.554***
	(0.024)	(0.025)
Large Landlord (20+)	−0.484***	−0.484***
	(0.064)	(0.065)
Sale-Year indicator	2.019***	2.580***
	(0.042)	(0.082)
Sale Year × Owner-Occupied		−0.854***
		(0.096)
Sale Year × Large Landlord (20+)		1.304***
		(0.459)

Table 33. Home Reinvestment by Large-Landlord Size – Minneapolis

	Dependent variable: Permit Activity			
	Plumbing Counts	*Mechanical Counts*	*Building Permit Counts*	*Building Permit Values*
	(1)	(2)	(3)	(4)
Large Landlord, 10+ properties	−0.158***	−0.382***	−0.361***	−0.443***
	(0.029)	(0.045)	(0.035)	(0.056)
Large Landlord, 20+ properties	−0.172***	−0.471***	−0.503***	−0.595***
	(0.034)	(0.054)	(0.043)	(0.064)
Large Landlord, 50+ properties	−0.196***	−0.808***	−0.781***	−0.857***
	(0.042)	(0.078)	(0.060)	(0.079)
Large Landlord, 100+ properties	−0.250***	−0.827***	−0.872***	−0.914***
	(0.047)	(0.086)	(0.068)	(0.085)
Observations	597,248	597,248	597,248	597,248
Housing Characteristics	✓	✓	✓	✓
Year F.E.	✓	✓	✓	✓
Neighborhood F.E.	✓	✓	✓	✓

Notes: Each row reports estimates from separate regressions that vary the definition of a large landlord using the ownership thresholds shown (10+, 20+, 50+, or 100+ properties). Columns (1)–(3) estimate negative-binomial models for permit counts; column (4) reports OLS estimates with the dependent variable as natural log of permit value. All specifications include housing characteristics (bedrooms, bathrooms, heated area, lot size, age, and age squared) as well as year and neighborhood fixed effects. Standard errors are in parentheses. Significance levels: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 34. Home Reinvestment by Large-Landlord Size – Charlotte

	Building Permits (Count)		*Building Permits (Dollars)*
	Neg. binomial	*Poisson*	*OLS*
	(1)	(2)	(3)
Large Landlord, 10+ properties	-0.353^{***}	-0.472^{***}	-0.067^{***}
	(0.026)	(0.025)	(0.005)
Large Landlord, 20+ properties	-0.530^{***}	-0.664^{***}	-0.082^{***}
	(0.033)	(0.032)	(0.006)
Large Landlord, 50+ properties	-0.746^{***}	-0.889^{***}	-0.095^{***}
	(0.043)	(0.042)	(0.007)
Large Landlord, 100+ properties	-0.840^{***}	-0.982^{***}	-0.099^{***}
	(0.048)	(0.047)	(0.007)
Observations	4,114,816	4,114,816	4,114,816
Housing Characteristics	✓	✓	✓
Year F.E.	✓	✓	✓
Municipality F.E.	✓	✓	✓

Notes: Each row reports estimates from separate regressions that vary the definition of a large landlord using the ownership thresholds shown (10+, 20+, 50+, or 100+ properties). Columns (1) and (2) report negative-binomial and Poisson regressions of building permit counts, respectively; column (3) reports OLS estimates with the dependent variable as natural log of permit value. All specifications include housing characteristics (bedrooms, bathrooms, heated area, lot size, age, and age squared) and year and municipality fixed effects. Standard errors are in parentheses. Significance levels: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 35. Home Reinvestment by Tenure Status and Landlord Scale: Different Fixed Effects (without Sale-Year indicator) – Minneapolis

	(1)	(2)
Panel A: Mechanical Permits		
Owner-Occupied	0.069***	0.081***
	(0.015)	(0.015)
Large Landlord (20+)	−0.471***	−0.481***
	(0.054)	(0.054)
Panel B: Plumbing Permits		
Owner-Occupied	−0.002	0.012
	(0.011)	(0.011)
Large Landlord (20+)	−0.172***	−0.178***
	(0.034)	(0.034)
Panel C: Building Permits		
Owner-Occupied	0.227***	0.242***
	(0.012)	(0.012)
Large Landlord (20+)	−0.503***	−0.512***
	(0.043)	(0.043)
Panel D: Building Permit Value		
Owner-Occupied	0.437***	0.465***
	(0.024)	(0.024)
Large Landlord (20+)	−0.595***	−0.610***
	(0.064)	(0.064)
Neighborhood fixed effects	✓	
Community fixed effects		✓
Year fixed effects	✓	✓
Housing characteristics	✓	✓
Observations	597,248	597,248

Notes: Results are from parcel-year level regressions of permit outcomes on an owner-occupied indicator and a large-landlord indicator, with small-landlord as the reference group. Panels A–C report negative binomial estimates for counts; Panel D reports OLS estimates for log permit value. Standard errors are in parentheses.

References

Davidoff, Thomas. 2004. "Maintenance and the Home Equity of the Elderly."

Appendix

Appendix for Chapter 3

Location of ULBs

We map the settlements with Urban Local Bodies (ULBs) in India and present them below in Figure 28.

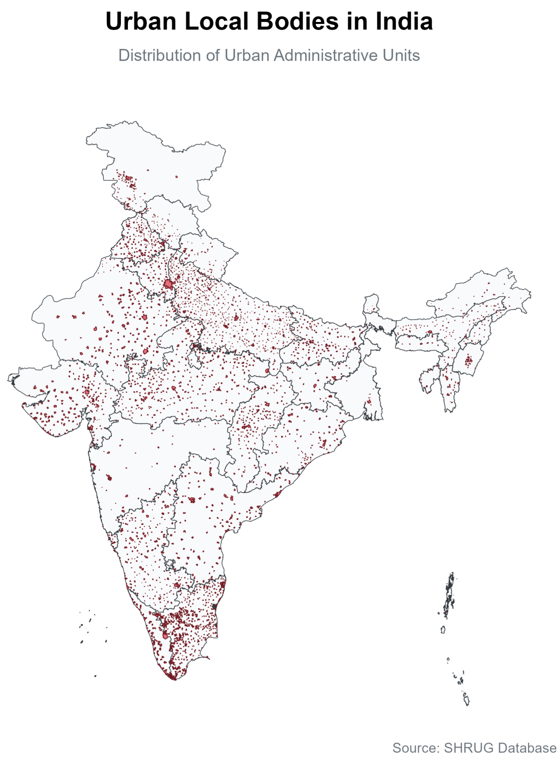

Figure 28. Urban Local Bodies (ULBs) in India

Reduced-Form Estimates

Table 36 reports the reduced-form effect of crossing the Census Town eligibility frontier on local public goods provision in the local sample. These estimates provide the numerator of the Wald ratio underlying the main IV results in the text. The sign pattern matches the IV estimates throughout: education, health, financial access, and selected community amenities improve, while sports infrastructure declines.

Table 36. Reduced-Form Effects of Census Town Eligibility on Public Goods Provision

Outcome	Estimate	SE	N
Education			
Primary schools	0.982***	(0.212)	37,143
Middle schools	0.572***	(0.121)	37,099
Secondary schools	0.346***	(0.066)	37,148
Health			
All hospitals	0.179***	(0.036)	37,091
Family welfare centers	0.212***	(0.041)	37,077
Financial access			
Cooperative banks	0.297***	(0.044)	37,089
Agricultural credit societies	0.202***	(0.061)	37,093
Community infrastructure			
Public libraries	0.074**	(0.030)	37,147
Public reading rooms	0.102***	(0.035)	37,147
Cinema halls	0.056*	(0.030)	37,146
Sports infrastructure	−0.405***	(0.057)	37,146
Controls: Yes	District FE: Yes	SEs clustered by district	

Notes: Reduced-form estimates from OLS regressions of each outcome on Census Town eligibility in 2001 (a binary indicator for meeting all three CT thresholds). The local sample restricts to settlements near the 2001 census thresholds (population ±5,000; density ±400; non-agricultural male share ±0.20). All specifications include controls for log population, log density, non-agricultural male workforce share, literacy rate, main worker share, and caste composition (SC and ST shares), as well as district fixed effects. Robust standard errors clustered at the district level are in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

In magnitude, the reduced-form estimates are smaller because only a subset of CT-eligible settlements is ultimately statutorily recognized. Even so, the effects remain consistent with IV estimates: crossing the eligibility frontier raises the number of primary schools, middle schools, and secondary schools, while also increasing hospitals, family welfare centers, cooperative banks, and agricultural credit societies. The decline in sports infrastructure is also visible in reduced form, consistent with the land-reallocation interpretation discussed in the main text.

Balance Checks

To assess the validity of our regression discontinuity design, we examine whether settlements just below and just above the Census Town thresholds are comparable on observable characteristics measured in 2001. Table 37 presents balance tests for key demographic and economic variables across both the full sample (Panel A) and the close-to-threshold sample (Panel B). The variables include log population, log density, share of non-agricultural male workers, literacy rate, main worker rate, and shares of Scheduled Caste and Scheduled Tribe populations.

As expected in the full sample (Panel A), settlements meeting CT criteria differ substantially from those that do not as they are considerably larger, denser, more economically non-agricultural, and have higher literacy rates. However, when we restrict attention to the close-to-threshold sample (Panel B), these differences narrow substantially. While some gaps remain, particularly in population and density, which are among the threshold criteria themselves, the key observation is that settlements just above and below the thresholds are much more similar than the full sample comparison would suggest. Neverthless, we control for these remaining differences in our regression models by adding these variables along with appropriate interaction terms and fixed effects.

Bandwidth Sensitivity

Table 38 presents bandwidth sensitivity analysis for the effect on government primary schools, the outcome with the largest absolute magnitude. We vary the bandwidth multiplier from 0.5× to 2.0× the baseline thresholds and report both the reduced-form (RF) and instrumental variable (IV) estimates at each bandwidth. At narrow bandwidths (0.5×, 0.75×), the first stage is too weak for reliable IV estimation as we lack number of treated observations, and

Table 37. Balance Checks: 2001 variables by CT Status

Variable	**Non-CT**	**CT**	**N**
Panel A: Full Sample			
Log Population	6.55	9.84	502,179
	(1.23)	(0.87)	
Log Density	5.69	7.77	502,179
	(1.45)	(1.12)	
Non-Ag Male Workers	0.26	0.90	502,179
	(0.24)	(0.15)	
Literacy Rate	0.47	0.67	502,179
	(0.18)	(0.12)	
Main Worker Rate	0.32	0.28	502,179
	(0.12)	(0.08)	
SC Share	0.17	0.13	502,179
	(0.21)	(0.15)	
ST Share	0.16	0.04	502,179
	(0.31)	(0.15)	
Panel B: Close-to-Threshold Sample			
Log Population	6.42	8.85	37,151
	(1.18)	(0.92)	
Log Density	5.60	6.37	37,151
	(1.38)	(1.25)	
Non-Ag Male Workers	0.71	0.83	37,151
	(0.18)	(0.14)	
Literacy Rate	0.55	0.63	37,151
	(0.16)	(0.13)	
Main Worker Rate	0.27	0.32	37,151
	(0.11)	(0.09)	
SC Share	0.19	0.14	37,151
	(0.22)	(0.17)	
ST Share	0.12	0.05	37,151
	(0.27)	(0.18)	

Notes: This table presents mean characteristics in 2001 for settlements that did not meet Census Town (CT) criteria versus those that did. Standard deviations are reported in parentheses below means. Panel A shows the full sample; Panel B restricts to settlements close to the CT thresholds (within specified bandwidth). All variables are measured in 2001, prior to potential statutory recognition by 2011.

the reduced-form estimates are imprecise. However, starting from the baseline bandwidth (1.0×) and above, the first stage becomes strong enough for valid IV estimation, and the reduced-form estimates stabilize around an increase of approximately 1 additional primary school due to crossing the CT eligibility frontier. The IV estimates at these bandwidths are also positive and statistically significant, though they vary in magnitude due to changes in the strength of the first stage and sample composition. Overall, the results suggest that our main findings are robust to reasonable variations in bandwidth choice, while very narrow bandwidths may lack sufficient power for precise estimation.

Table 38. Bandwidth Sensitivity: Primary School Estimates

Bandwidth	N	N(ST=1)	FS coef	RF est	RF SE	IV est	IV SE
0.5×	1,216	17	0.032	−0.190	(0.226)	−5.92	(8.38)
0.75×	6,434	57	0.013	0.041	(0.353)	3.11	(25.90)
1.0× (Baseline)	**37,143**	**127**	**0.071**	**0.982*****	**(0.212)**	**13.86*****	**(4.00)**
1.25×	51,198	234	0.076	0.899***	(0.145)	11.83***	(2.38)
1.5×	74,214	325	0.085	1.009***	(0.126)	11.84***	(2.06)
2.0×	112,856	535	0.110	1.027***	(0.124)	9.30***	(1.36)

Notes: Each row reports estimates from a different bandwidth specification. Bandwidths are multiples of the baseline thresholds: population $\pm k \times 5{,}000$, density $\pm k \times 400$, non-agricultural male share $\pm k \times 0.20$. "FS coef" is the first-stage coefficient on CT eligibility. "RF" is the reduced-form effect of CT eligibility on primary schools. "IV" is the 2SLS estimate of statutory recognition on primary schools. All specifications include the full set of controls and district fixed effects with district-clustered standard errors. At narrow bandwidths (0.5×, 0.75×), the first stage is too weak for reliable IV estimation. The reduced-form estimate is stable at approximately 1 additional primary school across all bandwidths with sufficient power (≥ 1.0×). $^{***}p < 0.01$.

Descriptive Outcome Levels

Table 39 reports unadjusted outcome means by statutory status in the full sample and in the close-to-threshold local sample. These comparisons are purely descriptive and are not part of the identification argument. Their value is diagnostic: they show how large the raw full-sample differences are, and how much those differences compress once we restrict attention to settlements near the Census Town thresholds used in the main analysis.

The descriptive comparison reveals a sharp contrast between the two samples. In the full sample, statutory settlements have far higher levels of nearly every public good: for example, they have 15.50 primary schools versus 1.27 in non-statutory settlements, 1.90

Table 39. Descriptive Statistics: Outcome Variables by Statutory Status in the Global and Local Samples

	Global sample			Local sample		
Outcome	Non-ST	ST	Difference	Non-ST	ST	Difference
Education						
Primary schools	1.27	15.50	14.23	1.24	4.96	3.72
Middle schools	0.55	7.64	7.09	0.58	2.65	2.07
Secondary schools	0.19	4.30	4.12	0.25	1.64	1.39
Health						
All hospitals	0.01	1.13	1.12	0.02	0.41	0.39
Family welfare centers	0.06	1.18	1.12	0.08	0.53	0.46
Financial access						
Cooperative banks	0.06	1.90	1.83	0.09	0.97	0.88
Agricultural credit societies	0.16	2.21	2.04	0.17	0.92	0.75
Community infrastructure						
Public libraries	0.13	1.02	0.89	0.11	0.43	0.31
Public reading rooms	0.15	1.03	0.88	0.13	0.45	0.32
Cinema halls	0.06	1.73	1.67	0.05	0.13	0.08
Sports infrastructure	0.42	1.06	0.64	0.46	0.49	0.03

Notes: Entries are unadjusted means of each outcome by statutory town status. "Difference" column is the raw difference in means between statutory and non-statutory settlements. The global sample uses the full regression sample with non-missing covariates and outcome data. The local sample uses the close-to-threshold sample from the main analysis. Outcome-specific sample sizes vary slightly because of missingness and match the corresponding samples used in Table 40. These differences are descriptive only.

cooperative banks versus 0.06, and 1.73 cinema halls versus 0.06. Once we move to the local sample, these gaps narrow substantially. The compression is especially stark for community amenities: the raw gap in cinema halls falls from 1.67 to 0.08, and the raw gap in sports infrastructure shrinks from 0.64 to 0.03. Note that these raw differences are not adjusted for any covariates or fixed effects, and thus should not be interpreted as causal estimates of local urban governance.

OLS versus IV Estimates

Table 40 takes the comparison one step further by moving from raw descriptive differences to adjusted regression estimates. Column (1) reports global OLS estimates from the full sample of 502,110 settlements. Column (2) restricts the OLS to the close-to-threshold local sample used in the main analysis. Column (3) reports the local 2SLS estimates from Tables 16–19, where statutory status is instrumented by Census Town eligibility. Reading alongside Table 39, the progression is straightforward: global raw differences are large, local raw differences are smaller, local OLS attenuates further after conditioning on observables, and the local IV estimates provide the paper's causal parameter.

Taken together, Tables 39 and 40 reveal three instructive patterns. First, both the raw means and the global OLS estimates in Column (1) show very large positive differences between statutory and non-statutory settlements for nearly all outcomes. Those differences are informative about broad urban-rural contrasts, but they are heavily contaminated by selection: larger, denser, and more economically developed settlements are both more likely to be statutory towns and more likely to have schools, hospitals, and banks, regardless of their governance status. The global comparison therefore conflates the effect of statutory recognition with pre-existing differences across places.

Second, restricting the comparison to the close-to-threshold local sample sharply compresses the differences. This is visible in the raw means above and again in the adjusted local OLS estimates in Column (2). For education outcomes, the OLS coefficient on primary schools drops from 11.20 to 2.08, and for middle schools from 5.22 to 1.10. For community infrastructure, the local OLS estimates are essentially zero — public libraries (-0.01), reading rooms (-0.08), and cinema halls (-0.01) show no conditional association with statutory status. This attenuation is exactly what we would expect if settlements near the thresholds are much more comparable than the full sample.

Third, Column (3) shows that the local IV estimates that are interpreted as our paper's causal estimates for settlements whose statutory recognition status was shifted by crossing the Census Town thresholds. These 2SLS coefficients are uniformly larger than the local OLS estimates, and for several outcomes they are larger than even the global OLS estimates. For example, the IV estimate for family welfare centers (3.00) is more than 3.5 times the global OLS (0.85), and the IV estimate for public libraries (1.05) is significant where the local OLS is essentially zero. This pattern is consistent with a LATE interpretation: statutory recognition has especially large effects for marginal settlements near the eligibility frontier, while naive cross-sectional comparisons either mix together incomparable places or understate the effect because of attenuation in the observed status measure.

Table 40. OLS versus IV Estimates of the Effect of Statutory Recognition

	(1) OLS Global	(2) OLS Local	(3) 2SLS Local
Education			
Primary schools	11.20***	2.08***	13.86***
	(0.95)	(0.49)	(4.00)
Middle schools	5.22***	1.10***	7.72***
	(0.42)	(0.22)	(2.25)
Secondary schools	3.14***	0.92***	4.89***
	(0.24)	(0.10)	(1.30)
Health			
All hospitals	1.00***	0.36***	2.53***
	(0.06)	(0.06)	(0.69)
Family welfare centers	0.85***	0.33***	3.00***
	(0.06)	(0.05)	(0.88)
Financial access			
Cooperative banks	1.46***	0.64***	4.09***
	(0.10)	(0.07)	(0.97)
Agricultural credit societies	1.47***	0.47***	2.84***
	(0.26)	(0.09)	(1.01)
Community infrastructure			
Public libraries	0.51***	−0.01	1.05**
	(0.06)	(0.06)	(0.50)
Public reading rooms	0.49***	−0.08	1.44**
	(0.06)	(0.07)	(0.63)
Cinema halls	1.49***	−0.01	0.79
	(0.11)	(0.04)	(0.49)
Sports infrastructure	−0.09	−0.46***	−5.71***
	(0.09)	(0.08)	(1.48)
Controls	Yes	Yes	Yes
District FE	Yes	Yes	Yes
Observations	502,110	37,148	37,148

Notes: Columns (1) and (2) report OLS estimates of the coefficient on statutory town status (ST_i). Column (3) reports 2SLS estimates where ST_i is instrumented by Census Town eligibility in 2001 (Z_i). All specifications include controls for log population, log density, non-agricultural male workforce share, literacy rate, main worker rate, and caste shares, as well as district fixed effects. Robust standard errors clustered at the district level are in parentheses. The "Global" sample includes all settlements with non-missing covariates. The "Local" sample restricts to settlements within ±5,000 of the population threshold, ±400 of the density threshold, and ±0.20 of the non-agricultural share threshold.
Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.